The Role of Macroeconomic Low Frequency Variables in Electricity Price Forecasting: a Reverse-MIDAS Approach

Lorenzo Rossi[1,2]

[1] *Department of Economics, Management, and Quantitative Methods, University of Milan, Milan, Italy*

[2] *Department of Economics and Management, University of Pavia, Pavia, Italy*

August 27, 2026

Abstract

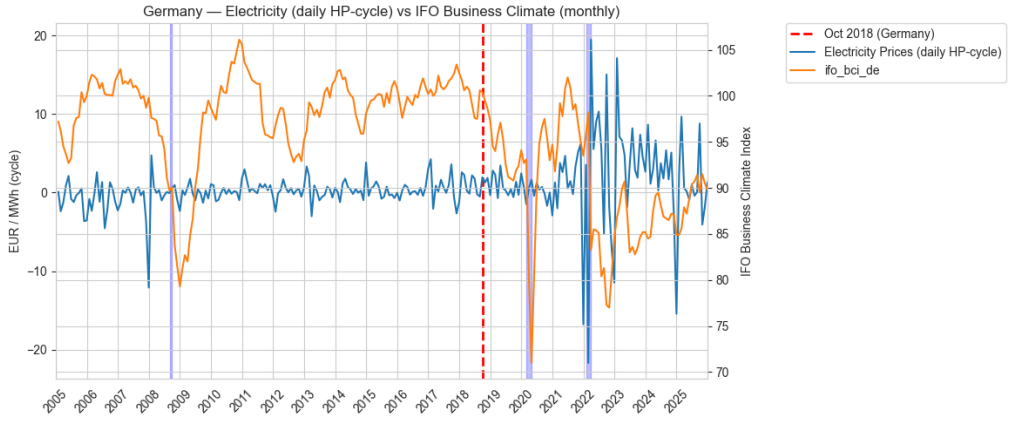
Daily electricity prices respond to macroeconomic conditions, but the indicators that summarise those conditions are released only monthly, while the prices and their usual predictors are daily. This paper forecasts daily baseload electricity prices for Germany and Italy by conditioning them on monthly macroeconomic information through a Reverse-MIDAS (R-MIDAS) model, which lets a high-frequency target depend on low-frequency regressors without ad hoc aggregation. The restricted variant compresses the daily autoregressive lags with an exponential Almon polynomial, which keeps the model estimable by frequentist methods despite the large daily-versus-monthly frequency mismatch. To our knowledge this is the first application of R-MIDAS to electricity price forecasting, and I enlarge the macroeconomic information set beyond industrial production to include forward-looking sentiment indicators (the PMI, the ifo Business Climate Index and the ISTAT Business Confidence Indicator). Using daily data from 2005 to 2025 and a rolling-window comparison at horizons of one to twenty-eight days ahead, I find that the R-MIDAS specifications produce significantly lower forecast errors than daily-only autoregressive models at every horizon in both markets. Industrial production is most informative at the short horizons, the business-climate surveys dominate at the longer ones, and gas is the single most important daily variable throughout. The results support the view that low-frequency macroeconomic information acts as a slow-moving anchor that makes high-frequency forecasts both more accurate and more stable, especially in turbulent periods.

1 Introduction

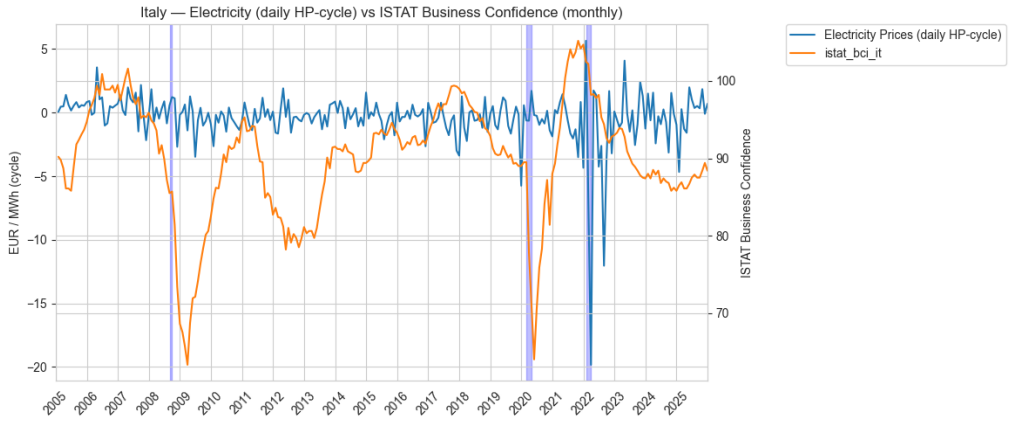
European electricity prices display an high degree of variability. Within a single trading day they can move, according to the Country, from negative values to several hundred euros per megawatt hour. This variability has direct consequences for utilities and generators, who set bidding strategies on the day-ahead and intraday markets, for large industrial consumers, who plan production around energy costs, and for policy makers, who monitor the pass-through of energy prices into headline inflation. According to recent estimates by the European Central Bank, electricity contributes around 9-10% of the Harmonised Index of Consumer Prices and 25-35% of the Producer Price Index, so even moderate improvements in forecast accuracy have a noticeable effect on firms and policy makers. The forecasting horizons that matter for these users are not limited to the day-ahead window: production scheduling for energy-intensive industries such as steel, cement and chemicals, and risk management for generators, both span weeks rather than hours.

The standard econometric toolkit for electricity price forecasting (EPF) has been built around daily, intraday and hourly data. Auto-regressive specifications with calendar dummies, fossil-fuel prices and load forecasts as exogenous regressors are well documented in the surveys of Weron (2014), Nowotarski & Weron (2018) and Lago et al. (2021), and they remain competitive benchmarks even against modern deep-learning models. Their predictive content, however, rests almost entirely on high-frequency information. In turbulent periods this information is not be sufficient. Energy markets are extremely volatile during macroeconomic and geopolitical events, and a model whose inputs are all sampled at the daily frequency can be disrupted by single-day spikes in fuel prices or in residual demand. The resulting forecast path inherits this short-run distortion, producing predictions whose volatility reflects the volatility of the high-frequency regressors rather than the slower and more persistent state of the economy. The episodes of 2020 to 2022, with the COVID-19 lockdowns followed by the invasion of Ukraine, are the most recent and the most visible example: gas and power prices reacted within days to news, but whether each shock turned out to be transient or persistent was determined by macroeconomic factors that move at a much lower frequency, such as industrial activity, business confidence and the inflation regime.

Figure 1 gives a first visual impression of this link. It plots, for Germany and Italy, the main monthly business-climate indicator against the cyclical component of the daily electricity price, that is, the price with its trend and seasonal pattern removed by the Hodrick-Prescott filter (Hodrick & Prescott, 1997). In both countries the survey indicator and the price cycle fall together in the downturns and recover together afterwards, and the link is clearest around the crisis bands of 2008 to 2009, 2020 and 2022. This is the slow-moving co-movement that the rest of the paper sets out to exploit for forecasting.



(a) Germany: ifo Business Climate Index.



(b) Italy: ISTAT Business Confidence Indicator.

Figure 1: Monthly business-climate and business-confidence indicators (right axis) against the cyclical component of the daily electricity price (Hodrick–Prescott filter, left axis), 2005–2025. Shaded bands mark the main crisis episodes.

This paper starts from the conjecture that low-frequency macroeconomic indicators carry information that high-frequency variables cannot reflect. Monthly variables such as the Purchasing Managers' Index, the ifo Business Climate Index, the ISTAT Business Confidence Indicator and the Industrial Production Index summarise firms' short and medium-term expectations together with their realised production. They evolve smoothly compared with daily fuel prices, but they encode something structural: whether the economy is entering or exiting a contraction, whether industrial demand for electricity is expected to grow or to fall, and whether the price pressures observed at high frequency are likely to be transient or persistent. Hüfner & Schröder (2002) and Lehmann (2023) document that German business surveys carry significant predictive information about future activity beyond what is contained in the production data themselves. Dokas et al. (2023) review the macro-energy literature and show that the link between macroeconomic and uncertainty shocks and energy prices has become more pronounced in the post-2008 period. Deniz & Stengos (2024) document that this transmission is asymmetric across optimistic and pessimistic phases of the cycle. The implication for EPF is direct. A

forecasting model that conditions daily electricity prices on monthly macroeconomic indicators, using as "anchor" the underlying state of the economy, dampens the influence of day-by-day noise on the multi-step forecasts, and can deliver predictions that are both more accurate and more stable exactly in the periods in which forecast quality matters most. Industrial and manufacturing firms in particular gain from being able to tell temporary price movements apart from more persistent ones when planning production, a distinction that is hard to draw with daily inputs alone.

