Revisiting Seasonality in Electricity Prices "Work in Progress"

Joel Sarkisyan
Prof. Florentina Paraschiv
Prof. Fred Espen Benth

Zeppelin University
Energy Finance Christmas Workshop 2025

December 12, 2025

Outline

Motivation

- Electricity must be consumed on delivery and precisely balanced by grid operators due to technical properties of the grid
- Due to this balancing and the demand-driven nature, the seasonality of electricity prices strongly reflects consumer behavior
- Previously, the seasonality was believed to be defined purely by the load. However, a strong change in today's seasonality shape of electricity prices is observed

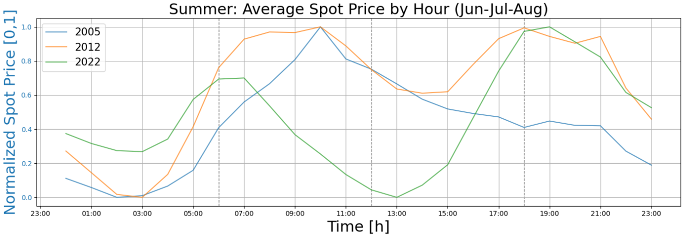

Figure: Hourly average PHELIX spot prices in summer of 2005, 2012, and 2022

Motivation

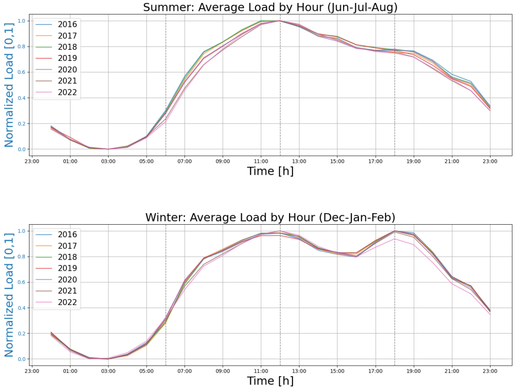

Figure: Average load for each hour of the day in summer and winter across the years 2016-2022

- Daily hour-averaged load profiles show that the shape of the load did not change significantly in the years 2016-2022.
- The characteristic load patterns still reflect the consumption behaviour of electricity in both summer and winter. This is because societal norms such as work hours, weekends, lunch/dinner times did not change over the years

Motivation

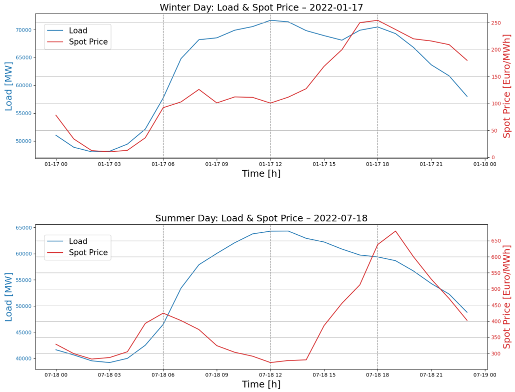

Figure: Hourly load and spot price comparison in winter and summer 2022

- The spot price seasonality is still defined by the load in very early and late sections of the day
- During the mid-day however, a vast discrepancy is observed which decouples the spot from the load
- This decoupling indicates that another source is responsible for the seasonality of electricity prices independent of the load

Seasonality in Electricity Price Forecasting (EPF): Literature Summary

How existing models represent the seasonal shape of electricity prices:

- **Calendar–Driven Seasonality**
  - Weekly seasonality via dummy variables (Mon/Sat/Sun, weekends, holidays)
  - Annual seasonality through seasonal dummies (Winter/Spring/Summer)
  - Intraday pattern handled implicitly through hour-specific models or AR lags
  - *Papers:* Nowotarski & Weron (2016), Maciejowska (2019), Sgarlato & Ziel (2023), Ghelasi & Ziel (2024)
- **Statistical Long-Term Seasonal Component (LTSC)**
  - LTSC extracted from prices using HP filters, wavelet smoothing, moving averages
  - Short-term seasonality captured by ARX/LASSO + weekly dummies
  - Entire seasonal shape is statistical, not physically derived
  - *Papers:* Nowotarski & Weron (2016), Chеć et al. (2025)

Seasonality in Electricity Price Forecasting (EPF): Literature Summary

- **Fundamental-Driven Seasonality**
- Seasonal variation enters through load, wind, and solar regressors
- Sometimes season-specific coefficients, but no continuous annual function
- Price seasonality remains indirect; fundamentals drive the shape
- *Papers:* Fuke & Ohashi (2025), Maciejowska (2019), de Marcos et al. (2019), Ghelasi & Ziel (2024)
- **Machine Learning (Implicit Seasonality)**
- Neural networks learn periodicity from calendar embeddings and historical patterns
- No explicit functional seasonality; purely implicit
- *Papers:* Wagner et al. (2022)

Problem Statement

Background.

