Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Seasonality in transition: Renewables, electricity price shapes and seasonal priors

Viviana Fanelli

viviana.fanelli@unifg.it

University of Foggia

Research Article

Keywords: Electricity prices, Seasonality, Forward curves, Renewable energy

Posted Date: November 11th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11651089/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Seasonality in transition: Renewables, electricity price shapes and seasonal priors

Viviana Fanelli[a,*]

[a] *Dipartiment of Economics, University of Foggia, Foggia, Italy*

Abstract

Seasonal price shapes underpin electricity forward curves, solar power purchase agreement valuation and flexible-asset scheduling, yet historical priors may become stale as renewable penetration rises. We develop transition-dependent seasonality, linking the Fourier representation of electricity price shapes to solar and wind penetration, and derive a decision-based criterion for updating a prior in linear optimisation. European evidence documents a broad decline in summer prices relative to winter, concentrated at midday. The solar channel is identifiable within the day, whereas the model's common-effect restriction fails for wind. Out of sample, transition-dependent seasonality reduces systematic overvaluation and average absolute valuation errors for solar power purchase agreements relative to historical shapes, although its benefits depend on penetration forecasts and do not consistently extend to longer horizons or futures-calibrated curves. Gains for monthly baseload shapes are not significant. For flexible assets, updated priors often leave optimal schedules unchanged, and average operating gains remain statistically insignificant. These results identify where transition-adjusted seasonality matters for valuation and where it adds less operational value.

Keywords: Electricity prices, Seasonality, Forward curves, Renewable energy

1. Introduction

Many operational and valuation decisions in power markets rest on a seasonal shape: the profile that distributes an annual price across months and, in hourly price forward curves, across the hours of the day. The forward-curve constructions of Fleten and Lemming [1], Benth et al. [2], Caldana et al. [3] and Kiesel et al. [4] impose seasonal shapes on traded annual and quarterly contracts to infer monthly delivery prices beyond the front months. Hirth [5] and Qorbanian et al. [6] show why the timing of production across cheap and expensive months and hours matters for the value of pay-as-produced power purchase agreements (PPAs). Löhndorf et al. [7] and Löhndorf and Wozabal [8] study scheduling and valuation of energy-limited assets, including reservoir hydro plants and seasonal stores, against expected seasonal price profiles. In practice, these shapes are priors built from historical averages, calendar dummies or sinusoids, as in Lucia and Schwartz [9] and Frau and Fanelli [10].

The energy transition calls this practice into question. Solar output peaks in summer and around noon, wind output in winter, so the seasonal profile of residual load, and with it that

*Corresponding author.
Email address: viviana.fanelli@unifg.it (Viviana Fanelli)

of prices, changes as their shares grow, and it already has: between 2017-2019 and 2023-2025 the ratio of summer to winter prices fell in all 24 European markets with three years of data in 2017-2019, from a median of 0.93 to 0.82; in Spain the April and May shape factors fell from 0.88 and 0.96 to 0.49 and 0.53 of the annual average. The change sits in the middle of the day: across markets, the median midday price of the summer months fell from 86-111% to 39-56% of the annual baseload price, whereas the seasonal shapes of night and evening prices did not fall. Priors built from history are therefore stale. Yet replacing a prior has risks of its own, and not every decision is sensitive to it. When does updating a seasonal prior for the energy transition change a decision, and when is a historical shape good enough?

We answer this question for three decisions: the valuation of solar PPAs with hourly price forward curves, the scheduling of energy-limited assets and the valuation of seasonal flexibility. The answer depends on the decision. The update matters most for contracts with fixed hourly profiles: history-based priors overvalued solar output by 22% of the baseload price in 2023-2025, and a transition-adjusted prior cut the overvaluation to 7.5%. It matters less for schedules, which stayed unchanged under 36-45% of the corrections, and it cannot close the gap between attainable and perfect-foresight values of flexibility. The paper makes four contributions.

First, a transition-dependent seasonality (TDS) model. A first-order expansion of the merit order around annual residual load makes the seasonal component of log prices affine in the annual penetration of each technology, with a loading equal to the technology's relative output profile; in Fourier form, it is the sinusoidal seasonal function of Lucia and Schwartz [9] and Frau and Fanelli [10] with harmonic coefficients that move with penetration. Applied hour by hour, it implies that solar penetration reshapes the seasonal profile of midday prices but not of night prices, and deepens the discount of midday prices relative to baseload. It also implies why historical shapes are biased and how seasonal value factors respond to penetration (Propositions 1-2 and the Online Supplement).

Second, criteria for when a prior matters. Following the linear-programming sensitivity framework of Wendell [11] and the predict-then-optimise approach of Elmachtoub and Grigas [12], Proposition 3 collects results for decisions that are linear programmes in prices: prior errors pass through to valuations, exactly for fixed profiles and to first order for flexible assets, but leave the schedule unchanged within a tolerance region, and intrinsic values fall short of expected perfect-foresight values by the value of perfect information. Proposition 4 applies them to the choice between an incumbent and a candidate prior: the realised gain from updating depends only on forecast errors along the change in the schedule, and updating pays in expectation if and only if the conditional mean of the shape endorses a share of the update above a threshold known at the decision date; the endorsed share itself must be estimated from past decisions.

Third, evidence on the drivers of the change. The decline of summer relative to winter prices holds net of the within-year path of fuel costs, and seventeen other commodities show no more significant changes than chance would produce. At monthly baseload resolution, wind, hydro and nuclear output and demand move prices as the model predicts, also net of common European shocks, but the solar effect rests on common variation; within the day, solar is identified from market-specific variation, a solar deviation of 1% of monthly demand lowering the midday price relative to the night price by 2.1% of the annual baseload price. The model's central restriction,

that penetration-driven and weather-driven output changes move prices alike, is not rejected for solar and fails for wind.

Fourth, an out-of-sample evaluation of shaping rules and decisions. History-based hourly forward curves overvalued one-year solar PPAs by 11.2 points of the annual baseload price on average over 2018-2025 and by 22.2 points in 2023-2025. The TDS adjustment, computed ex ante, cut the mean overvaluation to 2.8 points and the mean absolute error from 13.0 to 7.7 points, with $p = 0.09$, outperforming trend extrapolation with $p = 0.03$, mainly by deepening the midday discount, and that of three-year contracts from 12.6 to 6.9 points, with $p = 0.06$. In EUR/MWh the gain is not significant, because the adjustment underpriced solar output in 2022, when prices were highest, and at long horizons it is only as good as the penetration forecast. On curves calibrated to traded futures in three markets, it still lowered valuations, by 7 to 8 points in each, but its accuracy gain was insignificant. For monthly baseload shapes it brings no significant accuracy gain, nor do shapes implied by traded futures; for midday shapes it removes the mean summer bias. For a stylised reservoir and seasonal store, replacing a smoothed historical prior by its transition-adjusted version leaves the schedule unchanged in 35.5% and 44.5% of market-years; changed schedules gain in nearly two thirds of cases, specifically 64% for the reservoir and 65% for the store, and more when the value claimed ex ante is high, but the mean gain is not significant at 5%, with $p = 0.14$ for the reservoir and $p = 0.09$ for the store, and the estimated endorsed share, 1.7-2.1, is significantly above one only for the store. Rolling re-optimisation with month-ahead prices did not systematically help, and the perfect-foresight value of flexibility, which rose by 70-160%, is an upper bound far above attainable values.

For practice, the results suggest building the hourly forward curves used for solar PPAs by imposing on traded futures an updated within-month shape, transition-adjusted or at least recent, with the penetration forecast as a key input; relying on smoothing rather than structural adjustment for monthly baseload priors; re-optimising schedules on the adjusted prior; and treating perfect-foresight valuations of flexibility as upper bounds. Section 2 reviews the relevant literature. Section 3 develops the TDS framework, its valuation and operational implications, and the empirical specification. Section 4 presents the data and documents changes in electricity-price seasonality. Section 5 examines their drivers, cross-market effects and model restrictions. Section 6 evaluates seasonal priors out of sample across price blocks and forecast horizons. Section 7 assesses their implications for solar PPA valuation, reservoir and storage scheduling, and the value of flexibility. Section 8 concludes. Proofs appear in the appendix, while the Online Supplement provides data checks, robustness analyses, additional results and futures benchmarks. All data are public except proprietary EEX futures prices.

2. Related literature

Moreno et al. [13] combine a Fourier seasonal component with stochastic long-run swings. Earlier deterministic seasonal functions in spot and futures models are developed by Lucia and Schwartz [9], Sørensen [14], Borovkova and Geman [15], Benth et al. [16] and Janczura et al. [17]; Nowotarski and Weron [18] explain their relevance for long-term forecasting. Seasonality also enters volatility, as discussed by Back et al. [19]. Fanelli et al. [20] model seasonal path-dependent volatility of electricity futures, Fanelli and Schmeck [21] find that implied volatilities depend on

the delivery month, Frau and Fanelli [10] build seasonal stochastic volatility from sinusoids, and Schneider and Tavin [22] extend this approach to agricultural markets. García Mirantes et al. [23] and Lisi and Pelagatti [24] allow seasonality to evolve stochastically. Our seasonal component is time-varying too, but driven by an observable state, renewable penetration, which makes it testable and usable under penetration scenarios. Bufalo and Fanelli [25] model the seasonal component of solar production itself, which is our loading.

Structural models map demand and supply conditions, including renewable infeed, into spot and forward prices through a bid stack, as in Barlow [26], Coulon and Howison [27] and Carmona et al. [28]. Kiesel and Kusterman [29] extend this logic to coupled markets, so seasonal state dependence is implicit in these approaches. The TDS model is a deliberately reduced counterpart at seasonal frequency: a first-order projection of bid-stack logic with closed-form implications and few parameters, estimable on public monthly data for many markets.

Fleten and Lemming [1], Benth et al. [2], Caldana et al. [3] and Kiesel et al. [4] construct smooth forward curves by combining traded prices with a seasonal prior; Leoni et al. [30] note that shaping coefficients change with solar and wind. The declining value factors of variable renewables as penetration increases are analysed by Lamont [31], Hirth [5] and López Prol et al. [32]. For corporate renewable procurement, Qorbanian et al. [6] value PPAs with a fundamental price model, while Mohseni Taheri et al. [33] compare physical and virtual PPAs under demand and price uncertainty. Storage valuation combines intrinsic value under the forward curve with an extrinsic value from re-optimisation. Rolling-intrinsic policies are studied by Boogert and de Jong [34], Secomandi [35], Lai et al. [36], Wu et al. [37], Secomandi [38] and Löhndorf and Wozabal [8]; Löhndorf et al. [7] study hydro scheduling against seasonal price expectations. We contribute a seasonal prior conditioned on the generation mix, at monthly and hourly resolution, and show which valuations and schedules are sensitive to it.

Bertsimas and Kallus [39] and Elmachtoub and Grigas [12] emphasise decision regret rather than forecast error when forecasts enter optimisation problems. Wendell [11] characterises the region in which perturbations leave an optimal solution unchanged, and El Balghiti et al. [40] develops the corresponding margin conditions in predict-then-optimise. Smith and Winkler [41] describe the upward bias in optimised values when estimates are noisy, while Madansky [42] establishes the classical stochastic-programming bound between the expected perfect-foresight value and the value of the mean problem. Brown et al. [43] interpret that gap through an information relaxation. Propositions 3 and 4 collect these results for seasonal priors and apply them to the choice between two priors. The merit-order effects of wind and solar on price levels and variability analysed by Sensfuß et al. [44], Paraschiv et al. [45] and Wozabal et al. [46] concern hourly or daily price formation; our focus is the annual cycle and its variation across hours of the day.

3. Model

Consider market i, year y and month $m \in \{1, \dots, 12\}$, with monthly baseload price $P_{i,y,m}$ and annual average $\bar{P}_{i,y} = \frac{1}{12} \sum_m P_{i,y,m}$. The shape factor $s_{i,y,m} = P_{i,y,m}/\bar{P}_{i,y}$ distributes the annual price across months, and the log shape is $q_{i,y,m} = \ln P_{i,y,m} - \frac{1}{12} \sum_{m'} \ln P_{i,y,m'}$, with $s = 1 + q + O(q^2)$. Given the price F_Y of an annual contract and a shape forecast $\hat{s}_{Y,m}$, the

monthly forward curve is $F_{Y,m} = F_Y \hat{s}_{Y,m}$, so shape errors are forward-price errors in percent of the annual price. We normalise by the simple average of monthly prices, which differs from the hour-weighted average by 0.1% in the median market-year and by at most 0.6%. When quarterly contracts trade, the prior matters for the shape within each quarter. For a block b of hours of the day, $P^b_{i,y,m}$ is the mean price of those hours in month m, and the block shape factor $s^b_{i,y,m} = P^b_{i,y,m}/\bar{P}_{i,y}$ expresses it in units of the annual baseload price, the unit in which an hourly price forward curve shapes a baseload contract.

Let $D_{i,y,m}$ be demand and $Q^j_{i,y,m}$ generation from technology $j \in \mathcal{J} = \{S, W, H, N\}$ (solar, wind, hydro, nuclear). Annual penetration is $z^j_{i,y} = \sum_m Q^j_{i,y,m} / \sum_m D_{i,y,m}$. The relative seasonal profiles $\kappa^j_{i,y,m} = Q^j_{i,y,m}/\bar{Q}^j_{i,y} - 1$ and $\delta_{i,y,m} = D_{i,y,m}/\bar{D}_{i,y} - 1$ average zero within the year. The within-year deviation of technology j's output, in units of mean monthly demand, satisfies

$$x^j_{i,y,m} \equiv \frac{Q^j_{i,y,m} - \bar{Q}^j_{i,y}}{\bar{D}_{i,y}} = z^j_{i,y}\,\kappa^j_{i,y,m}, \tag{1}$$

so the seasonal footprint of a technology is its penetration times its seasonal profile.