It is worth making this stabilising mechanism explicit, because it is the central motivation for the paper. A model that uses only daily inputs has, in effect, only fast-moving information to work with. When it projects the price several days or weeks ahead, it carries the most recent daily shocks forward, so its forecast path moves almost as much as the raw data. For example, a single jump in the gas or brent price is enough to swing the whole forecast. Low-frequency macroeconomic indicators behave in the opposite way. They change slowly and they summarise the broad state of the economy, such as whether industrial demand is expanding or contracting, and whether firms expect activity to improve or to deteriorate over the coming months. What I propose is that when such a variable is added to the model, it acts as a slow-moving anchor, a background level around which the daily price fluctuates. The daily lags still capture the short-run dynamics, but the conditional mean of the forecast is now tied to a quantity that does not jump from one day to the next. The practical effect is twofold. The multi-step forecast inherits less of the day-to-day noise, so the forecast path is smoother and more stable; and at the same time it is centered on a level that reflects the underlying economic conditions rather than the latest spike. This is the sense in which low-frequency information can make a high-frequency forecast both more accurate and less volatile, and the gain should be largest precisely in turbulent periods, when the daily data alone send the most conflicting signals. The same idea, which is pairing a slow-moving component with a fast-moving one, has been shown to improve electricity-price forecasts when long-term and short-term models are pooled (Bordignon et al., 2013), and more generally when low-frequency indicators are brought into a forecasting equation through mixed-frequency regression (Clements & Galvão, 2008).

This stabilizing property is not only a statistical convenience: it maps directly onto the decisions that electricity-market participants face. As already mentioned, utilities and generators set bidding and hedging strategies on the day-ahead and forward markets, where the relevant question is not only the level of next-day prices but the level over the coming weeks. A forecast that is anchored to the macroeconomic state gives them a steadier view of where prices are heading once short-lived spikes are stripped away, which helps in deciding how much capacity to offer and how to hedge forward positions. Energy-intensive manufacturers face the same problem from the demand side: they schedule production and manage procurement around energy costs, and the value of a forecast lies in being able

to separate a temporary price movement from a more persistent one. A forecast centred on slow-moving fundamentals makes that distinction easier than one driven only by daily inputs, because a change in the forecast is then more likely to reflect a genuine shift in conditions than a passing shock. Policy makers and central banks, finally, track the pass-through of energy prices into inflation, and they too benefit from forecasts that tell the persistent component of electricity prices apart from the transient one, since it is the persistent component that matters for the medium-term inflation outlook.

Realizing concretely this conjecture requires a model that can combine a daily target with monthly regressors without resorting to ad hoc aggregation. Aggregating daily prices to the monthly frequency would discard most of the variation that EPF users care about, and disaggregating monthly indicators to the daily frequency would introduce artificial autocorrelation and would not respect the release calendar of the macro variables. The natural device is mixed-frequency regression. The MIDAS family of Ghysels et al. (2007) projects a low-frequency target onto high-frequency regressors with a parsimonious lag-weighting polynomial. Foroni et al. (2018) flip the setting and develop the Reverse-MIDAS framework, in which a high-frequency target is regressed on low-frequency variables together with its own intra-period lags. Within this framework, the restricted variant (R-MIDAS) compresses the high-frequency auto-regressive block through an exponential Almon polynomial with only two shape parameters, which keeps the parameter count manageable when the frequency mismatch is large. With a daily target and a monthly regressor the intra-month position takes $k = 28$ values, so this parsimony is not a stylistic choice but a practical necessity for frequentist estimation. The closest (and to my knowledge, the only one) precedent in EPF, Foroni et al. (2019), adopts the unrestricted Bayesian variant on daily German and Italian electricity prices conditioned on industrial production. This paper takes the restricted variant instead and enlarges the macroeconomic information set by adding forward-looking sentiment indicators (PMI surveys, ifo Business Climate Index and the ISTAT Business Confidence Indicator) alongside disaggregated industrial production. The emphasis on these macroeconomic indicators is deliberate. Industrial production records activity that has already taken place, whereas the business surveys measure what firms expect for the months ahead, so they carry a forward-looking signal about the underlying state of the economy. That state is latent and persistent: it moves slowly and it is not visible in the daily variables. By conditioning the daily price on these macroeconomic indicators, the model ties its forecast to this persistent component of the economy rather than to high-frequency noise, which should make the forecasts more reliable, and the gain should be clearest at the longer horizons.

The empirical exercise uses daily baseload electricity prices for Germany and Italy from 1 January 2005 to 31 December 2025. The two markets are chosen for their differing energy mixes and for the strong response of their power prices to global events such as COVID-19 and the invasion of Ukraine.

The forecasting comparison is run in a 10-year rolling-window setting at horizons $h \in \{1, 7, 14, 21, 28\}$ days ahead, with forecasts starting in January 2015, and uses the best practices of Lago et al. (2021) for comparing forecasting models: an AR plus calendar dummies as benchmark, ARX models with fuel prices and dummies as a daily-only competitor to the R-MIDAS, and the Diebold-Mariano test of Diebold & Mariano (1995) for equal predictive accuracy. The results show that R-MIDAS specifications enriched with sentiment indicators produce significantly lower forecast errors than the daily-only ARX models at all five horizons in both markets. The ifo BCI and the ISTAT BCI emerge as the dominant low-frequency signals at the longer horizons, gas remains the single most informative daily variable. The additional information carried by the monthly indicators which translates into lower forecast errors is consistent with the view that low-frequency macro information acts as a stabiliser for high-frequency forecasts in volatile periods.

The remainder of the paper is organised as follows. Section 2 reviews the three strands of literature on which this work builds: electricity price forecasting, the role of macroeconomic fundamentals and sentiment in energy markets, and mixed-frequency econometric methods. Section 3 describes the daily and monthly data. Section 4 sets out the R-MIDAS specification, the Almon weighting scheme and the rolling-window forecast design. Section 5 reports the results for Germany and Italy and discusses their economic interpretation. Section 6 concludes and outlines two extensions: probabilistic forecasting in the spirit of Nowotarski & Weron (2018) and Gneiting & Katzfuss (2014), and stability tests for forecast performance in unstable environments along the lines of Giacomini & Rossi (2010).

2 Literature Review

The literature this paper builds on intersects three strands. The first is the electricity price forecasting (EPF) literature, which has matured into a well-defined methodological field with established benchmarks for the day-ahead horizon and a growing branch devoted to medium- and long-term predictions. The second is the macro-finance literature on how aggregate activity, fundamentals and sentiment indicators are transmitted into energy and electricity prices. The third is the econometric literature on mixed-frequency time-series methods, in particular the MIDAS family and its reverse counterpart, which provides the formal device for letting daily electricity prices be conditioned on lower-frequency macroeconomic information. The remainder of this section reviews each strand in turn and then locates the contribution of this paper at their intersection.

The starting point of any EPF exercise is the survey by Weron (2014), which consolidated more than two decades of work on short-term electricity price modelling and laid out the canonical taxonomy of statistical, fundamental and computational-intelligence models that still organises the field. Nowotarski

& Weron (2018) extended that review to probabilistic EPF, arguing that point forecasts alone are not adequate for risk-sensitive decisions in volatile electricity markets and systematising the literature on quantile, interval and density forecasts. More recently, Lago et al. (2021) provided an open-access benchmark for day-ahead EPF, comparing classical econometric specifications such as the LEAR (lasso estimated auto-regressive) model with deep neural networks across several European power markets, and formalising a set of best practices (rolling-window evaluation, robust loss metrics, Diebold-Mariano testing) that this paper adopts.

A complementary strand widens the forecast horizon. Bello et al. (2017) combine a regime-switching market model with reduced-form econometrics to produce medium-term probabilistic forecasts, while Ziel & Steinert (2018) propose lasso-based econometric models for the mid- and long-term horizon. Ghelasi & Ziel (2025) push the same modelling toolkit from day-ahead all the way to monthly and yearly horizons, showing that with the appropriate fundamentals econometric models remain competitive well beyond the short term. These contributions motivate the explicit horizon design used in this paper, where forecast accuracy is evaluated at $h \in \{1, 7, 14, 21, 28\}$ days ahead rather than only one day ahead.

A separate question is which fundamentals matter for electricity prices. Karakatsani & Bunn (2008) document that load and fossil-fuel prices, together with time-varying coefficient structures, are essential drivers of UK day-ahead prices. Gianfreda et al. (2020) compare a suite of linear models with high renewable penetration in the Italian market, showing that the relative ranking of forecasting models changes with the penetration of intermittent sources. Bashiri Behmiri et al. (2025) confirm these findings for short-to-mid-term horizons and emphasise the role of renewable generation forecasts as predictors of price levels and price volatility. This strand makes clear that EPF performance hinges on a careful combination of demand-side variables, generation-mix information and fuel prices, all of which are themselves driven, at lower frequencies, by the macroeconomic environment.

The link between the macroeconomy and energy prices has been investigated both at the level of business cycles and at the level of policy and uncertainty shocks. Dokas et al. (2023) provide a comprehensive review of how macroeconomic and uncertainty shocks (including monetary, financial and geopolitical shocks) propagate to energy prices, and underline how their effects have become increasingly state-dependent after the 2008 financial crisis and the post-2020 energy crisis. Focusing on electricity specifically, Deniz & Stengos (2024) document asymmetric effects of global sentiment indicators on EU electricity prices: optimistic and pessimistic phases of the cycle do not transmit symmetrically to the power market. These contributions justify treating macroeconomic information not as background context but as a predictive signal for electricity prices.

