Energy Economics 126 (2023) 107006

Contents lists available at ScienceDirect

Energy Economics

journal homepage: www.elsevier.com/locate/eneeco

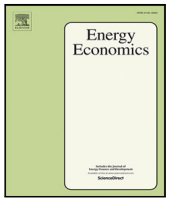

Nowcasting industrial production using linear and non-linear models of electricity demand

Giulio Galdi [a,*], Roberto Casarin [b], Davide Ferrari [c], Carlo Fezzi [a,d], Francesco Ravazzolo [c,e,f]

[a] Dipartimento di Economia e Management, Università di Trento, Trento, Italy
[b] Dipartimento di Economia, Università di Venezia, Venezia, Italy
[c] Facoltà di Economia e Management, Libera Università di Bolzano, Bolzano, Italy
[d] Land, Environmental, Economics and Policy Institute, University of Exeter, Exeter, United Kingdom
[e] BI Norwegian Business School, Oslo, Norway
[f] Rimini Centre for Economic Analysis (RCEA), Italy

ARTICLE INFO

Keywords:
Industrial production
Forecasting
Markov-switching
Mixed-data sampling
COVID-19

ABSTRACT

This article proposes different modelling approaches which exploit electricity market data to nowcast industrial production. Our models include linear, mixed-data sampling (MIDAS), Markov-Switching (MS) and MS-MIDAS regressions. Comparisons against autoregressive approaches and other commonly used macroeconomic predictors show that electricity market data combined with an MS model significantly improve nowcasting performance, especially during turbulent economic states, such as those generated by the recent COVID-19 pandemic. The most promising results are provided by an MS model which identifies two volatility regimes. These results confirm that electricity market data provide timely and easy-to-access information for nowcasting macroeconomic variables, especially when it is most valuable, i.e. during times of crisis and uncertainty.

1. Introduction

Most macroeconomic variables are released at monthly or quarterly intervals and with a significant delay, which typically goes from a month (e.g. Gross Domestic Product, GDP) to 40–45 days (e.g. industrial production index, IPI), and often further revised after release. This creates a substantial window of uncertainty for policy-makers and economic agents, particularly during periods of crisis, such as those created by the recent coronavirus (COVID-19) pandemic. In order to address this issue, researchers proposed several real-time indicators of economic activity, which are available in a more timely manner. This information has been used to develop so-called nowcasting approaches (e.g. Banbura et al., 2011; Andreou et al., 2013; Onorante and Raftery, 2016; Ravazzolo and Vespignani, 2020; Baumeister et al., 2022; Barbaglia et al., 2022). Recently, this literature has thrived in response to the need of tracking in real-time the various impacts of COVID-19. Several indicators were developed for this purpose, such as electricity consumption (Fezzi and Fanghella, 2020, 2021), consumers' transactions (Sheridan et al., 2020; Carvalho et al., 2021), mobile phone records (Goolsbee and Syverson, 2021), labour market trends (Forsythe et al., 2020; Kong and Prinz, 2020), or combinations of different information (Chetty et al., 2020; Foroni et al., 2020; Lewis et al., 2020).

This paper contributes to this growing literature by proposing and testing the forecasting performance of a set of models employing electricity market data for forecasting the seasonally and calendar adjusted IPI. IPI is a monthly indicator measuring the real output of the manufacturing, mining, electric and gas industries. It is commonly used to gauge the health of an economic system (Martínez-García et al., 2015), it is one of the most widely analysed macroeconomic variables (Heij et al., 2011; Schreiber and Soldatenkova, 2016; Chiu et al., 2017; Günay, 2018), and is a crucial input for forecasting the short-term evolution of GDP (Golinelli and Parigi, 2007; Baumeister and Guérin, 2021). Additionally, IPI is particularly difficult to forecast, since it reacts immediately to economic shocks and it is affected by both seasonal and cyclical variation (Bodo et al., 2000; Bruno and Lupi, 2004; Dendramis et al., 2020).

IPI forecasting methods proposed in the literature include linear regressions (Marchetti and Parigi, 2000; Franses and Van Dijk, 2005; Heij et al., 2011), neural networks (Heravi et al., 2004), singular spectrum analysis (Hassani et al., 2009, 2019), Markov-switching (MS) models (Billio et al., 2012), mixed-data sampling (MIDAS) models (Clements and Galvão, 2008) and vector autoregressive models (Bassetti et al., 2014; Schreiber and Soldatenkova, 2016; Chiu et al., 2017).