Assumption 1 (Merit order)**.** Within market-year (i, y), $\ln P_{i,y,m} = \mu_{i,y} + g_i(\phi_{y,m}) + \Psi_i(R_{i,y,m}) + \varepsilon_{i,y,m}$, where $\phi_{y,m}$ collects fuel and carbon costs, g_i is a market-specific pass-through, Ψ_i is twice continuously differentiable with $\Psi_i' > 0$, $\varepsilon_{i,y,m}$ has a mean given the generation and demand profiles that depends only on market and month, $\bar{\varepsilon}_{i,m}$, and $R_{i,y,m} = (D_{i,y,m} - \sum_j \omega_j Q^j_{i,y,m})/\bar{D}_{i,y}$ is effective residual load in units of mean demand. The weight $\omega_j \in (0, 1]$ measures how far technology j displaces marginal supply one-for-one; $\omega_j < 1$ when part of its timing responds to prices, as for reservoir hydro or scheduled nuclear maintenance.

By (1), $R_{i,y,m} = \bar{R}_{i,y} + \delta_{i,y,m} - \sum_j \omega_j z^j_{i,y} \kappa^j_{i,y,m}$ with $\bar{R}_{i,y} = 1 - \sum_j \omega_j z^j_{i,y}$. We expand Ψ_i to first order around $\bar{R}_{i,y}$ and treat the slope $\theta_D = \Psi_i'(\bar{R}_{i,y})$ as locally constant and common across markets. Both are approximations: the slope changes as penetration lowers $\bar{R}$ and may differ across markets, and the second-order remainder depends on penetration. Section 5 tests whether slopes vary with penetration.

3.1. Transition-dependent seasonality

Proposition 1 (Transition-dependent seasonality)**.** Under Assumption 1, to first order in the seasonal deviations of residual load and with a locally constant slope,

$$q_{i,y,m} = \alpha_{i,m} + \theta_D\,\delta_{i,y,m} - \sum_{j \in \mathcal{J}} \theta_j\, z^j_{i,y}\, \kappa^j_{i,y,m} + \tilde{g}_{i,y,m} + \xi_{i,y,m}, \tag{2}$$

where $\theta_j = \omega_j \theta_D > 0$, $\tilde{g}_{i,y,m}$ is the within-year deviation of $g_i(\phi_{y,m})$, $\alpha_{i,m}$ collects the within-year deviation of $\bar{\varepsilon}_{i,m}$ and the stable market-specific part of the approximation remainder, and $\xi_{i,y,m}$ is the rest, with zero mean.

(a) Loadings: $\partial q_{i,y,m}/\partial z^j_{i,y} = -\theta_j \kappa^j_{i,y,m}$. Solar penetration lowers prices in months of above-average solar output and raises them elsewhere, relative to the annual mean; the same holds for wind in its windy months.

(b) Fourier form: write stable profiles as $\kappa^j_m = \sum_{k=1}^{6} (c^j_k)' f_k(m)$, where the vector $f_k(m)$ stacks $\cos(2\pi k\tau_m)$ and $\sin(2\pi k\tau_m)$ with $\tau_m = (m-\frac{1}{2})/12$ (for $k=6$ only the sine term is non-zero), and let C^0_k collect the coefficients of the remaining terms. Then the expected log shape is $\mathbb{E}[q_m] = \sum_k C_k(z)' f_k(m)$ with $C_k(z) = C^0_k - \sum_j \theta_j z^j c^j_k$: a sinusoidal seasonal function whose harmonic coefficients are affine in penetration. Constant penetration returns fixed sinusoidal seasonality.

(c) Amplitude: $\partial \|C_k(z)\|^2 / \partial z^j = -2\theta_j\, C_k(z)' c^j_k$. For non-zero $C_k(z)$ and c^j_k, penetration of technology j widens the k-th seasonal cycle if and only if the phases of the prevailing price harmonic and of j's output harmonic differ by more than a quarter period.

Part (c) explains why the same build-out can have different effects across markets. Where prices peak in winter, solar output runs against the price cycle and widens the winter-summer spread; where cooling demand makes summer prices high, solar first flattens the annual cycle. When interconnected markets share marginal supply, effective residual load also contains neighbours' output; we add neighbours' and rest-of-Europe deviations as empirical extensions of (2) rather than deriving them.

The expansion holds hour by hour. For hour h of month m it gives (2) for the log price of that hour relative to the annual mean, with loadings $\kappa^j_{i,y,m,h}$, technology j's mean output at that hour relative to its annual hourly mean; averaging over the hours of a block b gives block loadings $\kappa^j_{i,y,b,m}$, and we allow block-specific slopes θ^b_j in estimation. For solar, $\kappa^S_{m,h} \approx \nu_{m,h}(1+\kappa^S_m) - 1$, where $\nu_{m,h}$, output at hour h relative to the month's hourly mean, exceeds two around noon and is near zero at night (the relation is exact when κ^S_m is measured on mean hourly rather than monthly output). Two implications follow. Solar penetration reshapes the seasonal profile of midday prices and leaves that of night prices unchanged; the log contrast between the two blocks absorbs every shock that hits all hours of a month alike, such as fuel costs, the level of demand or hydro conditions, and the level contrast used in estimation absorbs them to first order. And because the midday loading is positive on average over the year, with an annual mean of about 1.5-2 in every market despite negative values in winter months in northern markets, solar penetration also lowers the annual level of midday prices relative to baseload, the midday discount that determines the capture price of solar output. Regressions with market×year effects identify only the seasonal part, but the shaping rule of Section 3.1 applied hour by hour, with loadings $\kappa^j_{m,h}$, adjusts both (Section 7.1). For blocks we use shape factors in levels, s^b, first-order equivalent to log shapes, which accommodate the near-zero midday prices of recent springs.

A pay-as-produced contract on technology j pays the realised price for each unit of output. Relative to baseload, its seasonal value factor is $VF^j_{i,y} = \sum_m w^j_{i,y,m} s_{i,y,m}$ with output shares $w^j_m = (1+\kappa^j_m)/12$: the between-month component of the capture rate. To first order, $VF^j = 1 + \mathrm{Cov}_m(\kappa^j, q)$, the monthly analogue of the covariance form of Lamont [31], so by (2) a technology's penetration lowers its own seasonal value factor and, because the monthly profiles of solar and wind are negatively correlated in Europe, raises the other's (Online Supplement).

Let $\hat{q}^H_{Y,m} = \frac{1}{n_W} \sum_{y \in W} q_{y,m}$ average the log shapes of a window $W = \{Y-h-n_W+1, \ldots, Y-h\}$ of n_W years, so that the forecast for year Y is made at the end of year $Y-h$.

Proposition 2 (Staleness bias)**.** Suppose (2) holds with stable α, δ and κ^j, within-year fuel

deviations with zero mean, and linear penetration growth $z_y^j = z_0^j + \dot{z}_j y$. Then, to first order,

$$\mathbb{E}\big[\hat{q}^H_{Y,m} - q_{Y,m}\big] = \Big(h + \frac{n_W - 1}{2}\Big) \sum_j \theta_j \dot{z}_j \kappa^j_m. \tag{3}$$

The transition-adjusted rule

$$\hat{q}^T_{Y,m} = \hat{q}^H_{Y,m} - \sum_j \hat{\theta}_j \big(\hat{z}^j_Y \bar{\kappa}^j_m - \bar{x}^j_{W,m}\big), \tag{4}$$

with window averages $\bar{\kappa}^j_m$ and $\bar{x}^j_{W,m}$ and penetration forecasts $\hat{z}^j_Y$, has bias $\sum_j [\theta_j(z^j_Y - \hat{z}^j_Y) + (\theta_j - \hat{\theta}_j)(\hat{z}^j_Y - \bar{z}^j_W)]\kappa^j_m$.

Equation (3) is the familiar lag of a moving average behind a linear trend; its content here is the loading. Solar growth makes historical shapes overprice summer relative to winter, and wind growth the reverse. The adjusted rule removes the drift only as far as $\hat{\theta}_j$ and $\hat{z}^j_Y$ are right: with correct penetration forecasts and one technology, its bias is the fraction $(\theta_j - \hat{\theta}_j)/\theta_j$ of the historical bias. The rule keeps the historical shape, which embeds market-specific features that are hard to model, and corrects it for the change in residual-load composition. Applied to a two-harmonic sinusoidal fit, with the correction projected on the same harmonics, it gives a sinusoidal seasonal function with penetration-dependent coefficients, the empirical counterpart of Proposition 1(b). The adjustment is first order and assumes stable profiles and linear growth; whether the drift is linear in the data is an empirical question.

3.2. Seasonal priors and decisions

Many decisions that use a seasonal shape are linear programmes in monthly prices. A reservoir hydro plant with exogenous inflows ι_m (normalised to $\sum_m \iota_m = 12$), storage capacity $\bar{S}$ and monthly turbine limit $\bar{T}$ chooses releases r in

$$X^R = \{r : S_m = S_{m-1} + \iota_m - r_m - \sigma_m,\ 0 \le S_m \le \bar{S},\ 0 \le r_m \le \bar{T},\ \sigma_m \ge 0,\ S_0 = S_{12}\}, \tag{5}$$

with spill σ_m, and earns $\sum_m s_m r_m$, with prices in units of the annual baseload price. A store with capacity one, monthly charge and discharge limits $\bar{c}$ and round-trip efficiency $\eta = \eta_c \eta_d$ chooses net discharges $x_m = d_m - c_m$ in

$$X^B = \{d - c : L_m = L_{m-1} + \eta_c c_m - d_m/\eta_d,\ 0 \le L_m \le 1,\ 0 \le c_m, d_m \le \bar{c},\ L_0 = L_{12}\}. \tag{6}$$

A pay-as-produced contract or a run-of-river plant is the degenerate case of a single feasible profile. In each case the decision maker chooses x in a polytope X and is paid $s'x$; we write $V(s) = \max_{x \in X} s'x$.

Proposition 3 (Seasonal priors and decisions)**.** Let $X \subset \mathbb{R}^{12}$ be a non-empty polytope, $x^* \in \arg\max_{x \in X} s'x$ the schedule that is optimal at the realised shape s, and $\hat{x} \in \arg\max_{x \in X} \hat{s}'x$ the schedule chosen with the prior $\hat{s}$.

(a) Valuation: $x^{*\prime}(\hat{s} - s) \le V(\hat{s}) - V(s) \le \hat{x}'(\hat{s} - s)$. For a single feasible profile w, $V(\hat{s}) - V(s) = w'(\hat{s} - s)$.

(b) Operations: $0 \le V(s) - s'\hat{x} \le (\hat{s} - s)'(\hat{x} - x^*)$. If X has at least two extreme points and x^* is the unique optimum, with gap $\Delta = V(s) - \max\{s'x : x \text{ an extreme point of } X,\ x \neq x^*\} > 0$, and $\bar{d} = \max\{\|x - x^*\|_1 : x \text{ an extreme point of } X\}$, then every prior with $\|\hat{s} - s\|_\infty < \Delta/\bar{d}$ selects $\hat{x} = x^*$: the regret is zero, while the valuation error equals $x^{*\prime}(\hat{s} - s)$.

(c) Intrinsic valuation: if s is random with $\mathbb{E}\|s\| < \infty$ and mean $\bar{s}$, then $\mathbb{E}[V(s)] \ge V(\bar{s})$, with equality if and only if some $x \in X$ is optimal for almost every s.

The proposition collects standard results, in the form needed to read the evidence: the subgradient inequality of the support function V, the predict-then-optimise regret bound of Elmachtoub and Grigas [12], the stability region for optimal solutions studied by Wendell [11] and Jansen et al. [47] and Jensen's inequality for stochastic linear programmes as applied by Madansky [42]. By (a), valuations inherit prior errors one for one; with Proposition 2, a stale prior biases the value of a fixed profile by $w'\mathbb{E}[\hat{s} - s]$, which grows with the drift. By (b), schedules are robust: an error that does not change the optimal extreme point costs nothing, and one that does costs at most the error times the change in the schedule. By (c), the intrinsic value of an expected shape, the usual valuation of flexibility from a forward curve considered by Secomandi [35] and Lai et al. [36], falls short of the expected perfect-foresight value by the expected value of perfect information, $\mathbb{E}[V(s)] - V(\bar{s})$, which depends on the dispersion of shapes, not on the accuracy of the prior's mean, and bounds what re-optimising as information arrives can add, following the information-relaxation interpretation of Brown et al. [43]. Finally, an unbiased but noisy prior, $\mathbb{E}[\hat{s} \,|\, s] = s$, overstates on average both the perfect-foresight value and the value of its own schedule, $\mathbb{E}[V(\hat{s}) \,|\, s] \ge V(s) \ge s'\hat{x}$: the optimizer's curse of Smith and Winkler [41].

Proposition 3 compares a prior with the realised shape, which is unknown when the decision is made. The practical question is whether to re-optimise a schedule on a candidate prior, such as a transition-adjusted shape, instead of an incumbent, such as a smoothed historical one.

Proposition 4 (Updating a seasonal prior)**.** Let X be as in Proposition 3, $\hat{x}_0 \in \arg\max_{x \in X} \hat{s}_0' x$ the schedule of an incumbent prior $\hat{s}_0$, and $\hat{x}_1 \in \arg\max_{x \in X} \hat{s}_1' x$ that of a candidate prior $\hat{s}_1$, with $\hat{x}_1 = \hat{x}_0$ whenever $\hat{x}_0$ is optimal under $\hat{s}_1$. Write $\Delta s = \hat{s}_1 - \hat{s}_0$, $\Delta x = \hat{x}_1 - \hat{x}_0$, $U_1 = \hat{s}_1' \Delta x$, $U_0 = -\hat{s}_0' \Delta x$ and $K = \Delta s' \Delta x = U_1 + U_0$. Then $U_1, U_0, K \ge 0$, and:

(a) Tolerance: the set of $t \in [0, 1]$ for which $\hat{x}_0$ is optimal under $\hat{s}_0 + t\Delta s$ is an interval $[0, t^*]$. The function $u(t) = V(\hat{s}_0 + t\Delta s) - (\hat{s}_0 + t\Delta s)'\hat{x}_0$ is convex, piecewise linear and nondecreasing on $[0, 1]$, zero on $[0, t^*]$ and positive on $(t^*, 1]$, and $u(1) = U_1$. The schedule changes if and only if $t^* < 1$, equivalently $U_1 > 0$, equivalently $K > 0$; if $t^* = 1$, the update changes the valuation by $\hat{x}_0' \Delta s$ and leaves the schedule unchanged.