A natural channel through which the macroeconomy transmits to power prices is the industrial

sector. Industrial activity alone accounts for roughly a quarter of EU final electricity consumption, so the cyclical position of manufacturing is a first-order determinant of demand. This connects EPF to the broader literature on the forecasting power of business surveys. Hüfner & Schröder (2002) show that German sentiment indicators, notably the ifo Business Climate Index, carry significant predictive information about future industrial production beyond what is contained in the production data themselves. Lehmann (2023) updates this evidence using two decades of additional data, confirming that the ifo Business Survey remains among the most informative leading indicators for German activity. The implication for this paper is that sentiment indicators are not merely contemporaneous reflections of activity: they encode firms' expectations about the next quarter or semester, and should therefore be informative for medium-term electricity demand and, through it, for electricity prices.

Linking daily electricity prices to monthly or quarterly macroeconomic indicators raises an inherent sampling-frequency mismatch. The standard toolkit for this problem is the MIDAS (Mixed Data Sampling) family introduced by Ghysels et al. (2007), in which a low-frequency target is projected on high-frequency regressors whose lag polynomial is parsimonised via a flexible parametric weighting scheme (typically exponential Almon or Beta lag polynomials). The parsimony of the weighting function makes MIDAS estimable by nonlinear least squares even when the frequency mismatch is large. Foroni et al. (2015) relax this restriction by allowing unrestricted lag polynomials (U-MIDAS), and show that U-MIDAS is competitive when the frequency mismatch is small (for example quarterly versus monthly), but suffers from parameter proliferation when the mismatch is large.

The natural counterpart of MIDAS, when the target is high frequency and the regressors are low frequency, is the Reverse-MIDAS framework of Foroni et al. (2018). They develop two specifications: an unrestricted Reverse-MIDAS (RU-MIDAS), in which the daily lags of the target are parameterised individually for each intra-period position, and a restricted Reverse-MIDAS (R-MIDAS), in which the same daily lags are compressed through an Almon polynomial. The unrestricted variant is conceptually straightforward but parameter-heavy: with a daily target and a monthly regressor, the number of intra-month positions explodes and frequentist estimation becomes unreliable. The restricted variant, by contrast, replaces the long block of free coefficients with a smooth two-parameter Almon weighting scheme, which yields a much smaller parameter space and is therefore better suited to the daily-versus-monthly mismatch ($k = 28$) that is relevant for electricity-price applications.

The direct antecedent of this paper is the work of Foroni et al. (2019), who apply a Bayesian RU-MIDAS to daily German and Italian electricity prices conditioned on monthly macroeconomic variables (mainly industrial production). Their Bayesian approach is a practical response to the parameter proliferation issue mentioned above: by imposing informative priors, RU-MIDAS becomes feasible even with $k = 20$ to 30 intra-period positions. The present paper extends Foroni et al. (2019) in two

complementary directions. First, it adopts the restricted R-MIDAS of Foroni et al. (2018), which makes frequentist estimation tractable under a large frequency mismatch without resorting to priors. Second, it enlarges the macroeconomic information set by including forward-looking sentiment indicators (PMI Manufacturing, ifo Business Climate Index and the ISTAT Business Confidence Indicator) alongside the industrial-production indices used in the predecessor, motivated by the survey-forecasting evidence of Hüfner & Schröder (2002) and Lehmann (2023) that sentiment indicators carry information not contained in the production data themselves.

Two further mixed-frequency frameworks are worth mentioning even though their object of interest differs. The mixed-frequency VAR of Schorfheide & Song (2015) provides a system-based approach to joint forecasting of variables observed at different frequencies, with state-space filtering used to recover the latent high-frequency series; the focus there is on conditional and real-time forecasting rather than on parsimony under extreme frequency mismatches. The GARCH-MIDAS class developed by Asgharian et al. (2013) for stock-return variance and applied to oil prices by Pan et al. (2017) uses a low-frequency macroeconomic component to drive the long-run variance of a high-frequency asset, with applications focused on volatility rather than on the conditional mean of the high-frequency series.

The three strands above leave one specific gap. The day-ahead EPF literature has converged on a well-defined benchmark (Lago et al., 2021; Weron, 2014); the medium- and long-term EPF literature has shown that econometric models remain competitive once appropriate fundamentals are added (Bello et al., 2017; Ghelasi & Ziel, 2025; Ziel & Steinert, 2018); and the macro-energy literature has documented that activity and sentiment indicators contain genuine predictive content for energy prices (Deniz & Stengos, 2024; Dokas et al., 2023). On the methodological side, R-MIDAS (Foroni et al., 2018) provides the appropriate device for letting daily prices respond to monthly macro variables under a large frequency mismatch, but to my knowledge it has not been applied to electricity price forecasting before. The closest application (Foroni et al., 2019) uses the unrestricted Bayesian variant and a narrower macro information set focused on industrial production. The contribution of this paper is to bring the restricted R-MIDAS into EPF, to enlarge the macroeconomic information set with forward-looking sentiment indicators (Hüfner & Schröder, 2002; Lehmann, 2023), and to test the resulting model against an ARX benchmark in the spirit of the modern EPF best practices (Lago et al., 2021) across short, medium and longer horizons for the German and Italian markets.

3 Data

The empirical exercise uses daily and monthly data for Germany and Italy from 1 January 2005 to 31 December 2025. The two markets are chosen for three reasons. They are among the largest in

continental Europe; they rely on very different generation mixes, with Germany combining a large share of renewables and coal with gas-fired plants at the margin and Italy depending heavily on gas; and both saw their power prices react strongly to the global events of 2020–2022, which makes the sample well suited to study how forecasts behave in turbulent periods.

The forecast target is the daily baseload electricity price for each country, in euros per megawatt hour (EUR/MWh), obtained from LSEG. The baseload price is the simple average of the hourly day-ahead prices over the 24 hours of the day. Daily electricity prices are highly variable and can reach, especially in the German market where supply can exceed demand due to high renewable penetration, negative and near-zero values. Therefore, usual logarithmic transformation cannot be applied directly to the electricity price. Following the variance-stabilising approach recommended in the forecasting literature (Nowotarski & Weron, 2018; Uniejewski et al., 2018; Ziel & Weron, 2018), the target is modelled on the inverse hyperbolic sine scale,

$$\operatorname{asinh}(x) \;=\; \log\!\left(x + \sqrt{x^2+1}\right).$$

The asinh transform behaves almost exactly like the logarithm for the large positive values that dominate the series, so it compresses spikes in the same way, but it is defined for zero and negative prices as well. It therefore stabilises the variance of the target and keeps the spikes from dominating the estimation, without discarding the days on which prices are negative. All forecasts are produced on the transformed scale and then mapped back to EUR/MWh through the inverse transform, so that the error metrics reported later are expressed in the original units.

Two fuel prices enter as daily exogenous regressors. The first is the price of natural gas at the Title Transfer Facility (TTF) hub and the second is the price of Brent crude oil; both are daily series from LSEG. Gas is the single most important fuel for marginal electricity generation in both markets, so its price feeds almost directly into the cost of producing the next unit of power, while oil captures broader energy-market conditions and is often correlated with gas through indexed contracts. Both fuel prices are strictly positive and are transformed in logarithms. Finally, the daily specification includes a set of deterministic calendar dummies that absorb the strong and well-documented seasonality of electricity prices: day-of-week dummies for the weekly cycle and monthly dummies for the within-year seasonal pattern, as well as a holiday dummy for public holidays (Christmas, New Year's Eve, Easter and Workers Day).

The monthly information set combines forward-looking sentiment indicators with realised industrial production. The full list of the variables and their sources for each country is given in the appendix (Table 1); they are described in turn below.

The sentiment block has two layers. The Purchasing Managers' Indices (PMI) for manufacturing and for services, produced by S&P Global/HCOB, are survey-based diffusion indices that summarise firms' assessment of current and very near-term activity, roughly one month ahead. The business-climate and business-confidence indicators, the ifo Business Climate Index for Germany and the ISTAT Business Confidence Indicator for Italy, are also firm-side surveys, but they ask managers about their expectations for economic activity over a longer window, around the next six months. Including both layers spans the short and the medium-term part of the expectations spectrum: the PMIs carry the one-month signal, the climate and confidence indices carry the six-month signal. The economic logic is direct: when firms expect industrial activity to rise, they expect to consume more electricity, which puts upward pressure on the price; the survey indicators are leading signals of that demand. The evidence that German business surveys, and the ifo index in particular, carry information about future activity beyond what is contained in the production data themselves (Hüfner & Schröder, 2002; Lehmann, 2023) is the main reason for adding them to the information set used in the closest predecessor of this paper.