* Corresponding author.
E-mail addresses: giulio.galdi@unitn.it (G. Galdi), r.casarin@unive.it (R. Casarin), davide.ferrari2@unibz.it (D. Ferrari), carlo.fezzi@unitn.it (C. Fezzi), francesco.ravazzolo@unibz.it (F. Ravazzolo).

https://doi.org/10.1016/j.eneco.2023.107006
Received 16 May 2022; Received in revised form 7 August 2023; Accepted 28 August 2023
Available online 7 September 2023

Our approach recognises that most economic activities (including industrial production) require electricity as an input that is difficult to substitute away from, at least in the short-run. Indeed, electricity consumption (which, following the energy economics literature, we refer to as "load"), decreases significantly during weekends and public holidays, when many businesses are shut down (Fezzi and Bunn, 2010). Furthermore, information on load is publicly available in real time for almost all countries across the globe, since electricity is traded on an hourly or half-hourly basis in most developed economies. This feature makes our methodology widely applicable to different countries and contexts.

Here we compare different IPI forecasting models augmented with electricity market data against an AutoRegressive (AR) approach, which is commonly used as a benchmark (Heravi et al., 2004; Franses and Van Dijk, 2005), and a series of models employing other nowcasting predictors such as oil price, a stock market index, a business confidence indicator, and production expectations. Our baseline model specification is a simple linear regression using electricity consumption and its various determinants as explanatory variables, which we then enrich by introducing, in turn, a Markov-switching process (Hamilton, 1989) and a Mixed-Data Sampling specification. We then combine these two methodologies in an MS-MIDAS model.

We apply our approach to the Italian IPI and electricity data, and assess the robustness of our findings by extending the investigation to two other European countries: Spain and Germany. Our results can be briefly summarised as follows. First, most of our models which include electricity market information outperform the different benchmarks. Second, the most promising results are provided by the MS model which identifies two regimes of different volatility. In our application, the high volatility periods include both the 2008–2009 global financial crisis and the more recent recession triggered by the COVID-19 pandemic. It is during these turbulent periods that electricity market data give an edge to our approaches, which significantly outperform both the autoregressive specifications and the models using other predictors. On the other hand, during low volatility phases a simple AR model is not inferior to our proposed specifications. These results confirm that, in times of crisis, electricity data can provide crucial information for nowcasting macroeconomic indicators (Fezzi and Fanghella, 2020, 2021).

The remainder of the paper is organised as follows. Section 2 describes the data. Section 3 provides an overview of the different linear, nonlinear and mixed-frequency models that we consider for forecasting industrial production. Section 4 shows the results of our empirical study. Section 5 presents several extensions of our analysis. Section 6 concludes.

2. Data

Our data covers the period from January 2006 to December 2021. We downloaded our dependent variable, the latest release of IPI, from Eurostat, which publishes this information with roughly 40 days of delay.[1] For instance, the industrial production in January is available around the 10th of March of the same year. Electricity load data comes from the public platform of the European Network of Transmission System Operators for Electricity (ENTSO-E), where data is published on an hourly basis.[2] As mentioned in the introduction, our main case study is Italy, an important player in the European Union's industrial production landscape, accounting for 16% of the total output. Italy's power market is relatively young, having commenced operations in 2004. It primarily relies on two key sources to satisfy its annual demand of approximately 300 TWh (IEA, 2023): natural gas (about 50%), and renewables (40%). We construct our main explanatory variable by eliminating all weekends, in order to focus on the days in which most economic activity is carried out, and summing up the hourly information at the weekly level. In addition, we retain only the first four weeks to adjust for calendar effects (that is, the variation caused by months having a different number of working days).

Fig. 1 presents the time series of Italian IPI and load. The top panel displays the unadjusted IPI, which is characterised by strong seasonal patterns, with a noticeable dip in August, due to the summer holidays, and a smaller one in December, due to the winter festivities. The central panel shows the IPI after the seasonal and calendar adjustments performed by Eurostat.[3] This helps to highlight the slow but consistent fall in production following the global financial crisis in 2008 and 2009, and the V-shaped dramatic impact of the COVID-19 pandemic in 2020. As in Marchetti and Parigi (2000) and Bradley and Jansen (2004), we concentrate on this seasonal and calendar adjusted production index, which is significantly harder to forecast. Finally, the bottom panel of Fig. 1 presents the weekly average of hourly load. Seasonal effects are also visible in this time series, driven by both human activities (such as industrial production) and weather.