(b) Realised gain: at the realised shape s, the gain from updating is $G = s'\Delta x = U_1 - (\hat{s}_1 - s)'\Delta x = -U_0 - (\hat{s}_0 - s)'\Delta x$.

(c) Expected gain: suppose $\mathbb{E}\|s\| < \infty$, let $\mathcal{I}$ be the information at the decision date, with $\hat{s}_0$, $\hat{s}_1$, $\hat{x}_0$ and $\hat{x}_1$ determined by $\mathcal{I}$, and let $\mu = \mathbb{E}[s \,|\, \mathcal{I}]$. If the schedule changes, then $\mathbb{E}[G \,|\, \mathcal{I}] = \lambda K - U_0$ with $\lambda = (\mu - \hat{s}_0)'\Delta x / K$; $\lambda = \ell$ if $\mu = \hat{s}_0 + \ell \Delta s$. Updating raises the conditional expected realised value if and only if $\lambda > U_0/K$, a threshold in $[t^*, 1)$; in particular, a candidate with $(\mu - \hat{s}_1)'\Delta x = 0$ ($\lambda = 1$) always clears it.

Part (a) says how much of a correction a schedule absorbs: a share t^* of the update leaves $\hat{x}_0$ optimal, the objective-ranging interval of parametric programming along Δs introduced by Gass and Saaty [48], a one-directional counterpart of the tolerance region of Wendell [11]; if the whole update lies within it ($t^* = 1$), it moves the valuation but not the schedule, and the value the candidate claims for it, U_1, is zero. By (b), given the two schedules, the realised gain depends on the forecast errors only through their projection on the change in the schedule, a single direction among twelve monthly prices, not through their root mean squared errors. Part (c) gives the threshold the evidence must clear. The coefficient λ is the share of the update that the conditional mean endorses in the direction Δx: one if and only if the candidate is conditionally unbiased in that direction, $(\mu - \hat{s}_1)'\Delta x = 0$, zero if and only if the incumbent is, above one if the candidate corrects too little and negative if the update points the wrong way; unbiasedness in the sense $\mathbb{E}[\hat{s}_1 \mid s] = s$ does not imply $\lambda = 1$, because Δx is chosen with $\hat{s}_1$. Since $\mathbb{E}[G + U_0 \mid \mathcal{I}] = \lambda K$, the least-squares slope of realised $G + U_0$ on K through the origin, over past decisions that changed the schedule, estimates a K^2-weighted average of λ. If λ is common to the decisions, or at least does not vary with K and U_0/K and is stable over time, a schedule should be updated when the estimate exceeds U_0/K; since $U_0/K < 1$, an estimate above one then implies updating whenever the schedule changes.

3.3. Econometric specification

We estimate (2) with the identity (1):

$$\ln P_{i,y,m} = \mu_{i,y} + \alpha_{i,m} + \gamma_i^G \ln C_{y,m}^G + \gamma_i^K \ln C_{y,m}^K + \sum_{j \in \mathcal{J}} \beta_j \, x_{i,y,m}^j + \beta_D \, d_{i,y,m} + e_{i,y,m}, \qquad (7)$$

where x^j and $d = 100\,\delta$ are measured in percent of mean monthly demand, so that $\beta_j = -\theta_j/100$ and the coefficients reported $\times 100$ equal $-\theta_j$ and θ_D. Market$\times$year effects $\mu_{i,y}$ absorb price levels, so (7) models the log shape; market$\times$month effects $\alpha_{i,m}$ absorb the stable seasonal baseline; C^G and C^K are short-run cost proxies of gas- and coal-fired plants with market-specific pass-through. The coefficients are identified from year-to-year changes in each market's seasonal generation profile, due to penetration growth and weather, including changes common to all markets; we therefore add rest-of-Europe (leave-own-out) deviations, market$\times$month trends, and year$\times$month effects $\lambda_{y,m}$, which absorb every common seasonal shock that moves log prices equally in all markets. Coupled neighbours' deviations weight each neighbour's deviation by the share of hours in the previous year in which the two day-ahead prices were within 1 EUR/MWh or 2% of each other. A reduced form replaces the solar and wind terms with $\sum_{j \in \{S,W\}} z^j \sum_{k=1}^{2} b'_{jk} f_k(m)$, keeping the hydro, nuclear and demand deviations as controls: the empirical version of Proposition 1(b). For hour blocks, the dependent variable is the block shape factor or the midday-minus-night contrast $s^{mid} - s^{night}$ instead of the log price. Standard errors follow Driscoll and Kraay [49] to allow for cross-sectional and serial correlation, with year-month as the time dimension and four lags; we also report two-way clustered standard errors by market and year-month following Cameron et al. [50] and, for key coefficients, by market and year and a wild cluster bootstrap by year. Hydro and nuclear deviations partly respond to prices and measured demand includes net imports, so their coefficients are reduced-form responses rather than structural slopes.

4. Data and stylised facts

Hourly day-ahead prices come from the European wholesale electricity price data provided by Ember [51], compiled from ENTSO-E (with EMR for Great Britain and SEMOpx for the all-island Irish market) under a CC BY 4.0 licence. For countries with several bidding zones Ember publishes national series; for Italy its annual averages coincide with the national single price, and for Denmark, Norway and Sweden its series closely match fixed-weight averages of zonal prices (Online Supplement), so they are not traded prices. Daily and monthly baseload prices are averages of local-time hourly prices from 1 January 2015 to 31 December 2025, and block prices average local hours 00:00-05:59 (night), 10:00-15:59 (midday) and 18:00-21:59 (evening); our monthly averages match Ember's published ones (Online Supplement). The sample comprises 27 markets (listed in the Online Supplement). Austria enters in 2019, after the split of the joint German-Austrian zone, Ireland in 2019, Bulgaria and the United Kingdom in 2017 and Croatia in 2018, and Switzerland ends in 2024: 281 complete market-years and 3,372 market-months.

Monthly generation by technology and demand come from the monthly electricity data of Ember [52], available for 3,276 market-months; generation refers to countries (the United Kingdom including Northern Ireland, Ireland the Republic only). European penetration is a demand-weighted average over markets with generation data, chain-linked across years so that markets entering or leaving the data do not create breaks. Fuel costs use the World Bank Pink Sheet compiled by World Bank [53] for the Dutch TTF gas price and the Australian coal price, converted with the monthly euro exchange rate from Board of Governors of the Federal Reserve System [54], and the monthly average of EEX primary auction prices for EU allowances published by European Energy Exchange [55]. The cost proxies are $C^G = \mathrm{TTF}/0.50 + 0.37\,\mathrm{EUA}$ and $C^K = \mathrm{coal}/(6.978 \times 0.38) + 0.90\,\mathrm{EUA}$ in EUR/MWh of electricity. Coupling intensity is computed from the hourly prices of each pair of interconnected markets. The within-day solar profiles of Section 7.1 use hourly solar output from the Open Power System Data time series provided by Open Power System Data [56] for 2015-2020, and a check for Germany uses actual quarter-hourly solar output provided by Bundesnetzagentur [57]. For Germany, France and Italy, daily settlement prices of EEX baseload futures (monthly, quarterly and annual contracts, 2015-2025; LSEG Datastream and EEX) provide a market benchmark; unlike the other data, they are proprietary. The Online Supplement details sample construction and validation.

Figure 1 compares average shape factors in 2017-2019 and 2023-2025, equal three-year windows, for the eight largest markets by demand; all 27 are shown in the Online Supplement. The summer (April-September) to winter (January-March and October-December of the same year) price ratio fell by a median of 0.13 log points across all 24 markets with three years of data in 2017-2019, although Switzerland contributes only two years to the 2023-2025 window. The changes were marginal in Belgium, Italy and Switzerland, and declines also occurred in the three markets with shorter initial windows; relative to 2015-2019 the ratio fell in 25 of 27 markets. The decline holds in all 24 markets net of the within-year path of fuel costs, with a median of -0.12, and doubles when winters are defined as contiguous October-March periods, reaching a median of -0.25. The seasonal amplitude, the standard deviation of the average monthly log shape, rose in 18 of 24 markets, with its median increasing from 8.2% to 14.1%; after removing sampling noise, the corresponding medians rose from 4.9% to 7.3%. The largest increases

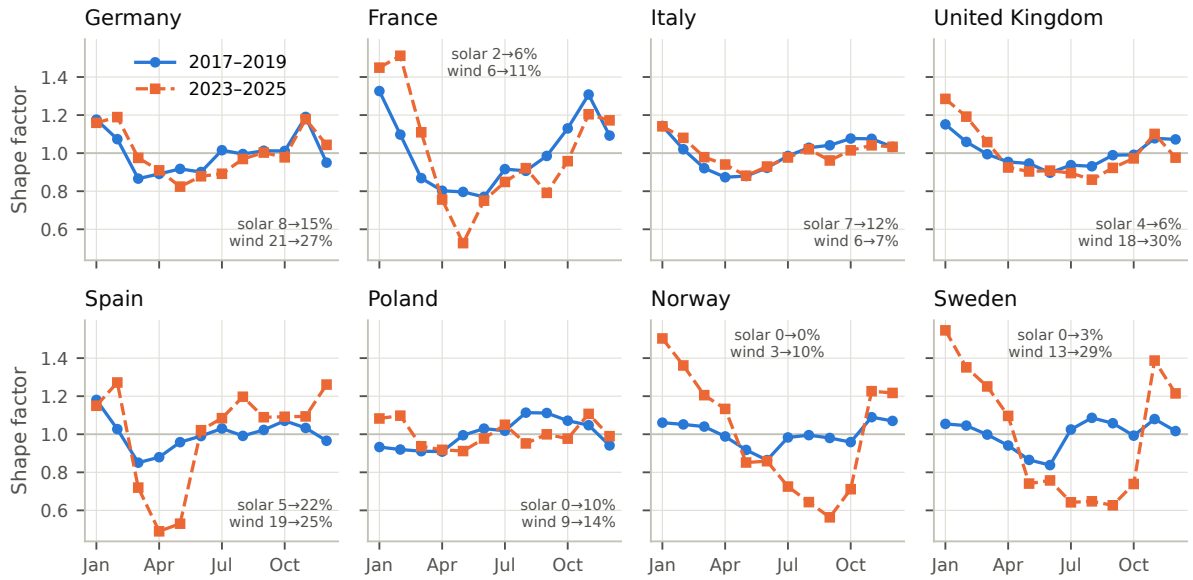


Figure 1: Average monthly shape factors (monthly baseload price relative to the annual average) in 2017-2019 and 2023-2025 for the eight largest markets by demand. Labels give average solar and wind penetration (share of annual demand) in the two windows.

occurred in Spain, Portugal, Sweden, Finland and Norway, where the amplitude rose four- to fivefold. In Spain the trough moved to April and May, while in the wind-heavy Nordic markets summer prices collapsed and winter prices rose.

Because the classical seasonal-stability tests of Canova and Hansen [58] behave poorly with eleven years of monthly data, we permute whole years under the null of an unchanged seasonal pattern. Pooling the 21 markets observed in all years 2015-2019 and 2023-2025, the average decline of the log summer/winter ratio is 0.15, with $p = 0.018$, the smallest value attainable with 56 permutations. Net of fuel paths, the decline is 0.14 and the associated p-value is 0.071. Market by market, seven of 22 tests reject at 5%. A test on the full twelve-month shape gives $p = 0.21$, which falls to $p = 0.054$ net of fuel paths: the change is concentrated in the summer-winter contrast.

The change is specific to electricity: the seasonal shapes of the gas- and coal-plant cost proxies did not change, with permutation p-values of 0.46 and 0.54, respectively. Among seventeen World Bank commodities, including TTF gas, only maize shows a change significant at 5%, about what chance alone would produce (Online Supplement). Across markets, however, the size of the change does not line up with own penetration growth: its coefficients on the changes in solar and wind penetration are $+0.007$ and -0.009, respectively, each with a standard error of 0.004, and have signs opposite to those implied by Proposition 1(a).

Figure 2 shows that the change is concentrated in the middle of the day. Measured as the difference between the summer and winter means of block shape factors, in percent of the annual baseload price, the summer/winter contrast of midday prices fell in all 24 markets, with a median change of -42, moving from -10 in 2017-2019 to -58 in 2023-2025, whereas that of night prices was essentially unchanged, with a median of $+1$ and declines in 9 of 24 markets, and that of evening prices rose, with a median of $+8$ and declines in 8 markets; for all hours together the median change is -12. In the pooled permutation test, the midday decline is significant with

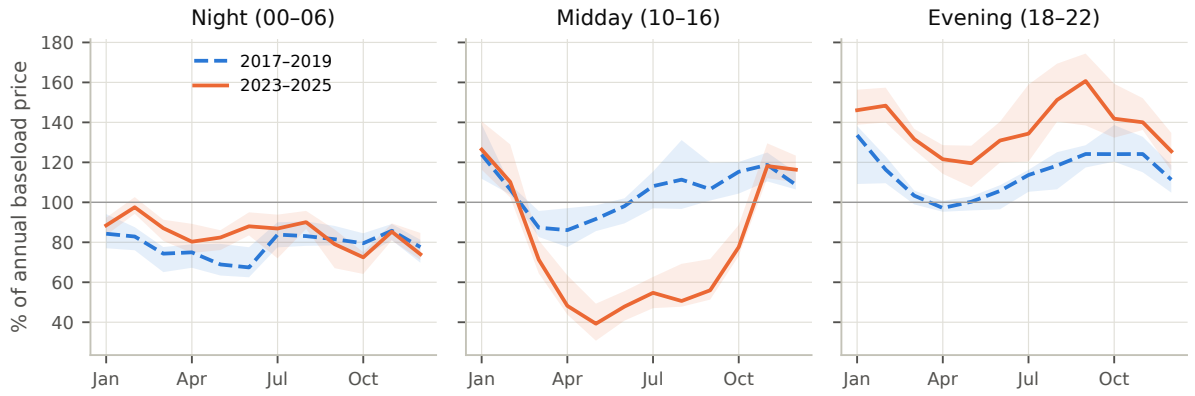


Figure 2: Median block shape factors in 2017-2019 and 2023-2025 (mean price of the block's local hours relative to the annual baseload average) across the 24 markets with three years of data in 2017-2019; shaded: interquartile range.