The second block is the Industrial Production Index (IPI) from Eurostat, broken down into three sub-indices: manufacturing, consumer goods, and the energy branch (electricity, gas and steam). Industrial production is included because it measures realised activity rather than expectations, and because industry alone accounts for roughly a quarter of EU final electricity consumption, so movements in production translate directly into movements in electricity demand. The IPI series are indices with a trend and seasonal patterns, so they are transformed into the first difference of their logarithms, which turns them into monthly growth rates; the sentiment indicators, which are already stationary bounded indices, enter in levels.

Combining a daily target with monthly regressors requires care on two points: the calendar alignment of the two frequencies, and the avoidance of any look-ahead, that is, of using information that would not have been available at the time the forecast is made.

The alignment follows the mixed-frequency convention used by the model and described in detail in Section 4. Each day of the sample is assigned an intra-month position, and a month is treated as having $k = 28$ day-positions, so that the mismatch between the daily target and the monthly regressor is fixed and the same Reverse-MIDAS structure applies to every month. The monthly value is held constant across the days of the month to which it is attached, and the daily distance from that value increases by one day at a time, which is exactly the structure exploited by the lag polynomial in the next section.

Look-ahead is avoided by attaching to each day only the monthly figures that had actually been released by that day. Macroeconomic indicators are published with a delay, so the figure for a given

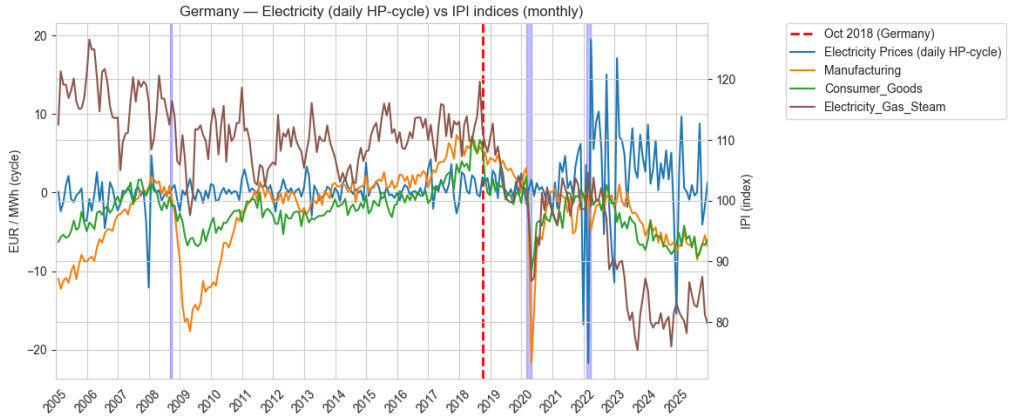
reference month becomes available only some weeks later. I therefore lag each indicator by its publication delay: the Industrial Production Index, which is released with the longest delay, is lagged by two months relative to the reference month, while the survey indicators (PMI and the business-climate indices), which are released earlier, are lagged by one month. As a result, when the model forms a forecast for a given day it conditions only on macroeconomic figures that a forecaster would genuinely have had in hand on that day.

Before turning to the model, it is useful to look at how the daily electricity price and the monthly indicators move together. Comparing a smooth monthly indicator with a noisy daily price directly is not very informative, because the trend and the seasonal pattern of the price hide the cyclical co-movement. To make the relationship visible, the figures below plot the cyclical component of the daily price, obtained with the Hodrick–Prescott filter (Hodrick & Prescott, 1997), against each monthly indicator; the shaded bands mark the main crisis episodes (the 2008–2009 financial crisis, the 2020 COVID-19 shock and the 2022 energy crisis). The co-movement between the business-climate indicators and the price cycle is the clearest of the three, and it is shown for both countries in the introduction (Figure 1): the ifo Business Climate Index for Germany and the ISTAT Business Confidence Indicator for Italy both fall with the price cycle in the downturns and recover with it afterwards, most visibly around the crisis bands, which is precisely where a slow-moving indicator is expected to be most informative. Figure 2 repeats the exercise here for the three industrial-production sub-indices, and the Purchasing Managers' Indices are shown in the appendix (Figure 3). The same broad pattern appears in every case: realised production and survey expectations co-move with the cyclical part of the electricity price, most strongly during the contractions. This descriptive evidence is consistent with the central conjecture of the paper, that the monthly indicators carry a slow-moving signal about the state of the economy that the daily price reflects but that daily data alone cannot isolate.

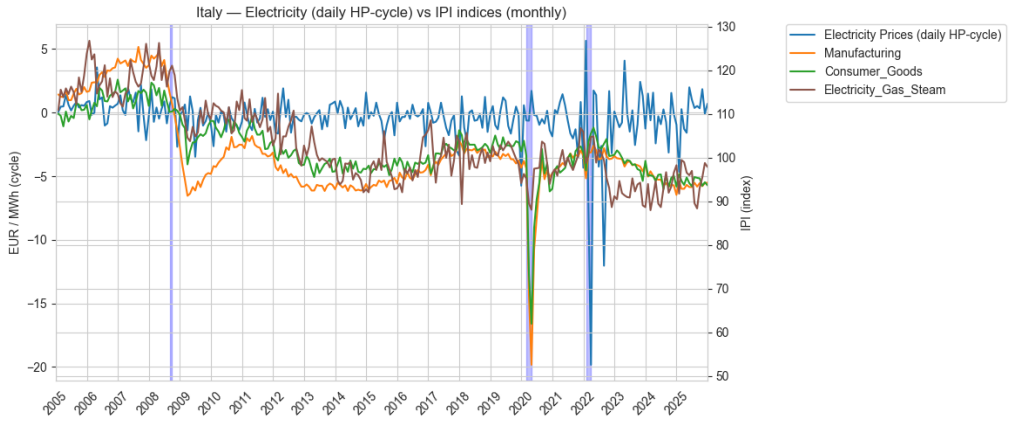
4 The Reverse-MIDAS Model

The MIDAS (Mixed Data Sampling) regression of Ghysels et al. (2007) predicts a low-frequency variable from a high-frequency one. Here the problem is the reverse: the target is the high-frequency series, the daily electricity price, and the additional predictor is low frequency, a monthly macroeconomic indicator. The framework built for this case is the Reverse-MIDAS of Foroni et al. (2018), which I follow throughout, swapping their variable names so that the daily target is y_t and the monthly regressor x_t.

Let t index the daily observations, and let $k = 28$ be the number of day-positions in a month, so



(a) Germany.



(b) Italy.

Figure 2: Eurostat Industrial Production Index (manufacturing, consumer goods and energy; monthly, right axis) against the cyclical component of the daily electricity price (Hodrick–Prescott filter, left axis), 2005–2025.

that k is the frequency mismatch between the daily target and the monthly regressor. The within-month position is $I = 0, 1, \ldots, k-1$. The lag operator L works at the daily frequency, $L^i y_t = y_{t-i/k}$, so that L^k shifts a series back by one month. The daily electricity price is the target y_t. The monthly indicator x_t is observed once a month and is constant within the month to which it refers, so that $L^{k+i} x_t$ returns the previous month's value at every within-month position i. In the empirical part x_t collects one or more of the indicators of Section 3, each lagged by its publication delay so that only genuinely available figures enter; when more than one indicator is used, x_t is replaced by the vector $\boldsymbol{x}_t$.

Following Foroni et al. (2018), the starting point is an autoregressive process for the daily price driven by the latent daily counterpart x_t^* of the monthly indicator, that is, the path the macro variable would trace if it were observed every day,

$$c(L)\, y_t \;=\; d(L)\, x_t^* \;+\; e_t, \tag{1}$$

with $c(L) = 1 - c_1 L - \cdots - c_p L^p$, $d(L) = d_1 L + \cdots + d_q L^q$ and white-noise e_t. The latent x_t^* is never observed: what is observed is the monthly value, which is its temporal aggregate $x_t = \omega(L)\, x_t^*$, where

$$\omega(L) = \omega_0 + \omega_1 L + \cdots + \omega_{k-1} L^{k-1}$$

is the aggregation choice (a monthly average, month-end value or sum). Multiplying (1) by a position-specific lag polynomial that isolates the lags aligned with the monthly dates and then substituting the aggregation rule eliminates x_t^*, leaving the daily price written in terms of observable quantities only. The polynomial depends on the day of the month, so the argument yields one equation per within-month position,

$$\tilde{c}_i(L)\, y_t \;=\; g_i\big(L^{\,k+i}\big)\, x_t \;+\; \tilde{e}_{i,t}, \qquad I = 0, 1, \ldots, k-1. \tag{2}$$

Two features carry the economic content. The monthly regressor enters at lag $L^{\,k+i}$: standing on the i-th day of the month, the most recent observable monthly figure is the previous month's value, which is $k + i$ days old, so its "staleness" grows by one day as it moves through the month. And the specification is periodic, since $\tilde{c}_i$ and g_i differ across positions i; this is a consequence of mixing a daily target with a monthly regressor, not a modelling choice.