Temperature is by far the most important weather variable affecting load, with demand peaking during both winters and summers in order to satisfy heating and cooling needs (Behmiri et al., 2023; Chang et al., 2016). We represent temperature for the entire country as the mean between the average daily temperature of the two most populous Italian cities, Milan and Rome, retrieved from the National Oceanic and Atmospheric Administration.[4] Fig. 2 shows that the relationship between temperature and load is non-linear and V-shaped. Following the energy economics literature (Møller, 2017; Ahmed et al., 2018; Durmaz et al., 2020; Fezzi and Fanghella, 2021), we address this issue by calculating heating degree days (HDD) and cooling degree days (CDD) as the absolute difference with respect to 64 °F (about 18 °C), which is the value of temperature for which demand is at its minimum, defined by visually inspecting the scatter plot in Fig. 2. CDD/HDD represent temperature when electricity is used for cooling/heating. We calculate both variables on a daily basis and then aggregate them at the weekly and monthly levels.

Finally, in line with the IPI forecasting literature (Franses and Van Dijk, 2005; Heravi et al., 2004; Heij et al., 2011; Schreiber and Soldatenkova, 2016; Chiu et al., 2017), we treat our series as non-stationary and, therefore, develop all our models on the first differences. This choice is confirmed by the Augmented Dickey–Fuller test (Dickey and Fuller, 1979) which does not reject the null hypothesis of a unit root on the levels. For example, the test with intercept, trend and 5 lags generates a statistic of 1.88 with a corresponding p-value of 0.63. The same test strongly rejects the null on the first differences (test statistic −7.08, p-value < 0.01). Such results remain consistent with different test specifications.

3. Methodology

In this section, we first describe the different models considered in our analysis and then illustrate how we compare their forecasting performances.

[1] https://ec.europa.eu/eurostat/cache/metadata/en/sts_esms.htm#timeliness_punct1638197521702.

[2] https://transparency.entsoe.eu/.

[3] The statistical procedure for producing seasonal and calendar adjusted figures is illustrated by Eurostat in the metadata for industrial production: https://ec.europa.eu/eurostat/cache/metadata/en/sts_esms.htm#stat_process1651761756762. Each country is encouraged to provide Eurostat with seasonal and calendar adjusted figures. In case a country provides only unadjusted figures, Eurostat calculates the adjusted figure using the TRAMO/SEATS method in JDemetra+ software (Maravall et al., 2015).

[4] https://www.ncei.noaa.gov.