$p = 0.018$, whereas the night and evening changes are not, with respective p-values of 0.77 and 0.80. The yearly cross-market median of the midday contrast was between -19 and $+1$ in 2015-2019 and fell to -38, -53 and -72 in 2023, 2024 and 2025.

5. Estimation results

Table 1 reports estimates of (7). In column (1), a within-year increase in solar output equal to 1% of mean monthly demand lowers the monthly price by 1.76% relative to the annual mean; the effects of wind, hydro and nuclear output are -1.41%, -0.80% and -0.55%, and a 1% increase in demand raises the price by 1.23%. All are significant with Driscoll-Kraay standard errors and with two-way clustering, although nuclear is only marginally significant at $t = 2.0$; the solar coefficient also survives clustering by market and year with a standard error of 0.42 and a wild cluster bootstrap by year with $p = 0.007$. Excluding 2021-2022 changes the solar coefficient little: it is -1.73% in column 2, and across nine robustness variants in the Online Supplement it ranges from -1.61% to -1.94% and the wind coefficient from -1.08% to -1.57%. The solar coefficient exceeds the demand coefficient by 0.53, with a standard error of 0.29, which Assumption 1 rules out for exogenous demand; measured demand includes pumping and net imports, which respond to prices.

Columns (3)-(6) show where the solar coefficient comes from. With rest-of-Europe deviations, the own solar coefficient falls to -0.10% while rest-of-Europe solar output carries -4.90% and rest-of-Europe demand is also significant, with a coefficient of 1.81%; with market×month trends the rest-of-Europe coefficient becomes -6.35%, while the own coefficient is -0.30%, so it does not reflect smooth drifts in each market's shape. Year×month effects (column 5) absorb every common shock: wind, hydro, nuclear and demand remain significant, while the solar coefficient becomes 0.11%, with a standard error of 0.28, although 42% of the within variation in solar deviations survives these effects, against 82% for wind and 91% for hydro. In column (6), coupled neighbours' wind and hydro output lower a market's price, with coefficients of -1.57% and -2.17%, respectively, consistent with transmission through market coupling, and they dominate uncoupled neighbours' deviations in a horse race (Online Supplement); neighbours' solar output does not, with a coefficient of -0.48% and a standard error of 0.41. At monthly baseload

Table 1: Monthly log price on seasonal deviations of residual-load components.

	(1)	(2)	(3)	(4)	(5)	(6)
Solar	−1.76***	−1.73***	−0.10	−0.30	0.11	−0.13
	(0.28) [0.33]	(0.30) [0.34]	(0.26) [0.37]	(0.49) [0.55]	(0.28) [0.41]	(0.26) [0.33]
Wind	−1.41***	−1.23***	−1.19***	−1.36***	−1.09***	−0.94***
	(0.17) [0.19]	(0.17) [0.21]	(0.16) [0.19]	(0.17) [0.25]	(0.16) [0.19]	(0.13) [0.15]
Hydro	−0.80***	−0.72***	−0.77***	−0.84***	−0.72***	−0.59***
	(0.09) [0.25]	(0.10) [0.23]	(0.10) [0.25]	(0.11) [0.31]	(0.10) [0.26]	(0.09) [0.19]
Nuclear	−0.55***	−0.56***	−0.55***	−0.47***	−0.55***	−0.56***
	(0.12) [0.27]	(0.14) [0.29]	(0.10) [0.29]	(0.10) [0.29]	(0.10) [0.26]	(0.10) [0.25]
Demand	1.23***	0.96***	0.89***	1.01***	0.78***	0.79***
	(0.17) [0.22]	(0.13) [0.21]	(0.14) [0.19]	(0.16) [0.22]	(0.15) [0.21]	(0.14) [0.21]
Rest-of-Europe solar			−4.90***	−6.35***		
			(0.56) [0.85]	(1.39) [1.44]		
Rest-of-Europe wind			−1.87***	−1.60***		
			(0.33) [0.35]	(0.34) [0.40]		
Rest-of-Europe demand			1.81***	1.51***		
			(0.41) [0.52]	(0.43) [0.55]		
Coupled neighbours' solar						−0.48
						(0.41) [0.64]
Coupled neighbours' wind						−1.57***
						(0.38) [0.58]
Coupled neighbours' hydro						−2.17***
						(0.52) [0.75]
Market×year, market×month FE	✓	✓	✓	✓	✓	✓
Fuel pass-through	✓	✓	✓	✓	✓	✓
Rest-of-Europe hydro, nuclear			✓	✓		
Market×month trends				✓		
Year × month FE					✓	✓
Excludes 2021–2022		✓				
Within R^2	0.633	0.434	0.676	0.718	0.182	0.217
Observations	3,276	2,628	3,276	3,276	3,276	3,264

Coefficients (×100) on within-year deviations of monthly generation and demand, in percent of mean monthly demand; −1.76 means that output equal to 1% of demand lowers the monthly price by 1.76%. Rest-of-Europe deviations are demand-weighted averages over all other markets with complete generation data in the year. Coupled neighbours' deviations weight each interconnected neighbour's deviation by its demand and by the share of hours in the previous year with day-ahead prices within 1 EUR/MWh or 2% of each other. Driscoll-Kraay standard errors (year-month time dimension, four lags) in parentheses, two-way clustered standard errors (market and year-month) in brackets. Within R^2 relative to the fixed effects of each column. ***, **, * denote significance at 1%, 5%, 10% (Driscoll-Kraay).

resolution, therefore, wind, hydro, nuclear and demand effects are identified from market-specific variation, whereas the solar effect rests on variation common to European markets, which coincides with the solar build-out but cannot be separated from other common seasonal shocks.

Proposition 1 applies the same coefficient to changes in a technology's monthly output deviation that come from penetration growth and from weather. We split each deviation into a penetration part, $z^j_{i,y}\bar{\kappa}^j_{i,m}$, the deviation that the market's average profile implies at the year's penetration, and a weather part, the year's departure from that profile. For solar, 88% of the within variation is penetration-driven, and the penetration and weather parts have similar coefficients of −1.19 and −0.84, respectively; their difference is −0.35, with a standard error of 1.06, although the confidence interval of the difference is wide. For wind, 90% of the within variation is weather-driven, and the parts have opposite signs: −1.63 for the weather part and +1.28 for the penetration part. The difference of 2.91 is significant with Driscoll-Kraay standard errors at $p = 0.001$ and with a wild cluster bootstrap by year at $p = 0.016$, but not with two-way clustering by market and year, for which $p = 0.11$; with year×month effects the difference is 2.66, and the three corresponding p-values are 0.004, 0.005 and 0.06. More wind capacity has not

Table 2: Hour-block prices on seasonal deviations of residual-load components.

	Midday − night spread			Block shape factor		
	(1)	(2)	(3)	Night	Midday	Evening
Solar	−4.73***	−2.05***	−2.02***	0.73***	−1.32***	0.81**
	(0.25) [0.56]	(0.17) [0.50]	(0.17) [0.52]	(0.24) [0.30]	(0.24) [0.48]	(0.41) [0.61]
Wind	0.08	−0.01	0.07	−1.07***	−1.08***	−1.12***
	(0.14) [0.17]	(0.11) [0.12]	(0.09) [0.12]	(0.17) [0.19]	(0.16) [0.21]	(0.21) [0.24]
Hydro	0.06	0.04	0.02	−0.63***	−0.59***	−0.74***
	(0.09) [0.09]	(0.07) [0.07]	(0.07) [0.07]	(0.09) [0.27]	(0.11) [0.26]	(0.12) [0.23]
Nuclear	−0.23***	−0.19***	−0.19***	−0.40***	−0.59***	−0.72***
	(0.08) [0.09]	(0.07) [0.08]	(0.07) [0.08]	(0.08) [0.19]	(0.11) [0.26]	(0.14) [0.31]
Demand	0.80***	0.69***	0.72***	0.42***	1.11***	0.85***
	(0.13) [0.19]	(0.13) [0.20]	(0.13) [0.19]	(0.15) [0.19]	(0.20) [0.27]	(0.25) [0.33]
Coupled neighbours' solar			−1.14***			
			(0.24) [1.14]			
Coupled neighbours' wind			−0.57**			
			(0.29) [0.37]			
Coupled neighbours' hydro			0.01			
			(0.24) [0.31]			
Year × month FE		✓	✓	✓	✓	✓
Within R^2	0.496	0.187	0.198	0.228	0.199	0.158
Observations	3,276	3,276	3,264	3,276	3,276	3,276

Dependent variables (×100, percent of the annual baseload price): difference between the mean prices of local hours 10:00-15:59 and 00:00-05:59 in the month (columns 1-3); mean price of local hours 00:00-05:59, 10:00-15:59 and 18:00-21:59 relative to the annual baseload average (night, midday, evening). Regressors as in Table 1; all columns include market×year and market×month effects and market-specific fuel pass-through. Driscoll-Kraay standard errors in parentheses, two-way clustered (market and year-month) in brackets; significance from Driscoll-Kraay. Wild cluster bootstrap by year for own solar: $p \leq 0.003$ in columns (1)-(3) and for the midday block.

lowered winter prices the way windy months do, possibly because reservoir hydro, interconnector flows and other flexible resources adapt to predictable changes in the generation mix but not to weather. The wind coefficient of the TDS rules is thus identified from weather; on samples ending before 2024 the difference is not significant (Online Supplement), so a forecaster could not have dropped the wind channel ex ante before 2024, or before 2023 for midday shapes. Estimated on samples ending in 2019 to 2025, the solar coefficient stays between −0.4% and −1.2% until 2023 and reaches −1.37% and −1.76% only with 2024 and 2025, years in which European solar penetration reached 10.3% and 12.4% of demand.

Table 2 repeats the analysis for hour blocks. Columns (1)-(3) use the within-day contrast between midday and night prices, which absorbs the month's fuel costs, demand level and hydro conditions. A solar deviation of 1% of demand lowers the midday price relative to the night price by 4.7% of the annual baseload price, and by 2.1% with year×month effects; the latter estimate has a two-way clustered standard error of 0.50 and a wild cluster bootstrap by year yielding $p < 0.001$. With coupled neighbours' deviations, own solar keeps a coefficient of −2.0, and neighbours' solar output has a coefficient of −1.1 in the same direction. The neighbour effect is significant with the wild bootstrap by year, yielding $p = 0.007$, but not with two-way clustering by market and year-month, which gives a standard error of 1.1. Wind and hydro deviations, which affect all hours alike, do not move the contrast. The last three columns show where the effect sits. With year×month effects, solar output lowers midday prices by 1.3% of the annual baseload price per 1% of demand and raises night and evening prices by 0.7% and 0.8%; for the monthly baseload average the coefficient is 0.01, with a Driscoll-Kraay standard error of 0.27 and a two-way clustered standard error of 0.34. Market-specific solar output therefore reallocates

value within the day rather than shifting the monthly level, consistent with flexible resources moving energy away from sunny hours, and the monthly baseload regressions of Table 1 cannot detect it. The TDS restriction is again not rejected for solar and fails for wind in the midday block (Online Supplement). Midday shapes are thus where the transition-dependent seasonality of Proposition 1 can be identified market by market. Results on intraday price variability, relevant for the seasonal volatility functions studied by Fanelli et al. [20], Fanelli and Schmeck [21] and Frau and Fanelli [10], are in the Online Supplement.

6. Out-of-sample shaping

At the end of each year $b = Y - h$ every method forecasts the twelve shape factors of year Y for every market, using data up to b only; the main evaluation uses $h = 1$ and target years 2018-2025. Benchmarks use price history only: FLAT sets all shape factors to one; HIST1 uses last year's shape; HIST5 averages the shapes of the last three to five years, and HIST5x skips the crisis years 2021-2022, which uses hindsight; SIN5 fits a two-harmonic sinusoidal function to the last five log shapes; EWMA, TREND-F, DRIFT-F and EUSHIFT overweight recent shapes or extrapolate own or cross-market drifts (Online Supplement). The TDS forecasts apply rule (4) for solar and wind, with coefficients re-estimated at every origin and target-year penetration extrapolated from the last two observed years: TDS-H adjusts HIST5 and TDS-F adjusts SIN5 with own and neighbours' penetration, and TDS-EU adjusts SIN5 with European penetration, own profiles and the sum of own and neighbours' coefficients. TDS-EU was added after the estimates of Section 5, TDS-F before them; the two differ by at most 0.11 points on any measure of Table 3. SIN5c and TDS-EUc remove the within-year fuel-cost path from the history and add back the average fuel profile of past years. Variants marked ‡ use realised penetration and are conditional forecasts. The midday evaluation applies the rules, in levels, to the seasonal shape of midday prices, with TDS coefficients estimated on the midday block shape factor.

We report the root mean squared error (RMSE) of the twelve shape factors in percent of the annual price; the within-quarter RMSE of log shapes, the part that quarterly futures do not fix; the bias and the absolute error of the summer/winter contrast; and the bias of the seasonal value factors of solar and wind output. Absolute errors measure accuracy; mean biases show the direction of errors, which crisis years of opposite sign partly cancel. Differences between methods are tested on paired market-year losses following Diebold and Mariano [59] with standard errors clustered by market and target year and $t(7)$ critical values, and with a wild cluster bootstrap by target year that recomputes the two-way statistic following Cameron et al. [60], MacKinnon et al. [61] and Webb [62]; with eight target years and strong common shocks, we also report the minimum detectable effect (MDE) at 5% size and 80% power. An evaluation on fuel-adjusted shapes is in the Online Supplement.

Table 3 shows that seasonal shapes are hard to forecast: RMSEs are 27-29% of the annual price for the better rules, and FLAT has an RMSE of 28.6%, close to the 28.8% of HIST5. Smoothing helps: SIN5 lowers the RMSE of HIST5 by 1.1 points, with $p = 0.04$, and the SIN5c variant, which removes fuel paths from the history, lowers it by a further 0.9 points, with $p = 0.25$. Rules that overweight or extrapolate recent shapes do worse: EWMA, TREND-F and DRIFT-F have RMSEs 2.9 to 7.1 points higher than SIN5, with respective p-values of 0.02, 0.02 and 0.07,

Table 3: One-year-ahead shaping accuracy on raw monthly baseload shapes, 2018-2025.