By having a separate regression for each of the $k = 28$ within-month positions, with up to three monthly indicators and the daily lags, the model would carry dozens to a few hundred free coefficients, too many to estimate reliably on a rolling window. This would be the unrestricted Reverse-MIDAS (RU-MIDAS) variant, that Foroni et al. (2019) keep feasible with Bayesian shrinkage priors; I take the restricted route instead.

The restricted Reverse-MIDAS (R-MIDAS) model replaces the free daily autoregressive coefficients by a smooth exponential-Almon weighting curve. Following the literature of Weron (2014) and Lago et al. (2021), we use a small set of daily lags, $\mathcal{J} = \{1, 2, 3, 7\}$, that is yesterday, the two days before it, and the same day one week earlier, and weight each lag by

$$b(j; \boldsymbol{\theta}) \;=\; \frac{\exp\big(\theta_1\, j + \theta_2\, j^2\big)}{\sum_{j \in \mathcal{J}} \exp(\theta_1\, j + \theta_2\, j^2)}, \qquad j \in \mathcal{J}, \tag{3}$$

which depends on two shape parameters $\boldsymbol{\theta} = (\theta_1, \theta_2)$ and sums to one. The restriction $\theta_2 \leq 0$ keeps the weights from exploding, so recent days receive most of the weight. Collecting them into the polynomial $B(L, \boldsymbol{\theta}) = \sum_{j \in \mathcal{J}} b(j; \boldsymbol{\theta})\, L^j$, where $L^j y_t = y_{t - j/k}$, the R-MIDAS equation, in the form of Foroni et al. (2018), is

$$y_t \;=\; \lambda_i\, L^{\,k+i}\, x_t \;+\; \delta\, B(L, \boldsymbol{\theta})\, y_t \;+\; \varepsilon_t. \tag{4}$$

The first term is the monthly regressor with its staleness L^{k+i} and loading λ_i; the second is the daily autoregressive block, whose four lag coefficients are now summarised by the single scale δ and the two Almon parameters $\boldsymbol{\theta}$, the implied coefficient on lag j being $\delta\, b(j;\boldsymbol{\theta})$. In the estimated specification δ and $\boldsymbol{\theta}$ are held common across within-month positions, so the only periodic element left is the staleness L^{k+i}, which simply selects the most recent available monthly value. As an illustration, on the first three days of a month the same monthly value enters as $L^{28}x_t$, then $L^{29}x_t$, then $L^{30}x_t$: the value is unchanged, only one day older each day, while the autoregressive block moves with the new prices.

In the estimated specification the Almon polynomial is applied only to the daily price lags. The fuel prices, natural gas and Brent oil, enter linearly with one daily lag each, since a single lag of the previous day's fuel price is the standard fundamental control in daily electricity price models and there is nothing to compress; the monthly indicators and the calendar dummies also enter linearly. Writing $z_t^{\mathrm{ar}}(\boldsymbol{\theta}) = \sum_{j\in\mathcal{J}} b(j;\boldsymbol{\theta})\, y_{t-j/k}$ for the Almon-weighted autoregressive block and $\mathrm{gas}_{t-1/k}$, $\mathrm{oil}_{t-1/k}$ for the one-day fuel lags, the full one-step model is

$$y_t \;=\; \alpha \;+\; \boldsymbol{\lambda}^{\top}\boldsymbol{x}_t \;+\; \delta\, z_t^{\mathrm{ar}}(\boldsymbol{\theta}) \;+\; \gamma^{g}\,\mathrm{gas}_{t-1/k} \;+\; \gamma^{o}\,\mathrm{oil}_{t-1/k} \;+\; \boldsymbol{\eta}^{\top}\boldsymbol{d}_t \;+\; \varepsilon_t. \tag{5}$$

The linear parameters are the intercept α, the macro loadings $\boldsymbol{\lambda}$, the autoregressive scale δ, the fuel coefficients γ^g and γ^o and the control loadings $\boldsymbol{\eta}$; the only nonlinear parameters are the two Almon shapes $\boldsymbol{\theta}$. Pooling across positions keeps the model to roughly two dozen parameters, against the few hundred of the unrestricted variant, which is what makes it estimable on a rolling window; the price is the standard bias-variance trade-off discussed by Foroni et al. (2015) and Andreou et al. (2010).

Estimation. Model (5) is linear in all of its parameters except the two Almon shapes $\boldsymbol{\theta} = (\theta_1, \theta_2)$, which enter nonlinearly through the weights (3). Concentrated nonlinear least squares exploits this separability, splitting the estimation on each rolling window W into an inner linear step and an outer two-dimensional search. The window W is the set of consecutive daily observations used for one estimation: it spans ten years of daily data (about 3,650 days) and is rolled forward one day at a time, re-estimating the model at each new origin.

In the first step the shape parameters $\boldsymbol{\theta}$ are held fixed. The weights $b(j;\boldsymbol{\theta})$ are then known numbers, so the Almon-weighted block $z_t^{\mathrm{ar}}(\boldsymbol{\theta})$ is known data and (5) is an ordinary linear regression: all the remaining coefficients, that is the intercept, the macro loadings, the autoregressive scale δ, the two fuel coefficients and the control loadings, are obtained by ordinary least squares.

In the second step these least-squares coefficients are substituted back into the model, which leaves

the residual sum of squares as a function of the two shape parameters alone, the concentrated criterion

$$\mathrm{RSS}(\boldsymbol{\theta}) \;=\; \sum_{t \in W} \Big(y_t - \hat{y}_t(\boldsymbol{\theta}) \Big)^2, \tag{6}$$

where $\hat{y}_t(\boldsymbol{\theta})$ is the fitted price at those coefficients (the notation $\sum_{t \in W}$ is just the sum over all the days contained in the current ten-year window). Nonlinear least squares then minimises $\mathrm{RSS}(\boldsymbol{\theta})$ over $\boldsymbol{\theta}$. The linear coefficients are then read off at the optimum. The whole procedure is repeated as the window rolls forward, giving one set of estimates per forecast origin. Because the autoregressive memory of the price is absorbed by the Almon block rather than left in an autocorrelated error, no error-correction term is needed at the forecast stage.

Forecasting and evaluation. Forecasts are produced at horizons $h \in \{1, 7, 14, 21, 28\}$ days ahead by iterating the one-step model (5) forward, in the iterated multi-step manner of Marcellino et al. (2006): one model is estimated per origin and rolled forward. At each step the calendar dummies are set at their true future dates; the monthly indicators and the one-day fuel lags are held at their last values available at the origin; and the autoregressive block is rebuilt from actual prices for the lags still observed and from previous-step forecasts for the rest. Writing $\widetilde{y}_{t+s}$ for the sequence that equals the realised price for $s \le 0$ and the forecast for $s \ge 1$,

$$\widehat{y}_{t+s} \;=\; \hat{\alpha} + \hat{\boldsymbol{\lambda}}^{\top} \boldsymbol{x}_t + \hat{\delta} \sum_{j \in \mathcal{J}} b(j; \hat{\boldsymbol{\theta}})\, \widetilde{y}_{t+s-j/k} + \hat{\gamma}^{g}\, \mathrm{gas}_t + \hat{\gamma}^{o}\, \mathrm{oil}_t + \hat{\boldsymbol{\eta}}^{\top} \widetilde{\boldsymbol{d}}_{t+s}, \tag{7}$$

run forward from $s = 1$ to $s = h$. The estimation uses the ten-year rolling window W described above, and the out-of-sample period begins in January 2015. The benchmark is an autoregression on the four daily lags in $\mathcal{J}$ together with the calendar dummies; it reaches back one week, which is why I label it AR(7), but it is a restricted autoregression on these four lags rather than a full seven-lag model. The second competitor is an ARX model that adds gas and Brent prices to that benchmark. The same four daily lags form the autoregressive core of all three model families, so they differ only in the predictors they add to it. Accuracy is summarised by the root mean squared error (RMSE) on the original EUR/MWh scale, reported as a ratio to the benchmark, so that a value below one means the model beats the benchmark. Whether a competitor is significantly more accurate is assessed pairwise with the Diebold-Mariano test of Diebold & Mariano (1995) The Diebold-Mariano test is the standard tool for comparing forecasts in the electricity price forecasting literature (Lago et al., 2021; Weron, 2014), and its outcome is reported as the significance stars attached to each entry of the results tables in Section 5.

5 Results

This section reports the out-of-sample forecasting results for Germany and Italy. The design is the one set out in Section 4: a long rolling window of daily data, iterated forecasts at horizons $h \in \{1, 7, 14, 21, 28\}$ days ahead, an out-of-sample period that starts in January 2015, and the AR(7) model with calendar dummies as the common benchmark. Every model is scored by its root mean squared error (RMSE) on the EUR/MWh scale; the figure in parentheses below each RMSE is the ratio to the benchmark RMSE at that horizon, so a value below one means the model beats the AR(7) benchmark. Stars report the Diebold & Mariano (1995) test of equal predictive accuracy with the Harvey et al. (1997) small-sample correction, marking the cases in which the competing model has a significantly lower expected loss than the benchmark. In each column the lowest RMSE is in bold. To keep the main text readable, the full set of RMSE figures, ratios and significance stars is collected in Tables 2 and 3 in the appendix, and the discussion below refers to them throughout.