- In both EPF and Price Forward Curve (PFC) construction, electricity prices are typically decomposed into a **deterministic seasonal component** and a **stochastic component**
- For PFCs in particular, the deterministic part must be a **smooth, continuous function** so that forward prices can be constructed and interpolated consistently

Problem.

- Rising solar penetration fundamentally changes the **intra-day shape** of spot prices
- Existing seasonality models in electricity pricing are mainly **load-driven** or rely on **categorical RES indicators**, which cannot capture the evolving intra-day shape caused by solar generation
- We aim at constructing a **more accurate deterministic seasonality** that improves the disentangling of deterministic and stochastic components in EPF and PFC construction

Data Overview

- Source: **ENTSO-E Transparency Platform**
  - Day-ahead spot prices (hourly)
  - Day-ahead load forecasts (hourly)
  - Day-ahead solar generation forecasts (15-min)
- Study region: **EPEX Germany (PHELIX)**
- Time span: **2016–2022**
- Preprocessing steps:
  - Solar forecasts aggregated to hourly resolution (average of four 15-min power values)
  - Harmonized timestamps and DST adjustments
  - Removed missing or negative forecast entries
  - Exclusively day-ahead information used (no realized data)

Methodological Framework

- Electricity spot prices decomposed into:

$$P(t) = S(t) + \varepsilon(t)$$

- Seasonal component modelled as:

$$S(t) = f_{\text{load}}(t) + f_{\text{solar}}(t)$$

- Strategy:
  - Extract baseline seasonal shape from the **load** (**demand-driven**)
  - Construct a **solar-driven** seasonal shape via a parametric PV model
  - Combine into a smooth, continuous seasonality function

Motivation: Solar Infeed is Highly Structured

- Daily PV infeed exhibits a **smooth, bell-shaped** intra-day pattern
- The pattern is **repeated day after day**, but
  - peak level changes,
  - width (day length) changes,
  - timing of the peak shifts slightly
- Suggests modelling PV as a **smooth wave** with **three main degrees of freedom**: amplitude, frequency, and phase

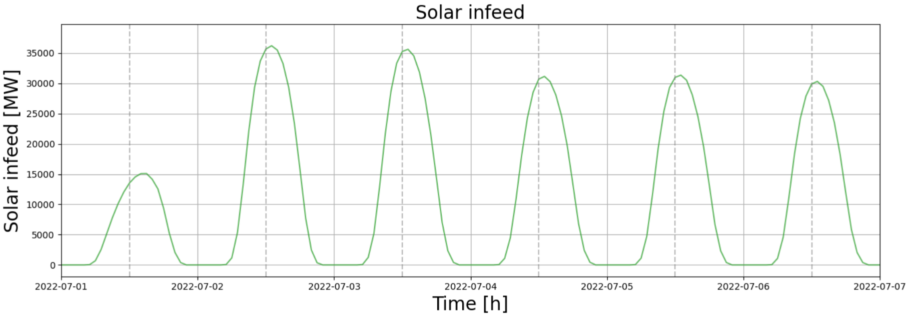

Motivation: Seasonal Contrast in PV Profiles

Summer 2022

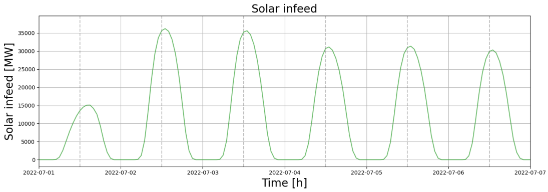

Winter 2022

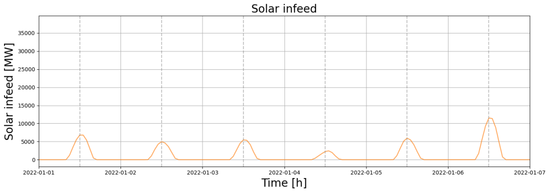

- Winter: lower peaks, much **narrower** generation window
- Summer: higher peaks, **wider** generation window
- Both are well-described by the **same functional form** with different:

$$\text{amplitude } A(t), \quad \text{frequency/width } T(t), \quad \text{phase } \phi(t)$$

Solar Seasonal Component: Truncated Cosine Model

- Daily PV generation profile approximated by:

$$PV(t) = \max\{0,\, A(t) \cdot \cos\left(\frac{2\pi(t - \phi(t))}{T(t)}\right)\}$$

- Components:
  - $T(t)$: daylength (frequency) — highly deterministic
  - $\phi(t)$: phase shift — determined by sunrise/sunset timing
  - $A(t)$: amplitude — stochastic, weather-driven, capacity-driven
- Empirical observation:
  - $T(t)$ and $\phi(t)$ follow a deterministic pattern across years
  - Only $A(t)$ introduces substantial variability