Method	RMSE	Within-quarter RMSE	S/W bias	|S/W| error	Solar VF bias	Wind VF bias
Price history						
FLAT	28.61	14.63	13.40	22.01	3.97	−0.40
HIST1	37.99	18.89	2.46	26.01	0.30	0.54
HIST5	28.85	15.00	1.63	18.97	0.45	0.82
HIST5x†	26.94	14.50	1.05	18.20	0.51	0.85
SIN5	27.78	14.36	2.81	18.87	1.01	0.61
EWMA	30.67	15.55	2.24	21.22	0.50	0.68
TREND-F	34.90	17.17	4.01	27.79	1.20	0.21
DRIFT-F	33.21	16.58	3.84	25.55	1.29	0.10
EUSHIFT	33.37	16.72	2.04	23.69	0.60	0.46
Transition-adjusted						
TDS-H	29.05	15.20	−0.44	17.99	−0.19	1.03
TDS-F	27.77	14.45	0.64	17.71	0.38	0.83
TDS-EU	27.74	14.39	0.70	17.78	0.27	0.86
Clean history, average fuel path						
SIN5c	26.87	14.08	3.79	18.80	1.31	0.52
TDS-EUc	26.75	14.10	1.68	17.66	0.57	0.77
Conditional on realised penetration						
TDS-EU‡	27.80	14.42	−0.35	17.62	−0.05	0.99
Paired differences: p-values, two-way clustered [wild bootstrap by year]						
SIN5 − HIST5	0.04 [0.04]	0.13 [0.13]	0.01 [0.00]	0.77 [0.78]	0.00 [0.00]	
TDS-EU − SIN5	0.89 [0.90]	0.69 [0.68]	0.17 [0.17]	0.37 [0.39]	0.12 [0.12]	
SIN5c − SIN5	0.25 [0.31]	0.11 [0.17]	0.08 [0.06]	0.90 [0.94]	0.10 [0.10]	
TDS-EUc − SIN5c	0.68 [0.68]	0.80 [0.80]	0.17 [0.17]	0.34 [0.36]	0.12 [0.12]	
MDE, TDS-EUc − SIN5c	0.91	0.25	4.49	3.67	1.37	
Market-years	200	200	200	200	181	181

RMSE of monthly shape factors and within-quarter RMSE of log shapes in percent of the annual price; summer/winter (S/W) bias and absolute error of the log summer/winter ratio in log points (×100; positive values mean summer is overpriced relative to winter); value-factor (VF) biases in percent of the baseload price. Common sample of market-years for which all methods are available. †Uses hindsight. ‡Uses realised penetration of the target year. p-values for paired differences, two-way clustered by market and target year with $t(7)$ critical values; wild cluster bootstrap by target year with Webb weights in brackets. MDE: minimum detectable effect at 5% size and 80% power.

and EUSHIFT is 5.6 points higher, with $p = 0.08$. The transition adjustment does not lower the RMSE, with a TDS-EU minus SIN5 difference of -0.04 and $p = 0.89$, or the within-quarter RMSE, and it lowers the absolute summer/winter error only from 18.9 to 17.8 log points, with $p = 0.37$. It shifts the mean errors in the direction of the transition, from 2.8 to 0.7 log points for the summer/winter error, with $p = 0.17$, and from 1.0 to 0.3 points for the solar value factor, with $p = 0.12$, but neither baseline bias is significant, with standard errors of 8.0 and 2.4, and most of the shift comes from 2025; excluding that year gives a difference of -1.0 log points, with $p = 0.36$. Realised penetration moves the mean errors further, but not the absolute errors. Only reductions of 0.9 points in RMSE or 3.7 log points in the absolute summer/winter error would be detected with 80% probability.

Figure 3(a) shows why the evaluation is noisy. Fuel-price shocks dominate the yearly biases: every rule underpredicted summer prices in 2022, when gas prices peaked in August (by 45 log points for HIST5), and overpredicted them in 2021 (by 22). The transition drift is visible only late in the sample: in 2025 the summer/winter bias was 21.5 log points for HIST5 and 22.8 for SIN5c, against 13.2 for TDS-EUc and 10.3 with realised penetration. On identical target years

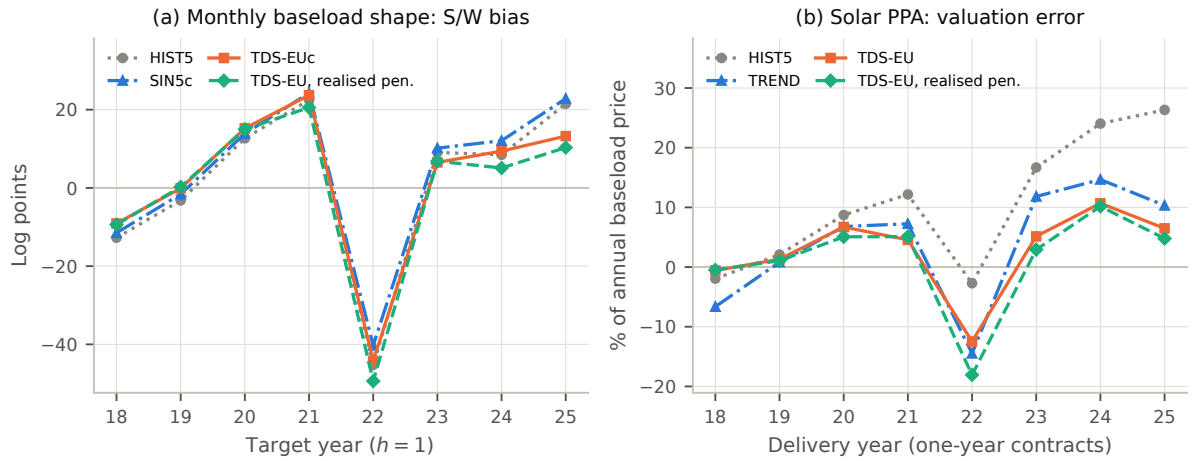


Figure 3: Out-of-sample errors by target year, one year ahead. (a) Summer/winter bias of monthly baseload shapes, log points; positive values mean summer is overpriced relative to winter. (b) Valuation error of one-year solar PPAs with hourly forward curves (Section 7.1): forecast minus realised capture rate in points of the annual baseload price; positive values overvalue solar output.

2020-2025, the bias of HIST5 grows neither with the horizon nor with the window length, and across markets it rises with regional wind growth rather than falling as Proposition 2 predicts (Online Supplement): these target years do not support its linear-drift premise.

One year ahead, the historical rules overprice summer middays by 11.5 points of the annual baseload price for HIST5 and 10.5 for SIN5 and overstate the seasonal component of solar capture values by about 4 points, mean biases that are not individually significant; for SIN5, the corresponding standard errors are 9.5 and 2.9, as detailed in the Online Supplement. The TDS adjustment removes the mean summer bias, leaving −0.1 for TDS-F and −0.2 for TDS-EU, and most of the capture-value bias, leaving 0.8 and 0.5 points, respectively. These shifts are significant, with p-values between 0.02 and 0.04 also under the wild bootstrap; excluding 2025, the two p-values for TDS-EU are 0.04 and 0.03, but absolute errors fall less and not significantly: the absolute summer/winter error declines by 18-20%, with p-values between 0.23 and 0.28. The gain is concentrated in 2023-2025, whereas in 2022 the adjustment added to the underprediction of summer prices caused by the gas crisis. Dropping the wind channel, ex post or by an ex ante rule, does not significantly improve on TDS-EU (Online Supplement).

Forecasting 2023-2025 from the end of 2019 across 71 market-years, HIST5 and SIN5 overprice summer relative to winter by 14.1 and 15.3 log points. With penetration extrapolated from 2017-2019, TDS-EU does worse, with a bias of 18.3 log points: the wind adjustment pushes in the wrong direction, as the full-sample restriction test suggests, and linear extrapolation missed most of the post-2020 solar boom. With the realised European penetration, the bias falls to 10.4 log points and the absolute error from 17.4 to 13.8, through the solar channel (Online Supplement). For midday shapes, the historical rules overprice summer middays by 43 points; the TDS prior hardly helps with extrapolated penetration (40.2) but removes three fifths of the bias with the realised one (16.8). With a single origin, these comparisons are descriptive.

For Germany, France and Italy, the futures traded at the end of each year imply a monthly shape for the next year (Online Supplement). Over 2018-2025, this market shape was about as accurate as the priors: its absolute summer/winter error was 20.3 log points, against 21.5 for

HIST5 and 20.0 for TDS-EU. Its summer/winter bias was more negative, at −4.1 log points against −0.5 for TDS-EU, a difference that comes from France and Germany, not Italy; in 2023-2025 the bias was smaller, at 7.2 log points against 14.4 for HIST5 and 12.0 for TDS-EU.

7. Decision consequences

We now feed the out-of-sample priors into three decisions and settle them at realised prices: the valuation of solar PPAs with hourly forward curves, the scheduling of energy-limited assets, including the choice between two priors, and the valuation of seasonal flexibility. The PPA valuations allow tests; with eight target years and common shocks, the schedules mainly illustrate Propositions 3 and 4.

7.1. Solar PPAs and hourly forward curves

A pay-as-produced PPA pays a fixed price for every unit of a plant's output. Per unit of output and in units of the annual baseload price, the contract's value at spot prices, the fixed price at which it breaks even, is the profile capture rate $CR = \sum_m w_m \sum_h (\nu_{m,h}/24)\, s_{m,h}$, where w_m is the share of annual output produced in month m, $\nu_{m,h}$ the output at local hour h relative to the month's hourly mean, and $s_{m,h}$ the mean price of hour h in month m relative to the annual baseload price. An hourly price forward curve supplies $\hat{s}_{m,h}$ and thereby $\widehat{CR}$. We take w_m from the market's solar fleet in the delivery year, so that errors isolate the price shape as in Section 6, and $\nu_{m,h}$ from the fleet's average hourly output in 2015-2020 from Open Power System Data [56], for the 22 markets where such data exist or can be borrowed from a nearby market (Online Supplement). The profile defines the contract and enters forecast and realised capture rates alike, so its 2020 data give the rules no information on prices. The annual level is set to its realised value, so that $\widehat{CR} - CR$, in points of the annual baseload price, is a shape-only valuation error; times the realised baseload price, it is the error in the break-even price in EUR/MWh. HIST1 and HIST5 take last year's hourly shape or the average of the last three to five years, a common construction of hourly forward curves; TREND adds to HIST5 the linear drift of each hour-month cell over the window, which, up to truncation at zero, extrapolates the capture rate linearly; and TDS-EU applies rule (4) hour by hour to HIST5, with the coefficients and extrapolated European penetration of the monthly TDS-EU rule and the hourly loadings of Section 3, which adjust both the seasonal shape and the level of each hour relative to baseload; TDS-F, specified before the estimates of Section 5, uses own and neighbours' penetration instead. We value one-year contracts for delivery in 2018-2025, covering 161 market-years; three-year contracts signed at the end of each year 2017-2022, covering 117 contracts in 22 markets; and a six-year contract signed at the end of 2019, covering 19 markets.

History-based hourly forward curves overvalued solar output (Table 4). One year ahead, HIST5 overpriced the capture rate by 11.2 points of the annual baseload price on average, equivalent to 8.9 EUR/MWh, with $p = 0.03$ for a zero mean and by 22.2 points in 2023-2025 (Figure 3b), when the median realised capture rate fell from 1.04 and 0.99 in 2018 and 2019 to 0.82, 0.69 and 0.59. The TDS adjustment cut the mean overvaluation to 2.8 points, not significantly different from zero, with $p = 0.34$. Because a difference of signed errors is a difference of forecasts,

Table 4: Valuation errors of pay-as-produced solar PPAs with hourly price forward curves.

	One-year contracts, 2018–2025 (161)					Three-year contracts (117)	
Rule	Bias	|Error|	RMSE	|Error|, EUR/MWh	Bias 2023–25	Bias	|Error|
FLAT	12.48	15.53	20.31	14.01	27.8	11.11	12.34
HIST1	5.47	10.13	12.68	11.40	13.9	9.77	10.57
HIST5	11.15	13.02	16.51	12.53	22.2	12.45	12.56
TREND	4.10	10.11	12.36	11.67	12.3	5.45	8.59
TDS-F	2.09	7.69	9.86	9.38	5.9	4.82	7.67
TDS-EU	2.77	7.67	9.49	9.17	7.5	5.12	6.91
TDS-EU, seasonal part only	9.31	11.64	14.43	11.59	18.6	11.02	11.18
TDS-EU, solar only[§]	2.10	7.32	9.16	8.94	6.7	4.30	6.42
TDS-EU, ex ante selection	3.83	8.66	10.52	10.22	7.3	7.22	8.37
TDS-EU[‡]	1.27	7.91	10.26	10.59	6.0	0.87	7.01
Paired differences: p-values, two-way clustered [wild bootstrap by year or origin]							
TDS-EU − HIST5	0.01 [0.01]	0.09 [0.09]	0.10 [0.11]	0.39 [0.39]		0.05 [0.06]	0.06 [0.06]
TDS-EU − TREND	0.39 [0.39]	0.03 [0.02]	0.02 [0.03]	0.02 [0.02]		0.88 [0.90]	0.15 [0.16]
TDS-EU − HIST1	0.19 [0.13]	0.18 [0.04]	0.18 [0.04]	0.23 [0.05]		0.15 [0.12]	0.20 [0.07]
TDS-F − HIST5	0.01 [0.02]	0.11 [0.12]	0.12 [0.14]	0.45 [0.48]		0.05 [0.07]	0.10 [0.07]
TDS-EU − HIST5, excl. 2025	0.02 [0.03]	0.19 [0.18]	0.20 [0.22]	0.67 [0.66]			
MDE, TDS-EU − HIST5	7.72	8.84		11.97		9.71	7.96

Forecast minus realised profile capture rate in points of the annual baseload price (positive values overvalue solar output); |Error|, EUR/MWh: absolute error times the realised annual baseload price. Output profile: realised monthly output shares of the market's solar fleet and its average within-day profile in 2015-2020. HIST1, HIST5: last year's and the last three to five years' average hourly shape; TREND: HIST5 plus the linear drift of each hour-month cell over the window; TDS-EU and TDS-F: HIST5 adjusted hour by hour by rule (4) with European or own and neighbours' penetration; seasonal part only: without the change in the annual level of each hour. Three-year contracts: signed at the end of each year 2017-2022 for delivery in the next three years, errors averaged over the tenor with equal annual volumes. [§]Chosen ex post. [‡]Realised penetration. Market-years or contracts in parentheses. *p*-values two-way clustered by market and delivery year (one-year contracts) or origin (three-year contracts), wild cluster bootstrap by year or origin in brackets; MDE as in Table 3.

the shift in the mean, which has $p = 0.009$, shows only that the adjustment moved valuations; accuracy is measured by absolute and squared errors. The mean absolute error fell from 13.0 to 7.7 points ($p = 0.09$; squared errors $p = 0.10$) and was 2.4 points below that of TREND ($p = 0.03$; squared errors $p = 0.02$), whose extrapolation cut the bias to 4.1 points; against HIST1 the gain is significant only with the wild bootstrap, and the pre-specified TDS-F performs like TDS-EU. The gain sits in 2023-2025, with absolute errors of 22.2 points for HIST5 and 9.5 for TDS-EU, against 6.9 and 6.5 in 2018-2022; the overvaluation of 20.2 and 6.4 EUR/MWh in these years amounts to about 2.0 and 0.6 million EUR a year on a contract for 100 GWh of annual output. In 2022, however, when the gas crisis raised summer and midday prices, the adjustment underpriced solar output by 12.5 points, against 2.7 for HIST5; as prices were then highest, the fall in the absolute error in EUR/MWh, from 12.5 to 9.2, is not significant, with $p = 0.39$. Most of the gain comes from the level of midday prices relative to baseload: the seasonal part of the adjustment alone lowers the bias only to 9.3 points. The solar channel carries the adjustment (2.1 points with solar only), and neither an ex ante selection of channels nor a recursive calibration of the size of the update (Proposition 4(c)) improves on TDS-EU (Online Supplement).