The first message of Tables 2 and 3 is that conditioning on fundamentals helps at every horizon. In both markets the daily-only ARX model, which adds gas (and, for Germany, Brent) to the benchmark, already cuts the forecast error substantially: for Germany the RMSE ratio falls to between 0.66 and 0.74 across horizons, and for Italy to between 0.88 and 0.96, all significant on the Diebold–Mariano test. The improvement is larger for Germany than for Italy, which is consistent with the more pronounced reaction of German power prices to the gas market over the sample.

The second and central message is that adding the monthly macroeconomic block through the R-MIDAS structure improves on the daily-only ARX model at every horizon in both markets. For Germany the best R-MIDAS ratio is around 0.66 at one day ahead and falls to roughly 0.57 at two and three weeks ahead, against the ARX figures of 0.74 and 0.66 at the same horizons; the gain from the monthly block is therefore not confined to the short end but, if anything, widens at the medium horizons. For Italy the R-MIDAS ratios fall to about 0.82–0.83 at three and four weeks ahead, clearly below the ARX figures of 0.88 around those horizons, and the significance of the Diebold–Mariano test strengthens from one or two stars for the ARX model to three stars for the R-MIDAS specifications at the longer horizons. In short, the monthly indicators carry predictive content that the daily fuel prices do not, and that content is most valuable precisely where the daily-only model is weakest, at the longer horizons.

The composition of the best-performing specifications tells a consistent story about which information matters and when. The industrial-production block is most useful at the short end and fades as the horizon lengthens: for Germany the IPI set in the best model thins from three sub-indices at $h = 1$ to none at $h = 28$, and for Italy the production sub-indices enter the best specification only at the $h = 7$ and $h = 14$ horizons. This is what one would expect of a coincident measure of realised activity, whose

informational advantage decays as the forecast reaches further into the future. The forward-looking sentiment indicators behave in the opposite way. Every best-performing German specification contains the combination of PMI Manufacturing and the ifo Business Climate Index, and for Italy the ISTAT Business Confidence Indicator takes over as the dominant monthly signal at the three- and four-week horizons. Because the business-climate surveys summarise firms' expectations about activity over the coming months rather than current output, they remain informative exactly where the production data lose their edge. Among the daily variables, gas enters every best-performing specification in both markets, while Brent adds value for Germany only at the one-day horizon; this confirms the central role of the gas market in setting marginal generation costs in these two economies.

Taken together, the results support the conjecture set out in the introduction. The monthly macroeconomic indicators act as a slow-moving anchor for the daily forecast: they lower the point-forecast error relative to a daily-only model, and the improvement is largest at the medium horizons, where a forecast driven only by recent daily shocks is least reliable. The division of labour between coincident and forward-looking indicators, with production carrying the short horizons and the business-climate surveys the longer ones, is the empirical counterpart of the idea that low-frequency information stabilises a high-frequency forecast by tying its conditional mean to the underlying state of the economy. The pattern holds across two markets with different generation mixes and different price dynamics, which suggests that it reflects a general feature of the relationship between macroeconomic conditions and electricity prices rather than a peculiarity of either country.

6 Conclusions

This paper asked whether daily electricity prices can be forecast more accurately by conditioning them on monthly macroeconomic information, and how to do so without distorting either frequency. The answer rests on the Reverse-MIDAS framework, which lets a daily target depend on monthly regressors together with its own daily lags. The restricted variant compresses the daily autoregressive lags with an exponential Almon polynomial, which keeps the model estimable by frequentist methods even though a month contains twenty-eight day-positions. To our knowledge this is the first application of R-MIDAS to electricity price forecasting, and it enlarges the macroeconomic information set beyond industrial production to include forward-looking sentiment indicators, the Purchasing Managers' Indices, the ifo Business Climate Index and the ISTAT Business Confidence Indicator.

The empirical results for Germany and Italy are clear and consistent across the two markets. The R-MIDAS specifications produce significantly lower forecast errors than daily-only autoregressive models at every horizon from one to twenty-eight days ahead, and the gain is largest at the medium

horizons, where a forecast driven only by recent daily shocks is least reliable. The composition of the best models shows a clear division of labour: industrial production is most informative at the short horizons and fades as the horizon lengthens, while the business-climate surveys, which summarise firms' expectations about activity over the coming months, take over as the dominant signal at the longer horizons. Among the daily variables, gas is the single most important predictor in both markets. These findings support the central conjecture of the paper, that low-frequency macroeconomic information acts as a slow-moving anchor that makes a high-frequency forecast both more accurate and more stable, and they carry directly over to the users who care about these forecasts: utilities and generators setting bidding and hedging strategies, energy-intensive manufacturers planning production, and policy makers tracking the pass-through of energy prices into inflation.

Several extensions would make these results more robust, and I leave them for future work. The first concerns the role of the 2021 to 2022 energy crisis. Because daily prices in those two years moved far more than in the rest of the sample, their squared errors dominate any full-sample root mean squared error, so a single figure mixes a calm regime with an extreme one. Repeating the comparison separately over the pre-crisis years (2015 to 2020), the crisis (2021 to 2022) and the post-crisis years (2023 to 2025) would show whether the gain from the monthly block is a general property or an artefact of the crisis period, and whether the ranking of the specifications is stable across regimes. A natural complement is to examine the stability of relative performance directly with the fluctuation test of Giacomini & Rossi (2010), which traces the Diebold-Mariano statistic through a rolling window across the out-of-sample period rather than averaging it into a single number, and compares its path against critical values that allow for the search across windows. This would reveal whether the advantage of the macro-augmented model holds throughout the sample or is concentrated in particular episodes.

A third extension addresses the fragility of selecting a single best specification after the fact, given that the best model differs across horizons and across regimes. Instead of picking one specification, the whole set can be combined, for example through an equal-weight average, a trimmed mean that discards the most extreme forecasts, or inverse-MSE weights in the spirit of Bates & Granger (1969), computed in real time from forecast errors already resolved at each origin. Forecast combination has a long record of success in electricity price forecasting (Bordignon et al., 2013; Weron, 2014), and it would provide a summary of the value of the monthly block that does not depend on having picked the right specification in advance. Finally, the analysis could move from point to probabilistic forecasts, in the spirit of Nowotarski & Weron (2018) and Gneiting & Katzfuss (2014). A single expected value is a thin summary for a market as volatile as electricity, and a full predictive distribution would let utilities and firms align their bidding and production decisions with the risks they actually face, which is where the stabilising role of low-frequency information is likely to matter most.

References

Andreou, E., Ghysels, E., & Kourtellos, A. (2010). Regression models with mixed sampling frequencies. *Journal of Econometrics*, 158(2), 246–261.

Asgharian, H., Hou, A. J., & Javed, F. (2013). The importance of the macroeconomic variables in forecasting stock return variance: A GARCH-MIDAS approach. *Journal of Forecasting*, 32(7), 600–612.

Bashiri Behmiri, N., Fezzi, C., & Ravazzolo, F. (2025). Renewable sources and short-to-mid-term electricity price forecasting. *Energy Economics*.

Bates, J. M., & Granger, C. W. J. (1969). The combination of forecasts. *Operational Research Quarterly*, 20(4), 451–468.

Bello, A., Reneses, J., Bunn, D. W., & Muñoz, A. (2017). Medium-term probabilistic forecasting of electricity prices: A hybrid approach. *IEEE Transactions on Power Systems*, 32(1), 334–343.

Bordignon, S., Bunn, D. W., Lisi, F., & Nan, F. (2013). Combination of long-term and short-term forecasts for electricity prices. *Energy Economics*, 35, 88–103.

Clements, M. P., & Galvão, A. B. (2008). Macroeconomic forecasting with mixed-frequency data: Forecasting output growth in the United States. *Journal of Business & Economic Statistics*, 26(4), 546–554.

Deniz, P., & Stengos, T. (2024). Asymmetric effects of global sentiments on EU electricity prices. *The Journal of Economic Asymmetries*.

Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263.

Dokas, I., Oikonomou, G., Panagiotidis, M., & Spyromitros, E. (2023). Macroeconomic and uncertainty shocks' effects on energy prices: A comprehensive literature review. *Energies*, 16(3), 1491.

Foroni, C., Guérin, P., & Marcellino, M. (2018). Using low frequency information for predicting high frequency variables. *International Journal of Forecasting*, 34(4), 774–787.

Foroni, C., Marcellino, M., & Schumacher, C. (2015). Unrestricted mixed data sampling (MIDAS): MIDAS regressions with unrestricted lag polynomials. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 178(1), 57–82.

Foroni, C., Ravazzolo, F., & Rossini, L. (2019). *Forecasting daily electricity prices with monthly macroeconomic variables* (Working Paper Series No. 2250). European Central Bank.

Ghelasi, P., & Ziel, F. (2025). From day-ahead to mid and long-term horizons with econometric electricity price forecasting models. *Renewable and Sustainable Energy Reviews*.