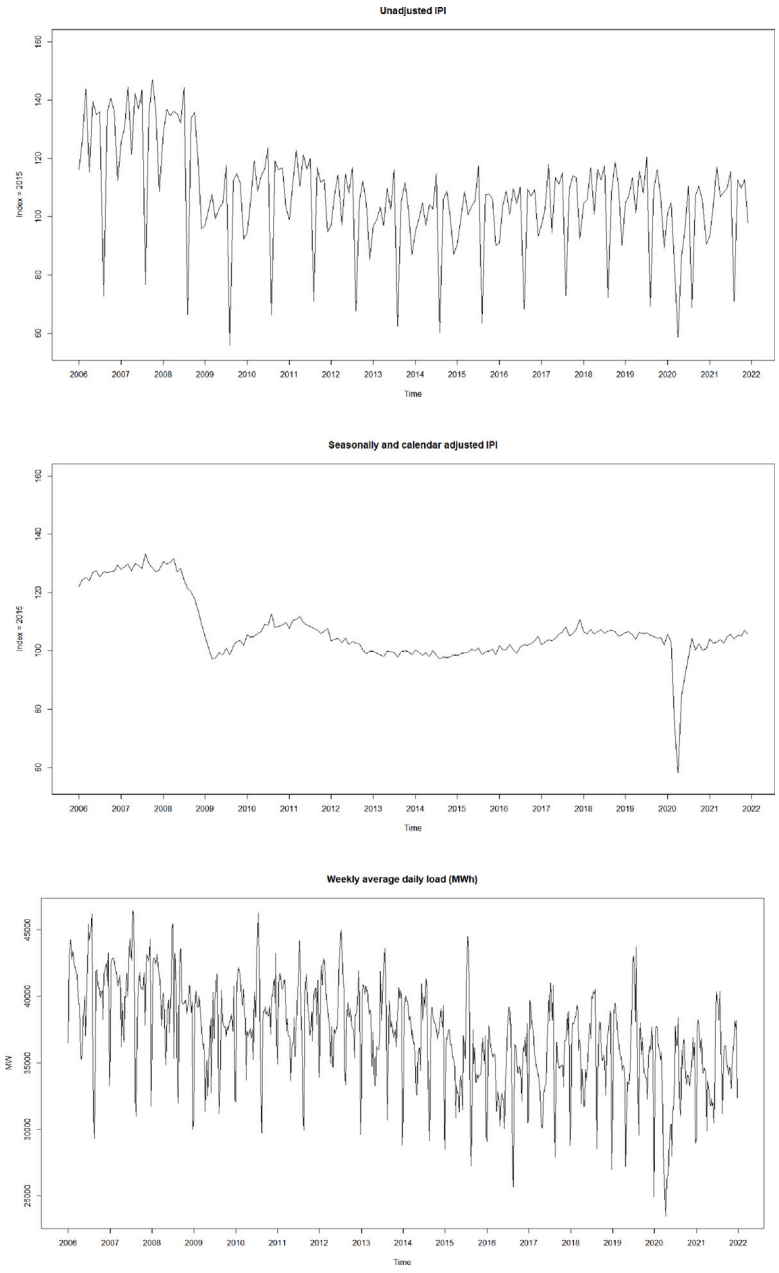

Fig. 1. Italian IPI and load time series. Notes: IPI in the second panel is seasonally and calendar adjusted by Eurostat using the TRAMO/SEATS method in the JDemetra+ software (Maravall et al., 2015). The bottom panel shows the load series, including only workdays.

3.1. Autoregressive benchmark

In line with previous research, we chose a simple AR model as the benchmark (Hassani et al., 2009; Heravi et al., 2004). The AR models assume that the value of a given variable depends on one or more of its past values, plus a stochastic error term that is assumed to be independent and identically distributed (*iid*). Indicating with y_t, $t = 1, \dots, T$ the sequence of IPI first differences, the AR(p) process can be written as:

$$y_t = \phi_0 + \phi_1 y_{t-1} + \cdots + \phi_p y_{t-p} + e_t, \tag{1}$$

where e_t, $t = 1, \dots, T$ is the *white noise* error term and $\phi_0, \dots, \phi_p$ are the autoregressive parameters that we estimate via Ordinary Least Squares (OLS). While we tested different values of p, $p = 1$ provides the best performance and, therefore, we choose an AR(1) as our benchmark.[5]

[5] In real applications, the lagged IPI would not be available at the end of the current month, for the reasons discussed at the beginning of Section 2. However, we also estimate an AR(2) with null first-lag coefficient, and it slightly underperforms the AR(1), so we decided to present the results of the AR(1) to give a slight advantage to the benchmark and to make our results more generalisable to other contexts.