Three-year contracts confirm the result: the adjustment cut the mean overvaluation by 7.3 points, to 5.1, and the absolute error from 12.6 to 6.9 points, with $p = 0.06$ under both the main test and wild bootstrap, against 8.6 for TREND, with $p = 0.15$. For the six-year contract signed at the end of 2019, the adjustment with penetration extrapolated from 2017-2019 lowers the average overvaluation only from 20.9 to 17.0 points, no better than FLAT, which yields 16.7

points, or TREND, which yields 16.3 points, because linear extrapolation missed the post-2020 solar build-out; with the realised European penetration, it falls to 7.7 points: at long horizons, the TDS model is worth what the penetration forecast is worth. Fixed profiles also ignore the covariance of daily output and prices: for Germany, for which Bundesnetzagentur [57] provide actual quarter-hourly output, the capture rate of actual output was 2.8 points below the profile capture rate on average, with a gap of 5.0 points in 2025, mostly because sunny days are cheap, which raises every rule's overvaluation alike (Online Supplement).

The baseline PPA exercise defines its hourly output profile using the 2015-2020 solar generation data, even for the earliest forecast origins. To check whether this affects the results, we estimate the within-day profiles recursively from observations available by each origin and recompute both the hourly TDS price-shape adjustment and the forecast and realised profile capture rates. With a contemporaneously available Slovak profile used for three early Hungarian observations lacking sufficient domestic history, the one-year results remain virtually unchanged: HIST5 overvalues solar output by 11.15 points of the annual baseload price and TDS-EU by 2.77 points; their respective absolute errors are 13.01 and 7.67 points. The difference in mean bias remains significant, with $p = 0.009$, whereas the reduction in absolute error relative to HIST5 is marginal, with $p = 0.089$; TDS-EU lowers absolute error relative to TREND, with $p = 0.025$. Excluding those three observations instead of borrowing a neighbouring profile yields the same conclusions (Online Supplement). As in the baseline exercise, annual baseload levels and delivery-year monthly solar production shares remain fixed at their realised values: the comparison isolates errors in the hourly price shape rather than representing an entirely ex ante PPA quotation.

In practice, hourly forward curves are calibrated to traded baseload futures, and a prior shapes only the hours within each month and the months within a traded quarter. For Germany, France and Italy, using 24 contracts for delivery in 2018-2025, we take the monthly levels from the EEX futures at the signing date and the within-month hourly shape from each rule (Online Supplement). At the futures without a within-month shape, solar output is overvalued by 8.4 points, and with the within-month shape of HIST5 by 6.4 points, rising to 14.9 points in 2023-2025; neither estimate is significantly different from zero, with respective p-values of 0.10 and 0.11. The TDS shape moves valuations down by 7 to 8 points in each market, to a mean error of -1.1 points, with an error of 1.2 points in 2023-2025 and an overcorrection of -5.4 points in Italy, and lowers the absolute error from 9.9 to 7.1 points, an insignificant gain with $p = 0.19$ that trend extrapolation and last year's shape also achieve here. Baseload futures price months, not hours, so the part of the adjustment that deepens the midday discount survives calibration to the market. In EUR/MWh, the error of the futures level, averaging 38% of the realised annual price in absolute value, dominates every rule's error.

7.2. Schedules and the choice between priors

The run-of-river plant and the reservoir share a stylised snowmelt inflow profile, identical across markets so that differences between markets and priors come only from price shapes. The reservoir stores half of annual inflows ($\bar{S} = 6$) with a turbine limit of 2.5 times mean inflow, and its value is the premium over run-of-river; the store charges or discharges at most a third of its capacity per month ($\bar{c} = 1/3$), with $\eta = 0.75$, and its value is per unit of capacity, both in

Table 5: Decision consequences of seasonal priors (percent of the annual baseload price).

	Linear valuation bias		Intrinsic − PF value		Operational regret			
					Reservoir		Store	
Prior	Solar PPA	Run-of-river	Reservoir	Store	Static	Rolling	Static	Rolling
A. One year ahead, targets 2018–2025 (200 market-years)								
FLAT	3.97	6.20	−31.44	−43.07	31.44	24.25	43.07	53.46
HIST1	0.30	0.83	−1.21	−1.59	21.21	24.28	55.92	61.38
HIST5	0.45	0.90	−11.84	−28.08	16.87	20.68	43.44	45.14
SIN5	1.01	1.67	−14.55	−31.76	15.94	18.83	42.35	41.70
SIN5c	1.31	2.34	−17.34	−36.46	15.83	19.35	40.62	43.55
TDS-EU	0.27	0.94	−13.35	−29.59	14.84	18.42	40.19	41.12
TDS-EUc	0.57	1.60	−15.99	−34.35	15.03	18.96	39.77	42.84
TDS-EU[‡]	−0.05	0.58	−12.83	−28.82	15.21	19.40	39.65	41.85
Ex post mean[§]			−11.64	−29.79				
p: SIN5 − HIST5	0.00	0.02	0.01	0.01	0.46	0.11	0.70	0.36
p: TDS-EU − SIN5	0.12	0.17	0.12	0.06	0.14	0.56	0.09	0.35
PF value			31.44	43.07				
B. Origin 2019, targets 2023–2025 (71 market-years)								
FLAT	6.29	8.31	−29.36	−34.93	29.36	15.57	34.93	37.79
HIST1	3.93	5.11	−13.19	−24.46	14.81	12.91	38.12	33.06
HIST5	3.89	4.78	−17.05	−31.11	14.38	14.14	32.67	34.33
SIN5	4.16	4.93	−18.22	−32.46	15.78	15.75	32.13	34.74
TDS-EU	4.95	5.94	−18.98	−32.31	17.04	16.42	34.56	34.75
TDS-EU[‡]	2.58	3.17	−15.41	−29.97	14.51	14.87	30.16	33.61
PF value			29.36	34.93				

Linear valuation bias: forecast minus realised value of a pay-as-produced solar contract at monthly baseload prices (181 and 68 market-years with solar output) and of a run-of-river plant. Intrinsic − PF value: value of the optimal schedule under the prior minus the perfect-foresight value. Operational regret: perfect-foresight value minus the realised value of the static schedule optimised on the prior, or of the rolling-intrinsic policy (monthly re-optimisation with the current month's price known and later months forecast by the prior rescaled to the observed price level). FLAT: run-of-river operation and no storage trades (static), half-full initial storage (rolling). [‡]Realised penetration. [§]Intrinsic value of each market's average realised shape over the target years minus its average perfect-foresight value, averaged over markets. p-values two-way clustered by market and target year.

percent of the annual baseload price. Table 5 reports, for each prior, the bias of linear valuations at monthly resolution, the gap between the intrinsic value $V(\hat{s})$ and the perfect-foresight value $V(s)$, and the operational regret of static schedules and of a rolling-intrinsic policy; the Online Supplement reports further outcomes for all priors of Table 3.

The switch from HIST5 to SIN5 illustrates Proposition 3. It changes the intrinsic values of the reservoir and the store by −2.7 and −3.7 points ($p = 0.006$ and 0.007) but not the regret of static schedules (−0.9 and −1.1 points, $p = 0.46$ and 0.70). The realised values of static schedules vary little across the averaged and smoothed priors of Table 5: the perfect-foresight premium is 31.4 points, of which reservoir schedules capture 46% under HIST5 and 53% under TDS-EU, against 31-42% under HIST1, EWMA and the drift rules. Their static store schedules trade in 74-88% of market-years but earn little at realised prices, ranging from −0.4 to 3.4% of the annual baseload price compared with a perfect-foresight value of 43%, and lose money in 38-49% of market-years: with a round-trip efficiency of 0.75, a mistimed cycle loses money.

Table 6 asks when replacing a prior changes the schedule and whether the change pays. Replacing SIN5 by its transition-adjusted version leaves the reservoir schedule unchanged in 35.5% of market-years and the store schedule in 44.5%: the correction stays within the tolerance of the

Table 6: When does updating the prior change the schedule? Static schedules, one year ahead (200 market-years).

		Schedule changes (%)						
Update	Unchanged (%)	Gain ≥ 1	$\lvert\text{Gain}\rvert < 1$	Loss ≥ 1	U_1	Gain G	p	$\hat{\lambda}$
Reservoir, one year ahead, 2018–2025								
HIST5 → SIN5	4.0	47.5	26.5	22.0	1.14	0.93	0.46 [0.46]	0.92 (0.21)
SIN5 → TDS-EU	35.5	29.0	24.5	11.0	0.22	1.10	0.14 [0.14]	1.72 (0.50)
SIN5 → TDS-EU‡	32.5	28.0	27.5	12.0	0.33	0.73	0.18 [0.20]	1.37 (0.42)
SIN5 → TDS-EU, 2023–2025	23.8	46.3	18.8	11.3	0.43	1.31		
Store, one year ahead, 2018–2025								
HIST5 → SIN5	15.5	39.0	6.5	39.0	2.28	1.10	0.70 [0.74]	0.77 (0.23)
SIN5 → TDS-EU	44.5	32.0	8.5	15.0	0.54	2.15	0.09 [0.11]	2.08 (0.20)
SIN5 → TDS-EU‡	42.5	32.0	9.5	16.0	0.79	2.70	0.11 [0.13]	1.85 (0.27)
SIN5 → TDS-EU, 2023–2025	23.8	48.8	12.5	15.0	1.12	4.41		

Update from the incumbent to the candidate prior (Proposition 4). Unchanged: the incumbent schedule remains optimal under the candidate ($U_1 = 0$). Schedule changes: shares of market-years in which the schedule changes with a realised gain G of at least one point, of less than one point in absolute value, or a loss of at least one point (rows for 2023-2025 may not sum to 100 because of rounding). U_1: value of the update claimed by the candidate; G: realised gain, both averaged over all market-years, in percent of the annual baseload price. p: for $G = 0$, two-way clustered [wild bootstrap by year]. $\hat{\lambda}$: slope of $G + U_0$ on K through the origin over market-years with changed schedules, two-way clustered standard error in parentheses. ‡Realised penetration.

incumbent schedule. Where the schedule changes, fractions of 0.41 and 0.35 of the correction would already have changed the reservoir and store schedules, respectively; these are the median values of t^*, and the change pays in nearly two thirds of cases, specifically 64% for the reservoir and 65% for the store; the mean gain, 1.1 and 2.2 points, is below the minimum detectable effects of 2.1 and 3.6 points, with respective p-values of 0.14 and 0.09. The value claimed for the update ex ante is small, 0.2 and 0.5 points on average, but informative: for the store, realised gains rise from 2.1 points, with 57% positive outcomes, when $0 < U_1 \leq 0.5$, to 5.0 points, with 74% positive outcomes, when $0.5 < U_1 \leq 2$, and to 10.4 points, with 85% positive outcomes, above that range (Online Supplement). The estimate of λ is 1.72 with a standard error of 0.50 for the reservoir and 2.08 with a standard error of 0.20 for the store; the latter is above one with $p = 0.001$ under the main test and $p = 0.07$ under the wild bootstrap: in the direction that matters for schedules, the transition-adjusted prior corrects too little, as it does for shapes. If λ is common across market-years, the point estimates imply that updating paid in expectation whenever it changed the schedule, since the break-even ratio U_0/K is below one for every change and averages 0.49, although for the reservoir the confidence interval for λ, from 0.5 to 2.9, includes values below the larger break-even ratios; a rule that updates only when $\hat{\lambda}K > U_0$, with $\hat{\lambda}$ estimated on earlier target years, earns less than always updating (Online Supplement). Smoothing behaves differently: replacing HIST5 by SIN5 changes most schedules, affecting 96.0% and 84.5% of reservoir and store market-years, respectively, with $\hat{\lambda}$ close to one at 0.92 and 0.77. The value of updating grew with the transition: in 2023-2025 the TDS adjustment changed 76% of schedules and raised the store's realised value by 4.4 points on average. A long-horizon update can backfire: at the 2019 origin, with extrapolated penetration, it lowered realised values in 2023-2025 by 1.3 and 2.4 points; with realised penetration it raised them by 1.3 and 2.0 points.