Decomposition of PV Amplitude $A(t)$

- Model for the log-amplitude:

$$\log A(t) = \underbrace{\alpha_0 + \alpha_1 t}_{\text{linear trend}} + \underbrace{b_1 \sin(2\pi t/365) + b_2 \cos(2\pi t/365)}_{\text{annual cycle}} + \underbrace{\text{AR}(1)}_{\text{short-term persistence}}$$

- Residual variance modelled by GARCH(1,1):

$$\sigma_t^2 = \omega + \beta \sigma_{t-1}^2 + \gamma \varepsilon_{t-1}^2$$

- Interpretation:
  - Trend: increasing PV capacity
  - Annual cycle: summer–winter irradiance differences
  - AR(1): short-term weather persistence
  - GARCH: winter volatility $>$ summer volatility (log-transformed amplitude)

Solar-Adjusted Seasonal Price Model

- Proposed structure:

$$P(t) = \widehat{L}(t) + f_{\text{solar}}(t) + \varepsilon(t)$$

- Interpretation:
  - Load determines the baseline intra-day and intra-week pattern
  - Solar modifies this pattern by reshaping the daytime curve
  - In the absence of solar infeed:

$$f_{\text{solar}}(t) = 0 \quad \Rightarrow \quad P(t) = \widehat{L}(t) + \varepsilon(t)$$

  $\rightarrow$ model collapses to classical load-driven seasonality

- Key novelty:
  - Seasonality is still load-driven but is **adjusted for Solar infeed**
  - Solar acts as a structural modifier of the seasonal shape
  - Accounts for dynamic structural seasonality changes within summer/winter

Preliminary Results: PV Shape Characteristics

- Daily PV generation profiles are well approximated by the truncated cosine form
- The **width** (daylight duration $T(t)$) varies deterministically between ~ 8 and ~ 16 hours
- The **phase shift** $\phi(t)$ tracks sunrise/sunset and is highly stable across years

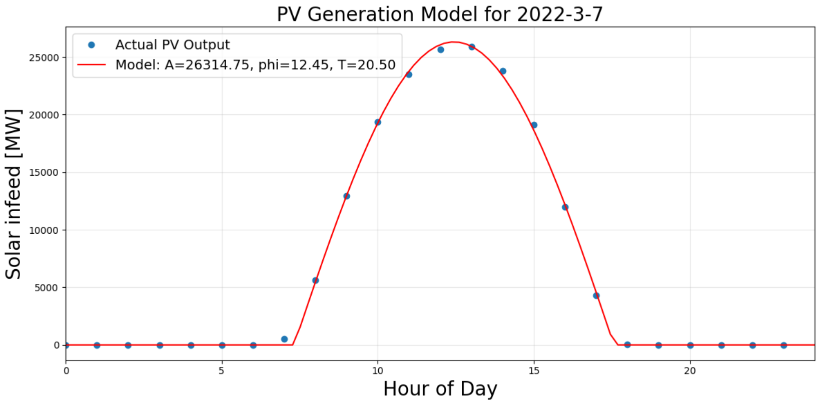

Preliminary Results: Estimated Frequency

- Estimated daily frequency $1/T(t)$ follows an annual cycle determined by daylight hours
- rounding the frequency to hours to match the hourly data reveals a reoccurring range of daylight hours for each day of year
- The frequency is a physical parameter that can be deterministically extracted for each day of the year

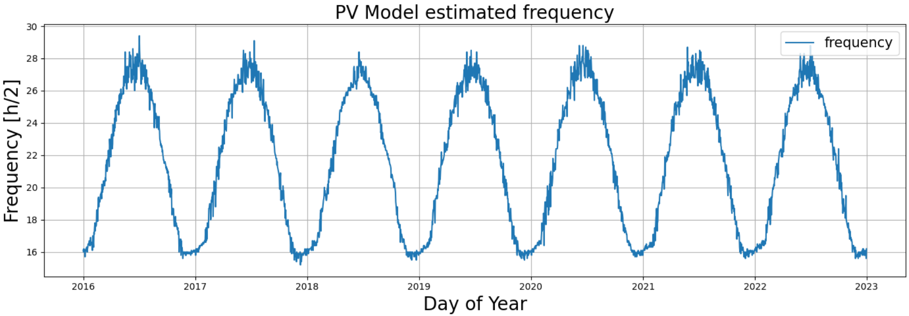

Preliminary Results: Estimated Phase Shift

- The phase shift governs the position of the maximum amplitude
- The variance of the phase within summer/winter is negligible in the context of hourly data points
- The jump of the phase captures the daylight saving time