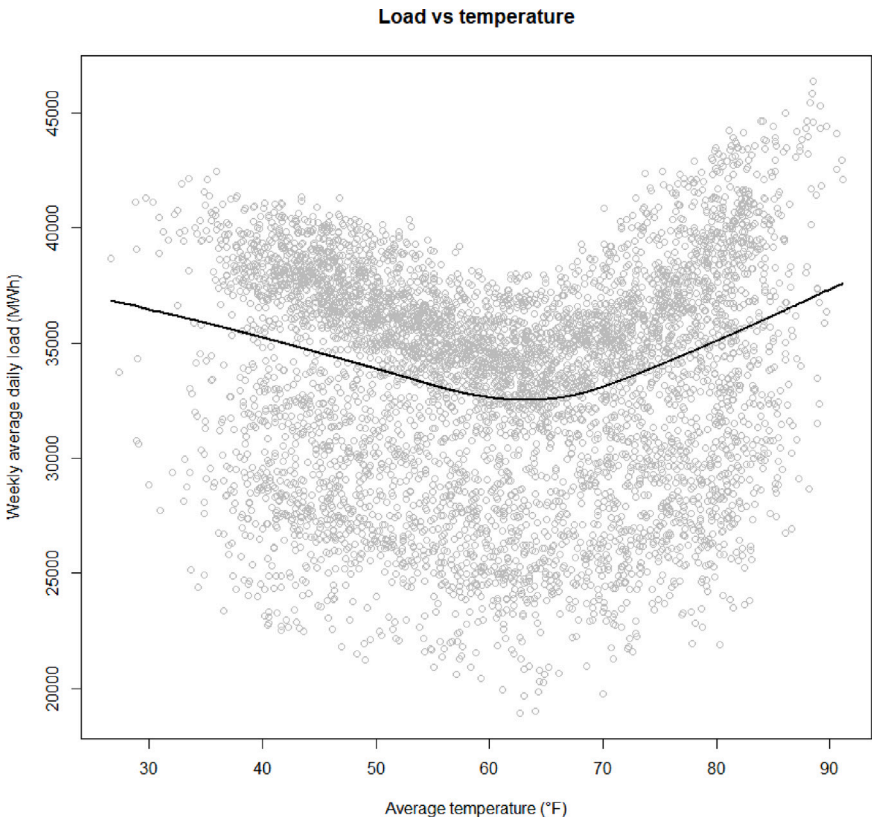

Fig. 2. Load vs. temperature data. Notes: The solid line is computed using a Lowess smoother (Cleveland, 1981), with a smoother span of $\frac{2}{3}$.

3.2. Linear load model

This approach assumes that IPI can be defined as a linear function of short term-adjusted load, i.e. load once the short-term drivers that do not depend on production have been controlled for. We control for such drivers, (i.e. seasonal, calendar, and temperature effects), because unadjusted load would not qualify as a good predictor. As the main determinants of load we include HDD, CDD, quarterly fixed-effects ($\beta_Q, Q = 1, \dots, 4$), and August and December fixed effects ($\beta_M, M = 8, 12$). Following Bulligan et al. (2010), the last two variables capture the effect of the typical periods of low activity in the Italian economy, due to the summer and Christmas vacations to account for the seasonality pattern in the Italian IPI series. In addition, we control for calendar effects by using only the first twenty working days for each month. Indicating with x_t the first differences of the load, we write the linear model as: $x_t = \beta_{Q,t} + \beta_{M,t} + \beta_1 HDD_t + \beta_2 CDD_t + u_t$, where HDD and CDD are also in first differences and u_t are the changes in load which are not caused by changes in temperature, seasonality, or months' length and, therefore, can be used to nowcast sudden changes in IPI. Assuming a linear relationship between y_t and u_t, we estimate the following linear load model (LLM):

$$\begin{aligned} y_t &= \gamma_0 + \gamma_1 u_t + e_t = \gamma_0 + \gamma_1 [x_t - \beta_{Q,t} - \beta_{M,t} - \beta_1 HDD_t - \beta_2 CDD_t] + e_t \\ &= \gamma_0 + \gamma_1 x_t - \theta_{Q,t} - \theta_{M,t} - \theta_1 HDD_t - \theta_2 CDD_t + e_t, \end{aligned} \tag{2}$$

where e_t, $t = 1, \dots, T$ is a *white noise* error term and $\gamma_0, \gamma_1, \theta_{Q,t}, \theta_{M,t}, \theta_1$, and θ_2 are the parameters we estimate with OLS. We also consider a specification of (2) which imposes $\theta_1 = \theta_2$, which, in effect, assumes the impacts of HDD and CDD to be the same. We refer to this specification as LLM-TDD, where TDD stands for Total Degree Days, that is the sum of HDD and CDD. Finally we also tried adding AR terms to these specifications, obtaining models that we indicate respectively with the acronyms AR-LLM and AR-TDD.

3.3. Markov-switching model

In this approach we allow the parameters in (2) to change according to a latent regime variable. We expect this variable to be able to separate "low- volatility" (normal and/or expansion periods) from "high-volatility" observations (crisis and/or recession periods). Similarly to Clements and Krolzig (1998), we assume that the coefficients and the standard deviation of the error term e_t change across regimes. The resulting Markov-switching model (MS) can be written as:

$$y_t = \gamma_0(s_t) + \gamma_1(s_t) x_t - \theta_{Q,t}(s_t) - \theta_{M,t}(s_t) - \theta_1(s_t) HDD_t - \theta_2(s_t) CDD_t + \sigma(s_t) e_t, \tag{3}$$

where all the parameters depend on the regime switching process s_t, $t = 1, \dots, T$ which is a m-states ergodic and aperiodic Markov-chain process. Since our framework recognises that industrial production is a function of load changes that do not depend on either seasonality or temperature effects, we let not only the parameter of load, but also those of temperature and seasonality to vary in each regime. We specify this model with two possible regimes, each representing a different state of the economy: expansion/low-volatility vs. crisis/high-volatility.[6] This latent process takes integer values and has transition probabilities $\mathbb{P}(s_t = j | s_{t-1} = i) = p_{ij}$, with $i, j \in \{1, 2\}$. In our applications, we assume that the initial values, $(y_{-p+1}, \dots, y_0)$, and s_0 are known. For estimation, we assume that the error terms are normally distributed and we estimate this model via maximum likelihood using the expectation–maximisation algorithm provided by the MSwM package (Sanchez-Espigares and Lopez-Moreno, 2021) for R (R Core

[6] Assuming more than two regimes increases proportionally the number of parameters to estimate, thus raising issues of overparametrisation. We tried to estimate a model with three regimes and the inference procedure had serious convergence issues.