A rolling-intrinsic operator re-optimises at the start of every month, following Lai et al. [36], Wu et al. [37] and Secomandi [38], with the current month's price known, an idealisation of the month-ahead forward price, later months forecast by the prior anchored to the price

level observed so far, and a return to the initial storage level by the end of the year; without month-ahead information this policy reproduces the static schedule. With it, operations do not systematically improve: for every averaged or smoothed prior the reservoir's regret rises by 2.9 to 4.2 points, from 14.8 to 18.4 for TDS-EU, and the store's changes by between −0.6 and +3.1 points. The result hinges on the level forecast (Online Supplement): anchoring later months to last year's annual price, a naive annual forward, does worse, producing regrets of 22.7 and 46.1 points for TDS-EU, whereas anchoring them to the realised annual price, which only a perfect annual forward would supply, cuts the regrets to 11.0 and 19.2 points. Adaptive operation pays when the annual level is well forecast, and the choice between SIN5 and TDS-EU then matters little: under each anchor, their regrets differ by at most 0.6 points, with $p > 0.18$. Static schedules optimised on year-end futures curves for Germany, France and Italy had higher regrets than those on the priors over 2018-2025, insignificantly and mainly because of 2021, and lower regrets in 2023-2025, mostly in 2025; re-optimised monthly on the futures curve, without any realised price, they did about as well as the rolling policy on the priors (Online Supplement).

Turning from schedules to the value of seasonal flexibility, intrinsic valuations on the averaged and smoothed priors of Table 5 fall short of the perfect-foresight value by 12-17 points for the reservoir and 28-36 points for the store, 38-55% and 65-85% of it. Most of the gap is the difference between the perfect-foresight value and the intrinsic value of each market's ex post mean shape (−11.6 and −29.8 points), two thirds or more of each prior's gap: the expected value of perfect information of Proposition 3(c), which no correction of the prior's mean can remove. Last year's shape, as dispersed as a realised shape, gives nearly unbiased intrinsic valuations, with errors of −1.2 and −1.6 points, but they exceed what its schedules realise by 20.0 and 54.3 points, against 1.4-5.0 and 10.6-15.4 for HIST5, SIN5 and TDS-EU: the optimizer's curse. The perfect-foresight value has nevertheless risen sharply: averaged over the 24 markets with three years of data in 2017-2019, the reservoir premium rose from 17.3% of the annual baseload price in 2017-2019 to 29.5% in 2023-2025 and the store value from 13.5% to 35.0% (also on fuel-adjusted shapes), and across markets and years it rises with wind rather than solar penetration. Attainable values rose by fewer points: the static TDS-EU reservoir premium by 5.5 points between 2018-2019 and 2023-2025, against 12.3 points under perfect foresight, while the static store lost money in both periods (Online Supplement).

8. Conclusions

The paper asked when updating a seasonal price prior for the energy transition changes a decision. On prices, the evidence is consistent: summer prices fell relative to winter prices in every market observed before and after the energy crisis, mostly at midday and also net of fuel-cost paths. The pooled permutation test gives $p = 0.018$, increasing to $p = 0.071$ after adjusting for fuel paths, and individual tests are significant in seven of 22 markets. The seasonal shapes of fuel costs did not change, and among seventeen other commodities only that of maize changed significantly. A first-order merit-order model, in which seasonal Fourier coefficients are affine in penetration hour by hour, organises the within-day solar evidence. Solar's effect on monthly baseload shapes cannot be separated from common European shocks, but within the day it is identified at the market level; the model's restriction that penetration-driven and weather-driven

output changes move prices alike is not rejected for solar and fails for wind.

On decisions, three conclusions follow. First, valuations of contracts with fixed profiles inherit prior errors one for one, and the relevant prior depends on the hours a contract covers. For solar PPAs valued with hourly forward curves, history-based shapes overvalued output by 22% of the baseload price in 2023-2025; the transition adjustment, computed ex ante, removed most of the mean overvaluation and two fifths or more of the absolute error of one- and three-year contracts, with p-values of 0.09 and 0.06, respectively, mainly through the level of midday prices; on curves calibrated to traded baseload futures, which do not price the hours of the day, it still lowered valuations by a similar amount, but without a significant accuracy gain. It underpriced output in 2022, however, and at long horizons is only as good as the penetration forecast. For monthly baseload shapes the gains were small and insignificant. Second, the realised values of energy-limited assets are robust to the choice among averaged or smoothed priors: their schedules stayed unchanged under 36-45% of the transition corrections, and where they changed, updating paid in nearly two thirds of cases, more when the claimed value was large, but without a mean gain significant at 5%. The tolerance and the claimed value can be computed before a prior is replaced, whereas the endorsed share must be estimated, and an estimated threshold rule did not beat always updating. Rolling re-optimisation with month-ahead prices did not systematically help unless later months were anchored to the realised annual level: for adaptive operation, the forecast of the annual level matters more than the seasonal prior. Third, the perfect-foresight value of seasonal flexibility rose by 70-160% between 2017-2019 and 2023-2025, but it is an upper bound: intrinsic valuations on averaged or smoothed priors understate it by 38-85%, mostly because of the value of perfect information, and attainable values rose by fewer points.

The analysis has limitations. Hourly solar output enters only through fixed within-day profiles from 2015-2020, which ignore the day-to-day covariance of output and prices; in Germany, this covariance lowered actual-output capture rates by a further 2.4 points on average; the Nordic price series are averages of zonal prices; and priors are evaluated against realised spot prices with the annual level fixed, except for a futures benchmark in three markets. The model is first order, and its wind restriction fails. The out-of-sample period contains eight target years dominated by an energy crisis, and the decision analysis uses stylised assets and simple policies with an idealised month-ahead price. Natural extensions include penetration scenarios for long-dated PPAs, joint calibration to futures and options in the spirit of Fanelli and Schmeck [21] and Frau and Fanelli [10], structural identification of coupling effects from interconnector expansions, and stochastic dynamic valuation of seasonal flexibility under transition-dependent price dynamics.

Declarations

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Competing interests

The author declares that she has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

All data except the EEX futures prices are public: Ember price and generation data and World Bank prices (CC BY 4.0), FRED exchange rates, Open Power System Data zonal prices and solar output, SMARD solar output for Germany (Bundesnetzagentur | SMARD.de, CC BY 4.0), and EEX auction results, which are available from EEX. The EEX futures settlement prices for Germany, France and Italy are proprietary (LSEG Datastream, EEX).

Author contributions

Viviana Fanelli: Conceptualization, Methodology, Formal analysis, Investigation, Writing - original draft, Writing - review and editing.

Appendix A. Proofs

Proof of Proposition 1. By Assumption 1 and (1), $R_{i,y,m} = \bar{R}_{i,y} + \rho_{i,y,m}$ with $\rho_{i,y,m} = \delta_{i,y,m} - \sum_j \omega_j z^j_{i,y} \kappa^j_{i,y,m}$ and $\sum_m \rho_{i,y,m} = 0$. A first-order expansion gives $\Psi_i(R_{i,y,m}) = \Psi_i(\bar{R}_{i,y}) + \theta_D \rho_{i,y,m} + r_{i,y,m}$, with a remainder $r = O(\rho^2)$ that depends on penetration through $\bar{R}$ and ρ. Subtracting the within-year mean of $\ln P_{i,y,m}$ removes $\mu_{i,y}$ and $\Psi_i(\bar{R}_{i,y})$ and leaves $q_{i,y,m} = \theta_D \rho_{i,y,m} + \tilde{g}_{i,y,m} + \tilde{\varepsilon}_{i,y,m} + \tilde{r}_{i,y,m}$, where tildes denote within-year deviations. The conditional mean of $\tilde{\varepsilon}_{i,y,m}$ is the within-year deviation of $\bar{\varepsilon}_{i,m}$; collecting it and the stable market-specific part of $\tilde{r}$ in $\alpha_{i,m}$ and the remainder in $\xi_{i,y,m}$, and writing $\theta_j = \omega_j \theta_D$, gives (2); treating $\tilde{r}$ as part of α and ξ is the first-order approximation. Part (a) follows by differentiation with profiles and θ_D held fixed. For (b), take expectations of (2), substitute the Fourier representations, which are exact for twelve-month vectors with zero mean, and use linearity: $C_k(z) = C^0_k - \sum_j \theta_j z^j c^j_k$, where C^0_k collects the coefficients of α, $\theta_D \delta$ and $\mathbb{E}[\tilde{g}]$. For (c), $\partial \|C_k\|^2 / \partial z^j = 2 C_k' \partial C_k / \partial z^j = -2\theta_j C_k' c^j_k$. For non-zero vectors, $C_k' c^j_k = \|C_k\| \|c^j_k\| \cos\psi$, where ψ is the phase difference between the two harmonics, so the derivative is positive if and only if $\cos\psi < 0$, i.e. the phases differ by more than a quarter of the harmonic's period. □

Proof of Proposition 2. Under the assumptions, $q_{y,m} = \alpha_m + \theta_D \delta_m - \sum_j \theta_j z^j_y \kappa^j_m + \zeta_{y,m}$ with $\mathbb{E}[\zeta_{y,m}] = 0$, where ζ collects $\tilde{g}$ and ξ. Then $\hat{q}^H_{Y,m} - q_{Y,m} = \sum_j \theta_j (z^j_Y - \bar{z}^j_W) \kappa^j_m + \bar{\zeta}_{W,m} - \zeta_{Y,m}$, where $\bar{z}^j_W$ and $\bar{\zeta}_{W,m}$ are window averages. Linear growth gives $z^j_Y - \bar{z}^j_W = \dot{z}_j (Y - \bar{y}_W)$ with $\bar{y}_W = Y - h - \frac{n_W - 1}{2}$, hence (3). For the adjusted rule, stable profiles imply $\bar{\kappa}^j_m = \kappa^j_m$ and $\bar{x}^j_{W,m} = \bar{z}^j_W \kappa^j_m$, so $\hat{q}^T_{Y,m} - q_{Y,m} = \sum_j [\theta_j (z^j_Y - \bar{z}^j_W) - \hat{\theta}_j (\hat{z}^j_Y - \bar{z}^j_W)] \kappa^j_m + \bar{\zeta}_{W,m} - \zeta_{Y,m}$, whose expectation is $\sum_j [\theta_j (z^j_Y - \hat{z}^j_Y) + (\theta_j - \hat{\theta}_j)(\hat{z}^j_Y - \bar{z}^j_W)] \kappa^j_m$. With $\hat{z}^j_Y = z^j_Y$ and one technology, this is $(\theta_j - \hat{\theta}_j)/\theta_j$ times the historical bias. The empirical rules average shape factors rather than log shapes, which is equivalent to first order. □

Proof of Proposition 3. (a) By optimality of $\hat{x}$ under $\hat{s}$, $V(\hat{s}) = \hat{s}'\hat{x} \geq \hat{s}' x^* = V(s) + (\hat{s} - s)' x^*$. By optimality of x^* under s, $V(s) \geq s'\hat{x} = V(\hat{s}) - (\hat{s} - s)'\hat{x}$. If $X = \{w\}$, both bounds equal $w'(\hat{s} - s)$. (b) $V(s) - s'\hat{x} \geq 0$ by optimality of x^*. Since $\hat{s}'\hat{x} \geq \hat{s}' x^*$, $V(s) - s'\hat{x} = s' x^* - s'\hat{x} \leq s' x^* - s'\hat{x} + \hat{s}'\hat{x} - \hat{s}' x^* = (\hat{s} - s)'(\hat{x} - x^*)$. For the second statement, a unique optimum of a linear function over a polytope is an extreme point; let $x \neq x^*$ be another extreme point of X, which exists by assumption. Then $\hat{s}'(x^* - x) = s'(x^* - x) + (\hat{s} - s)'(x^* - x) \geq \Delta - \|\hat{s} - s\|_\infty \|x^* - x\|_1 > \Delta - (\Delta / \bar{d})\, \bar{d} = 0$. Every point of X other than x^* is a convex combination of extreme points that puts positive weight on some extreme point other than x^*, so $\hat{s}' x < \hat{s}' x^*$ for all $x \in X \setminus \{x^*\}$.

Hence $\hat{x} = x^*$, the regret is zero, and by (a) $V(\hat{s}) - V(s) = (\hat{s} - s)'x^*$. (c) V is the support function of X and hence convex and finite, and $\mathbb{E}\|s\| < \infty$ makes $\mathbb{E}[V(s)]$ and $\bar{s}$ well defined, so $\mathbb{E}[V(s)] \geq V(\bar{s})$ by Jensen's inequality. Let $\bar{x}$ be optimal at $\bar{s}$; then $\mathbb{E}[V(s)] - V(\bar{s}) = \mathbb{E}[V(s) - s'\bar{x}]$, and $V(s) - s'\bar{x} \geq 0$, so equality holds if and only if $\bar{x}$ is optimal for almost every s. Conversely, if some x is optimal for almost every s, then $\mathbb{E}[V(s)] = \bar{s}'x \leq V(\bar{s})$, so equality holds. □

Proof of Proposition 4. $U_1 \geq 0$ and $U_0 \geq 0$ by the optimality of $\hat{x}_1$ under $\hat{s}_1$ and of $\hat{x}_0$ under $\hat{s}_0$, and $K = \hat{s}_1'\Delta x - \hat{s}_0'\Delta x = U_1 + U_0$. (a) The set $N = \{p : p'\hat{x}_0 = V(p)\}$ of price vectors under which $\hat{x}_0$ is optimal, the normal cone of X at $\hat{x}_0$, is closed and convex: for $p, q \in N$ and $\alpha, \beta \geq 0$, $(\alpha p + \beta q)'\hat{x}_0 \leq V(\alpha p + \beta q) \leq \alpha V(p) + \beta V(q) = (\alpha p + \beta q)'\hat{x}_0$, and N is closed because V is continuous. The set of $t \in [0, 1]$ with $\hat{s}_0 + t\Delta s \in N$ is therefore a closed interval, which contains 0. The function u is convex and piecewise linear because V is the maximum of finitely many linear functions (the values at the extreme points of X) and the subtracted term is affine in t; $u \geq 0$, with equality exactly when $\hat{s}_0 + t\Delta s \in N$. For $0 \leq t_1 < t_2 \leq 1$, convexity and $u(0) = 0$ give $u(t_1) \leq (t_1/t_2)u(t_2) \leq u(t_2)$. By the optimality of $\hat{x}_1$, $u(1) = \hat{s}_1'\hat{x}_1 - \hat{s}_1'\hat{x}_0 = U_1$. If $t^* = 1$, $\hat{x}_0$ is optimal under $\hat{s}_1$, so $\hat{x}_1 = \hat{x}_0$ by convention, $U_1 = K = 0$, and $V(\hat{s}_1) - V(\hat{s}_0) = \Delta s'\hat{x}_0$. If $t^* < 1$, $u(1) > 0$, so $U_1 > 0$, $\hat{x}_1 \neq \hat{x}_0$ and $K \geq U_1 > 0$. (b) $G = \hat{s}_1'\Delta x - (\hat{s}_1 - s)'\Delta x = \hat{s}_0'\Delta x - (\hat{s}_0 - s)'\Delta x$. (c) G is integrable because X is bounded and $\mathbb{E}\|s\| < \infty$, and Δx is bounded and determined by $\mathcal{I}$, so $\mathbb{E}[G \,|\, \mathcal{I}] = \mu'\Delta x = \hat{s}_0'\Delta x + (\mu - \hat{s}_0)'\Delta x = -U_0 + \lambda K$ almost surely. If $\mu = \hat{s}_0 + \ell\Delta s$, then $(\mu - \hat{s}_0)'\Delta x = \ell K$. Hence $\mathbb{E}[G \,|\, \mathcal{I}] > 0$ if and only if $\lambda > U_0/K$. For the bounds, $U_0/K = 1 - U_1/K < 1$ because $U_1 > 0$; and since $\hat{x}_1 \in X$, $u(t) \geq (\hat{s}_0 + t\Delta s)'\Delta x = tK - U_0$, which is positive for $t > U_0/K$, so $t^* \leq U_0/K$. Finally, $\lambda = 1$ when $(\mu - \hat{s}_1)'\Delta x = 0$, and $1 > U_0/K$. □

References

[1] Stein-Erik Fleten and Jacob Lemming. Constructing forward price curves in electricity markets. *Energy Economics*, 25(5):409–424, 2003.