Ghysels, E., Sinko, A., & Valkanov, R. (2007). MIDAS regressions: Further results and new directions. *Econometric Reviews*, 26(1), 53–90.

Giacomini, R., & Rossi, B. (2010). Forecast comparisons in unstable environments. *Journal of Applied Econometrics*, 25(4), 595–620.

Gianfreda, A., Ravazzolo, F., & Rossini, L. (2020). Comparing the forecasting performances of linear models for electricity prices with high RES penetration. *International Journal of Forecasting*, 36(3), 974–986.

Gneiting, T., & Katzfuss, M. (2014). Probabilistic forecasting. *Annual Review of Statistics and Its Application*, 1, 125–151.

Harvey, D., Leybourne, S., & Newbold, P. (1997). Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, 13(2), 281–291.

Hodrick, R. J., & Prescott, E. C. (1997). Postwar U.S. business cycles: An empirical investigation. *Journal of Money, Credit and Banking*, 29(1), 1–16.

Hüfner, F. P., & Schröder, M. (2002). Forecasting economic activity in Germany – how useful are sentiment indicators? *ZEW Discussion Paper*, (02-56).

Karakatsani, N. V., & Bunn, D. W. (2008). Forecasting electricity prices: The impact of fundamentals and time-varying coefficients. *International Journal of Forecasting*, 24(4), 764–785.

Lago, J., Marcjasz, G., De Schutter, B., & Weron, R. (2021). Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark. *Applied Energy*, 293, 116983.

Lehmann, R. (2023). The forecasting power of the ifo business survey. *Journal of Business Cycle Research*, 19(1), 43–94.

Marcellino, M., Stock, J. H., & Watson, M. W. (2006). A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series. *Journal of Econometrics*, 135(1–2), 499–526.

Nowotarski, J., & Weron, R. (2018). Recent advances in electricity price forecasting: A review of probabilistic forecasting. *Renewable and Sustainable Energy Reviews*, 81, 1548–1568.

Pan, Z., Wang, Y., Wu, C., & Yin, L. (2017). Oil price volatility and macroeconomic fundamentals: A regime switching GARCH-MIDAS model. *Journal of Empirical Finance*, 43, 130–142.

Schorfheide, F., & Song, D. (2015). Real-time forecasting with a mixed-frequency VAR. *Journal of Business & Economic Statistics*, 33(3), 366–380.

Uniejewski, B., Weron, R., & Ziel, F. (2018). Variance stabilizing transformations for electricity spot price forecasting. *IEEE Transactions on Power Systems*, 33(4), 2219–2233.

Weron, R. (2014). Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting*, 30(4), 1030–1081.

Ziel, F., & Steinert, R. (2018). Probabilistic mid- and long-term electricity price forecasting. *Renewable and Sustainable Energy Reviews*, 94, 251–266.

Ziel, F., & Weron, R. (2018). Day-ahead electricity price forecasting with high-dimensional structures: Univariate vs. multivariate modeling frameworks. *Energy Economics*, 70, 396–420.

A Monthly macroeconomic variables

Table 1 lists the monthly macroeconomic variables used in the empirical part, together with their sources for each country, completing the description given in Section 3.

Table 1: Monthly macroeconomic variables and their sources.

Indicator	Germany	Italy
Sentiment indicators (forward-looking, in levels):		
PMI Manufacturing	S&P Global/HCOB PMI Mfg DE	S&P Global/HCOB PMI Mfg IT
PMI Services	S&P Global/HCOB PMI Svc DE	S&P Global/HCOB PMI Svc IT
Business climate / confidence	ifo Business Climate Index	ISTAT Business Confidence Indicator
Industrial Production Index (Eurostat, first difference of logs):		
Manufacturing	IPI Manufacturing DE	IPI Manufacturing IT
Consumer goods	IPI Consumer Goods DE	IPI Consumer Goods IT
Energy (elec./gas/steam)	IPI Energy DE	IPI Energy IT

B Forecast accuracy tables

Tables 2 and 3 report the full forecast comparison discussed in Section 5, for Germany and Italy respectively.

Table 2: Forecast accuracy for Germany. Daily forecasts on the EUR/MWh scale. Benchmark = AR(7) + dummies. Each cell: RMSE on top, ratio to the benchmark in parentheses below. Stars indicate Diebold–Mariano significance when the competing model has lower expected loss ($^{*}p < .10$, $^{**}p < .05$, $^{***}p < .01$). Bold = lowest RMSE at that horizon. PMI = PMI Manufacturing; ifo = ifo Business Climate Index; IPI sub-indices: CG = Consumer Goods, Mfg = Manufacturing, En = Energy.

Model	h=1	h=7	h=14	h=21	h=28
AR(7)+dum (benchmark)	34.4	64.1	75.4	80.0	82.4
AR with daily controls:					
AR(7)+dum+gas+brent	25.4 (0.739)***	44.1 (0.689)***	49.9 (0.662)***	53.3 (0.667)**	56.6 (0.687)**
R-MIDAS:					
R-MIDAS+gas+brent+PMI+ifo+IPI$_{\text{CG,Mfg,En}}$	**22.6** (0.657)***	39.2 (0.611)***	43.6 (0.578)***	48.2 (0.602)**	53.1 (0.645)*
R-MIDAS+gas+PMI+ifo+IPI$_{\text{CG,Mfg}}$	22.8 (0.662)***	**38.8** (0.605)***	42.8 (0.567)***	47.0 (0.588)**	51.8 (0.629)*
R-MIDAS+gas+PMI+ifo+IPI$_{\text{Mfg}}$	22.9 (0.665)***	38.9 (0.606)***	**42.6** (0.566)***	**46.9** (0.586)**	51.7 (0.627)*
R-MIDAS+gas+PMI+ifo	23.0 (0.668)***	38.9 (0.606)***	42.7 (0.567)***	47.0 (0.587)**	**51.6** (0.626)*

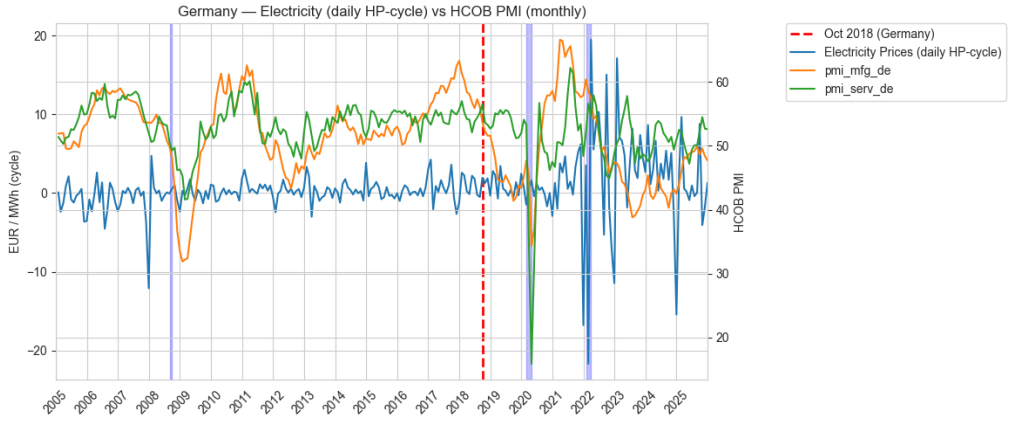
C Additional descriptive figures

This appendix completes the descriptive evidence of Section 3. Figure 3 shows the Purchasing Managers' Indices against the cyclical component of the daily electricity price, in the same Hodrick–Prescott form used for the business-climate and industrial-production figures in the main text. The remaining

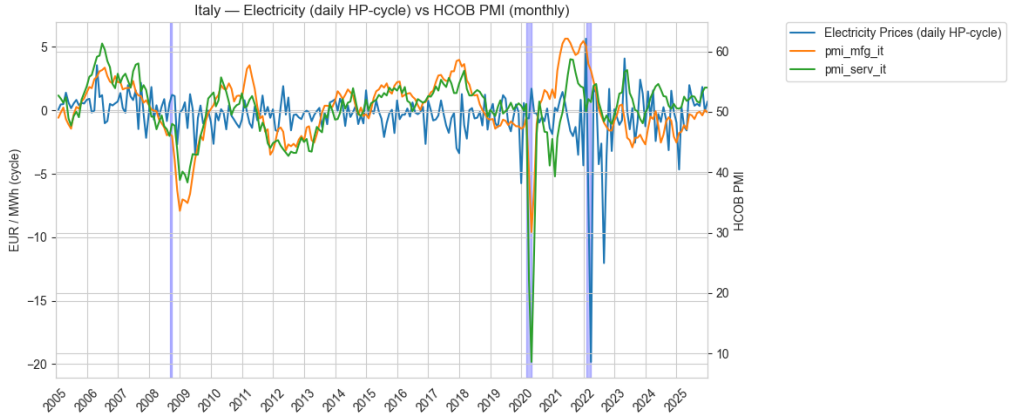
Table 3: Forecast accuracy for Italy. Daily forecasts on the EUR/MWh scale. Benchmark = AR(7) + dummies. Cell and significance conventions as in Table 2. ISTAT = ISTAT Business Confidence Indicator; IPI sub-indices: CG = Consumer Goods, En = Energy.