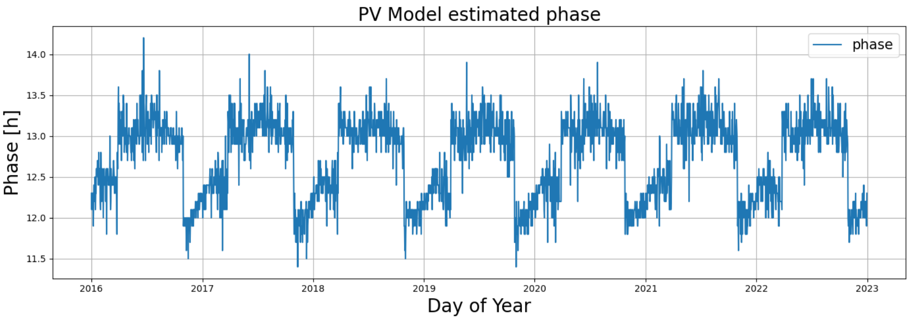

Preliminary Results: Amplitude Dynamics

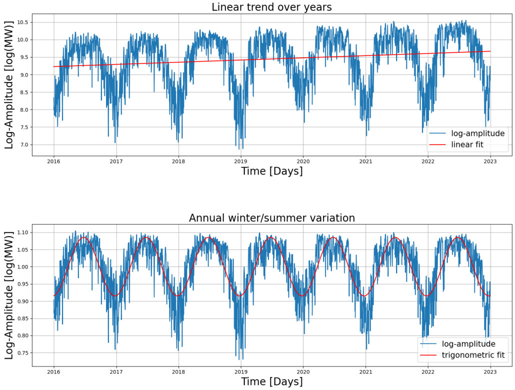

Figure: log-transformed amplitude with linear and trigonometric fit.

- The log-transformed amplitude exhibits a linear **upward trend** over years
- The amplitude displays a **sinusoidal annual pattern**, corresponding to summer–winter irradiance differences
- The seasonal decomposition indicates:
  - Summer: stable PV output
  - Winter: highest volatility and largest deviations

Preliminary Results: Residual Structure of $\log A(t)$

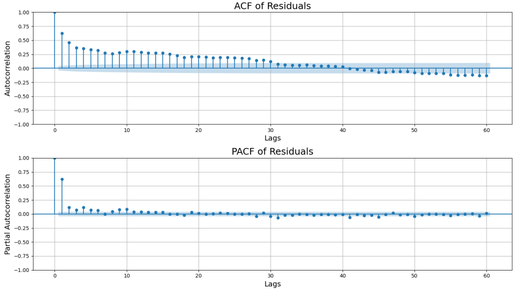

Figure: ACF and PACF of the residuals after linear and trigonometric fitting

- The trigonometric fit captures the summer/winter cycle but significant serial correlation remains
- An AR(1) component captures short-term dependence: yesterday's amplitude influences today's
- Residuals show strong **volatility clustering**, especially in winter $\rightarrow$ Indicates a GARCH (S-GARCH)

Preliminary Results: GARCH Structure of the Residuals

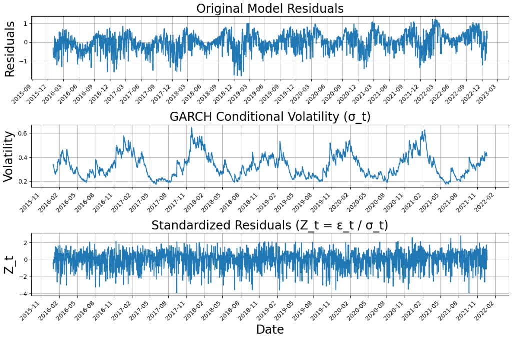

Figure: Residuals of the GARCH(1,1) fit

- AR(1)-GARCH(1,1) yields i.i.d. standardized residuals
- Captures both the mean dynamics and the heteroskedasticity of $A(t)$
- In conlcusion the amplitude has the following dynamics:
  - a positive **linear trend** over years accounting for growing PV capacity
  - a **sinusoidal movement** with yearly frequency between summer and winter
  - A **short-term memory** of weather persistence
  - **seasonal variation** predominantly between summer and winter

Intended Contribution

- Preliminary evidence suggests that load-based seasonality alone may not **fully reflect** recent **intra-day price shapes**
- Solar infeed appears to introduce **systematic daytime effects** that could be represented through a smooth functional component
- The truncated–cosine formulation provides a **candidate structure** for modeling PV-related seasonal patterns, with:
  - largely deterministic day length and phase behavior,
  - amplitude dynamics driven by trend, annual cycles, and short-term persistence
- Combining load-based seasonality with a solar-driven adjustment may offer a more **flexible continuous seasonal component**
- The approach is intended to support improved separation of **deterministic** and **stochastic** components in EPF and PFC settings

Thank you for your attention!