[2] Fred Espen Benth, Steen Koekebakker, and Frode Ollmar. Extracting and applying smooth forward curves from average-based commodity contracts with seasonal variation. *The Journal of Derivatives*, 15(1):52–66, 2007. doi: 10.3905/jod.2007.694791.

[3] Ruggero Caldana, Gianluca Fusai, and Andrea Roncoroni. Electricity forward curves with thin granularity: Theory and empirical evidence in the hourly EPEXspot market. *European Journal of Operational Research*, 261(2):715–734, 2017. doi: 10.1016/j.ejor.2017.02.016.

[4] Rüdiger Kiesel, Florentina Paraschiv, and Audun Sætherø. On the construction of hourly price forward curves for electricity prices. *Computational Management Science*, 16(1):345–369, 2019. doi: 10.1007/s10287-018-0300-6.

[5] Lion Hirth. The market value of variable renewables: The effect of solar wind power variability on their relative price. *Energy Economics*, 38:218–236, 2013. doi: 10.1016/j.eneco.2013.02.004.

[6] Roozbeh Qorbanian, Nils Löhndorf, and David Wozabal. Valuation of power purchase agreements for corporate renewable energy procurement. *European Journal of Operational Research*, 326(3):530–543, 2025. doi: 10.1016/j.ejor.2025.05.054.

[7] Nils Löhndorf, David Wozabal, and Stefan Minner. Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming. *Operations Research*, 61 (4):810–823, 2013.

[8] Nils Löhndorf and David Wozabal. Gas storage valuation in incomplete markets. *European Journal of Operational Research*, 288(1):318–330, 2021. doi: 10.1016/j.ejor.2020.05.044.

[9] Julio J. Lucia and Eduardo S. Schwartz. Electricity prices and power derivatives: Evidence from the Nordic Power Exchange. *Review of Derivatives Research*, 5(1):5–50, 2002. doi: 10.1023/A:1013846631785.

[10] Carme Frau and Viviana Fanelli. Seasonality in commodity prices: New approaches for pricing plain vanilla options. *Annals of Operations Research*, 336(1–2):1089–1131, 2024. doi: 10.1007/s10479-022-05128-x.

[11] Richard E. Wendell. The tolerance approach to sensitivity analysis in linear programming. *Management Science*, 31(5):564–578, 1985.

[12] Adam N. Elmachtoub and Paul Grigas. Smart "predict, then optimize". *Management Science*, 68(1):9–26, 2022.

[13] Manuel Moreno, Alfonso Novales, and Federico Platania. Long-term swings and seasonality in energy markets. *European Journal of Operational Research*, 279(3):1011–1023, 2019. doi: 10.1016/j.ejor.2019.05.042.

[14] Carsten Sørensen. Modeling seasonality in agricultural commodity futures. *Journal of Futures Markets*, 22(5):393–426, 2002.

[15] Svetlana Borovkova and Hélyette Geman. Seasonal and stochastic effects in commodity forward curves. *Review of Derivatives Research*, 9(2):167–186, 2006. doi: 10.1007/s11147-0 07-9008-4.

[16] Fred Espen Benth, Jūratė Šaltytė Benth, and Steen Koekebakker. *Stochastic Modelling of Electricity and Related Markets*. World Scientific, Singapore, 2008.

[17] Joanna Janczura, Stefan Trück, Rafał Weron, and Rodney C. Wolff. Identifying spikes and seasonal components in electricity spot price data: A guide to robust modeling. *Energy Economics*, 38:96–110, 2013. doi: 10.1016/j.eneco.2013.03.013.

[18] Jakub Nowotarski and Rafał Weron. On the importance of the long-term seasonal component in day-ahead electricity price forecasting. *Energy Economics*, 57:228–235, 2016.

[19] Janis Back, Marcel Prokopczuk, and Markus Rudolf. Seasonality and the valuation of commodity options. *Journal of Banking & Finance*, 37(2):273–290, 2013.

[20] Viviana Fanelli, Lucia Maddalena, and Silvana Musti. Modelling electricity futures prices using seasonal path-dependent volatility. *Applied Energy*, 173:92–102, 2016. doi: 10.1016/ j.apenergy.2016.04.003.

[21] Viviana Fanelli and Maren Diane Schmeck. On the seasonality in the implied volatility of electricity options. *Quantitative Finance*, 19(8):1321–1337, 2019. doi: 10.1080/14697688.2019.1582792.

[22] Lorenz Schneider and Bertrand Tavin. Seasonal volatility in agricultural markets: Modelling and empirical investigations. *Annals of Operations Research*, 334(1):7–58, 2024. doi: 10.1007/s10479-021-04241-7.

[23] Andrés García Mirantes, Javier Población, and Gregorio Serna. The stochastic seasonal behaviour of natural gas prices. *European Financial Management*, 18(3):410–443, 2012. doi: 10.1111/j.1468-036X.2009.00533.x.

[24] Francesco Lisi and Matteo M. Pelagatti. Component estimation for electricity market data: Deterministic or stochastic? *Energy Economics*, 74:13–37, 2018. doi: 10.1016/j.eneco.2018.05.027.

[25] Michele Bufalo and Viviana Fanelli. A seasonal two-factor model for solar energy production: A climate extreme events analysis. *Energy Economics*, 148:108611, 2025. doi: 10.1016/j.eneco.2025.108611.

[26] Martin T. Barlow. A diffusion model for electricity prices. *Mathematical Finance*, 12(4):287–298, 2002.

[27] Michael Coulon and Sam Howison. Stochastic behaviour of the electricity bid stack: From fundamental drivers to power prices. *Journal of Energy Markets*, 2(1):29–69, 2009.

[28] René Carmona, Michael Coulon, and Daniel Schwarz. Electricity price modeling and asset valuation: A multi-fuel structural approach. *Mathematics and Financial Economics*, 7(2):167–202, 2013.

[29] Rüdiger Kiesel and Michael Kusterman. Structural models for coupled electricity markets. *Journal of Commodity Markets*, 3(1):16–38, 2016.

[30] Peter Leoni, Pieter Segaert, Sven Serneels, and Tim Verdonck. Multivariate constrained robust M-regression for shaping forward curves in electricity markets. *Journal of Futures Markets*, 38(11):1391–1406, 2018. doi: 10.1002/fut.21958.

[31] Alan D. Lamont. Assessing the long-term system value of intermittent electric generation technologies. *Energy Economics*, 30(3):1208–1231, 2008.

[32] Javier López Prol, Karl W. Steininger, and David Zilberman. The cannibalization effect of wind and solar in the California wholesale electricity market. *Energy Economics*, 85:104552, 2020. doi: 10.1016/j.eneco.2019.104552.

[33] Seyed Danial Mohseni Taheri, Selvaprabu Nadarajah, and Alessio Trivella. Physical vs virtual corporate power purchase agreements: Meeting renewable targets amid demand and price uncertainty. *European Journal of Operational Research*, 320(1):256–270, 2025. doi: 10.1016/j.ejor.2024.08.002.

[34] Alexander Boogert and Cyriel de Jong. Gas storage valuation using a Monte Carlo method. *The Journal of Derivatives*, 15(3):81–98, 2008.

[35] Nicola Secomandi. Optimal commodity trading with a capacitated storage asset. *Management Science*, 56(3):449–467, 2010.

[36] Guoming Lai, François Margot, and Nicola Secomandi. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations Research*, 58(3):564–582, 2010.

[37] Owen Q. Wu, Derek D. Wang, and Zhenwei Qin. Seasonal energy storage operations with limited flexibility: The price-adjusted rolling intrinsic policy. *Manufacturing & Service Operations Management*, 14(3):455–471, 2012.

[38] Nicola Secomandi. Merchant commodity storage practice revisited. *Operations Research*, 63 (5):1131–1143, 2015.

[39] Dimitris Bertsimas and Nathan Kallus. From predictive to prescriptive analytics. *Management Science*, 66(3):1025–1044, 2020.

[40] Othman El Balghiti, Adam N. Elmachtoub, Paul Grigas, and Ambuj Tewari. Generalization bounds in the predict-then-optimize framework. *Mathematics of Operations Research*, 48 (4):2043–2065, 2023.

[41] James E. Smith and Robert L. Winkler. The optimizer's curse: Skepticism and postdecision surprise in decision analysis. *Management Science*, 52(3):311–322, 2006.

[42] Albert Madansky. Inequalities for stochastic linear programming problems. *Management Science*, 6(2):197–204, 1960.

[43] David B. Brown, James E. Smith, and Peng Sun. Information relaxations and duality in stochastic dynamic programs. *Operations Research*, 58(4):785–801, 2010.

[44] Frank Sensfuß, Mario Ragwitz, and Massimo Genoese. The merit-order effect: A detailed analysis of the price effect of renewable electricity generation on spot market prices in Germany. *Energy Policy*, 36(8):3076–3084, 2008. doi: 10.1016/j.enpol.2008.03.035.

[45] Florentina Paraschiv, David Erni, and Ralf Pietsch. The impact of renewable energies on EEX day-ahead electricity prices. *Energy Policy*, 73:196–210, 2014. doi: 10.1016/j.enpol.2014.05.004.

[46] David Wozabal, Christoph Graf, and David Hirschmann. The effect of intermittent renewables on the electricity price variance. *OR Spectrum*, 38(3):687–709, 2016. doi: 10.1007/s00291-015-0395-x.

[47] Benjamin Jansen, J. J. de Jong, Cornelis Roos, and Tamás Terlaky. Sensitivity analysis in linear programming: Just be careful! *European Journal of Operational Research*, 101(1): 15–28, 1997. doi: 10.1016/S0377-2217(96)00172-5.

[48] Saul Gass and Thomas Saaty. The computational algorithm for the parametric objective function. *Naval Research Logistics Quarterly*, 2(1–2):39–45, 1955. doi: 10.1002/nav.3800020106.

[49] John C. Driscoll and Aart C. Kraay. Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics*, 80(4):549–560, 1998. doi: 10.1162/003465398557825.

[50] A. Colin Cameron, Jonah B. Gelbach, and Douglas L. Miller. Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, 29(2):238–249, 2011. doi: 10.1198/jbes.2010.07136.

[51] Ember. European wholesale electricity price data [dataset]. https://ember-energy.org/data/european-wholesale-electricity-price-data, 2026. Hourly day-ahead prices from ENTSO-E, EMR and SEMOpx; licence CC BY 4.0; release of February 2026.

[52] Ember. Monthly electricity data [dataset]. https://ember-energy.org/data/monthly-electricity-data, 2026. European monthly generation, demand and emissions; licence CC BY 4.0; release of February 2026.

[53] World Bank. Commodity price data (The Pink Sheet), monthly prices [dataset]. https://www.worldbank.org/en/research/commodity-markets, 2026. Release of 3 February 2026; licence CC BY 4.0.

[54] Board of Governors of the Federal Reserve System. U.S. dollars to euro spot exchange rate, monthly [EXUSEU], retrieved from FRED, Federal Reserve Bank of St. Louis [dataset]. https://fred.stlouisfed.org/series/EXUSEU, 2026.

[55] European Energy Exchange. EU ETS primary auction results, Emission Spot Primary Market Auction Report 2012–2026 [dataset]. https://www.eex.com/en/market-data/market-data-hub/environmentals/eex-eua-primary-auction-spot-download, 2026.

[56] Open Power System Data. Data package time series, version 2020-10-06. https://doi.org/10.25832/time_series/2020-10-06, 2020. Primary data from ENTSO-E and national transmission system operators.

[57] Bundesnetzagentur. SMARD market data: Actual generation, photovoltaics, Germany, quarter-hourly resolution [dataset]. https://www.smard.de/en/downloadcenter/download-market-data, 2026. Bundesnetzagentur | SMARD.de; licence CC BY 4.0; data for 2015–2025.

[58] Fabio Canova and Bruce E. Hansen. Are seasonal patterns constant over time? A test for seasonal stability. *Journal of Business & Economic Statistics*, 13(3):237–252, 1995. doi: 10.1080/07350015.1995.10524598.

[59] Francis X. Diebold and Roberto S. Mariano. Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):253–263, 1995.

[60] A. Colin Cameron, Jonah B. Gelbach, and Douglas L. Miller. Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, 90(3):414–427, 2008.

[61] James G. MacKinnon, Morten Ørregaard Nielsen, and Matthew D. Webb. Wild bootstrap and asymptotic inference with multiway clustering. *Journal of Business & Economic Statistics*, 39(2):505–519, 2021.

[62] Matthew D. Webb. Reworking wild bootstrap-based inference for clustered errors. *Canadian Journal of Economics*, 56(3):839–858, 2023.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementES.pdf