Model	h=1	h=7	h=14	h=21	h=28
AR(7)+dum (benchmark)	18.2	33.9	39.9	42.3	43.6
AR with daily controls:					
AR(7)+dum+gas	17.4 (0.958)*	31.6 (0.931)*	36.5 (0.915)*	37.3 (0.882)**	38.3 (0.879)*
R-MIDAS:					
R-MIDAS+gas+$\text{IPI}_{\text{CG,En}}$+ISTAT	17.3 (0.952)**	**30.4** (0.896)**	**34.4** (0.862)**	35.3 (0.834)***	36.2 (0.831)***
R-MIDAS+gas+ISTAT	**17.3** (0.951)**	30.5 (0.899)***	34.5 (0.865)***	**35.2** (0.831)***	**35.9** (0.823)***

figures report the same comparisons under two alternative views of the price. The first set replaces the filtered daily price with the monthly median price, which removes the within-month noise and makes the level co-movement easier to read. The second set uses the raw daily price, which keeps the full short-run variation and shows how the daily series behaves around the slow-moving indicators, in particular during the 2020–2022 crisis episodes. In every figure the indicator is on the right axis and the electricity price on the left axis, and the shaded bands mark the main crisis episodes.

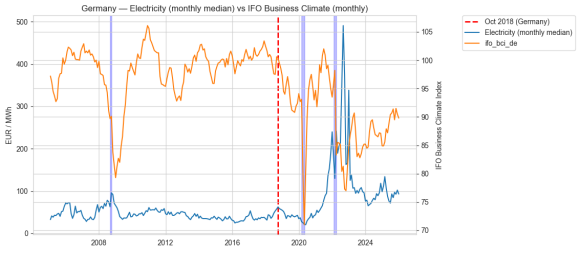


(a) Germany.

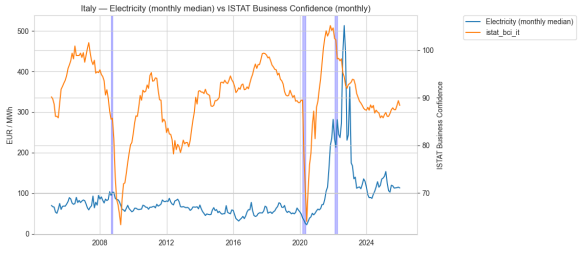


(b) Italy.

Figure 3: Purchasing Managers' Indices for manufacturing and services (monthly, right axis) against the cyclical component of the daily electricity price (Hodrick–Prescott filter, left axis), 2005–2025.

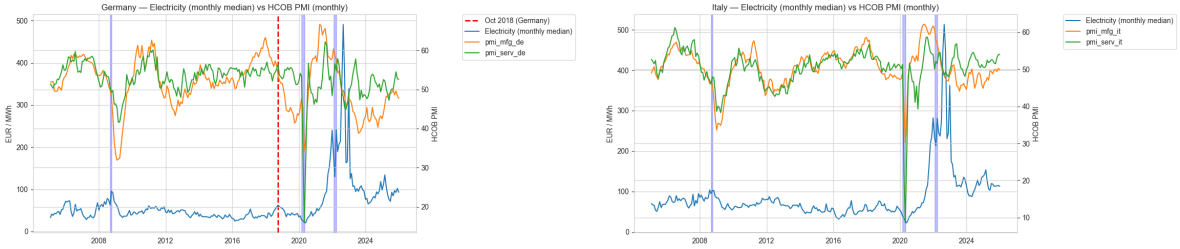


(a) Germany: ifo Business Climate Index.



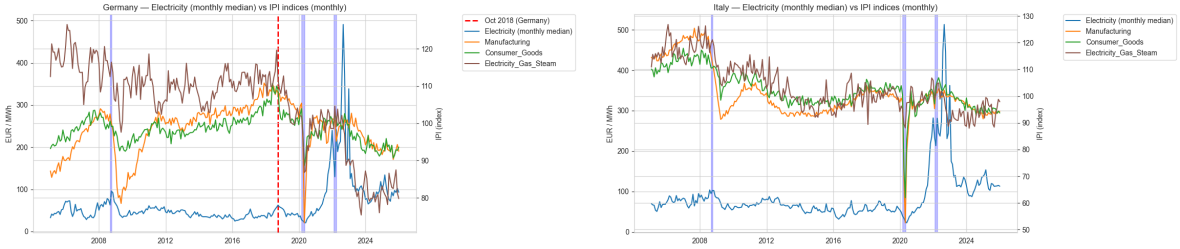
(b) Italy: ISTAT Business Confidence Indicator.

Figure 4: Business-climate and business-confidence indicators (monthly, right axis) against the monthly median electricity price (left axis), 2005–2025.



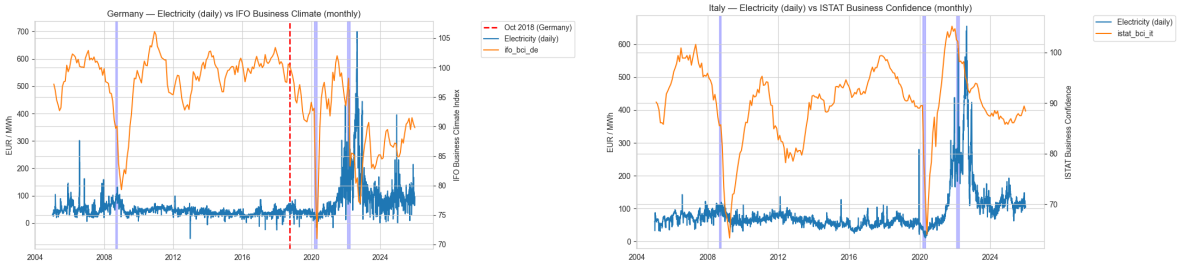
(a) Germany. (b) Italy.

Figure 5: Purchasing Managers' Indices for manufacturing and services (monthly, right axis) against the monthly median electricity price (left axis), 2005–2025.



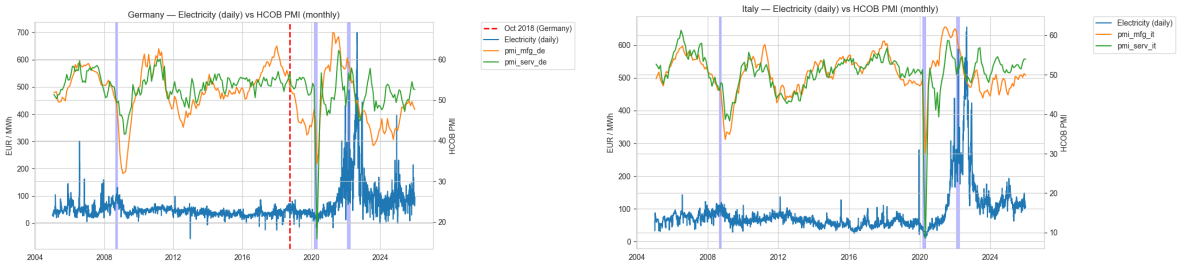
(a) Germany. (b) Italy.

Figure 6: Eurostat Industrial Production Index (manufacturing, consumer goods and energy; monthly, right axis) against the monthly median electricity price (left axis), 2005–2025.



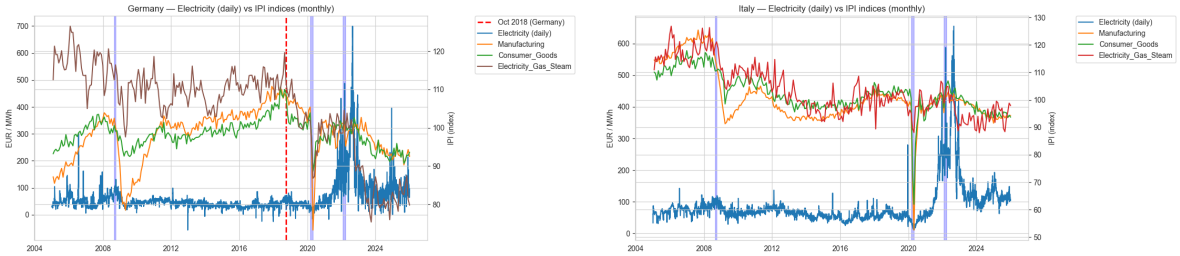
(a) Germany: ifo Business Climate Index. (b) Italy: ISTAT Business Confidence Indicator.

Figure 7: Business-climate and business-confidence indicators (monthly, right axis) against the raw daily electricity price (left axis), 2005–2025.



(a) Germany. (b) Italy.

Figure 8: Purchasing Managers' Indices for manufacturing and services (monthly, right axis) against the raw daily electricity price (left axis), 2005–2025.



(a) Germany.

(b) Italy.

Figure 9: Eurostat Industrial Production Index (manufacturing, consumer goods and energy; monthly, right axis) against the raw daily electricity price (left axis), 2005–2025.