Team, 2021). The algorithm first assigns a value to the latent variable, by randomly allocating each observation to one of the two regimes, effectively dividing our estimation sample in two sub-samples. After this step, the starting values for each regime are estimated via OLS and then the EM algorithm proceeds to iteratively maximise the likelihood following the approach illustrated by Perlin (2015). However, since the likelihood of this model is not globally concave, the algorithm can converge to a local rather than a global maximum. In order to address this issue, at each step of our forecasting comparison we run 30 estimates of model (3) with different starting values, and then select, at each step, the model which provides the highest likelihood. We choose the number of runs in order to make sure that our results were stable and reproducible, while maintaining a reasonable requirement of computational power.[7]

3.4. MIDAS model

Mixed-frequency models take advantage of the fact that the explanatory variables are released at much higher frequency than IPI. Electricity demand and temperature are, in fact, available on a daily or even hourly basis. Clearly, including such high-frequency information would require estimating too many parameters. However, we can still test if using weekly demand and temperature data (rather than monthly ones) leads to an improvement in forecasting performance. Following Foroni et al. (2015), in our context we can write the unrestricted MIDAS model as:

$$\omega(L)y_t = \gamma_0 + \gamma_1(L)\, x_t + \theta_1(L)\, HDD_t + \theta_2(L)\, CDD_t + \theta_{Q,t} + \theta_{M,t} + e_t, \quad (4)$$

where t indicates the weekly time periods, L is the lag operator, $\omega(L)$ is the aggregation scheme equal to $1+L+L^2+L^3$, $\gamma_1(L) = \gamma_{1,0}+\gamma_{1,1}L+\gamma_{1,2}L+\gamma_{1,3}L$, $\theta_1(L) = \theta_{1,0}+\theta_{1,1}L+\theta_{1,2}L+\theta_{1,3}L$, $\theta_2(L) = \theta_{2,0}+\theta_{2,1}L+\theta_{2,2}L+\theta_{2,3}L$, and all other symbols are defined as previously. This specification corresponds to (2) with 4 weekly load and temperature variables and 6 seasonal fixed effects, for a total of 18 parameters, which drop to 14 in the restricted version (MIDAS-TDD) where HDD and CDD have symmetric effects.

Finally, our last specification combines the MIDAS and the MS frameworks in an MS-MIDAS model with two regimes, as in Guérin and Marcellino (2013). We also consider an MS-MIDAS-TDD version.

3.5. Alternative predictors

Different IPI predictors have been proposed in the literature, including oil price (Schreiber and Soldatenkova, 2016), stock market indexes (Heij et al., 2011; Schreiber and Soldatenkova, 2016), business and consumer confidence indicators (Bruno and Lupi, 2004; Heij et al., 2011; Costantini, 2013), and production expectations (Bruno and Lupi, 2004; Lemmens et al., 2005). In order to compare the nowcasting performance of electricity market data against these other real-time indicators, we also estimate our models using the above-mentioned predictors. For each predictor, we estimate both linear specifications and Markov-switching ones. We retrieved information on the European Brent oil spot price from the US Energy Information Administration,[8] the FTSE-MIB stock market index from the Wall Street Journal,[9] and the business confidence indicator and the production expectations, both seasonally adjusted, are provided by Eurostat.[10]

[7] With 30 runs, estimating all our models with load requires about 50 h using a 2.40 GHz Intel i5 processor with 16 GB of RAM.

[8] https://www.eia.gov/dnav/pet/pet_pri_spt_s1_m.htm.

[9] https://www.wsj.com/market-data/quotes/index/IT/MTAA/I945/historical-prices.

[10] https://ec.europa.eu/eurostat/data/database.

Table 1
Forecasting performance.

	Overall 01/2014–12/2021		Calm 01/2014–02/2020 01/2021–12/2021		Turbulent 03/2020–12/2020	
	MAE	RMSE	MAE	RMSE	MAE	RMSE
Linear						
AR(1)	2.20	5.92	1.06	1.31	11.99	17.94
AR(1)-LLM	2.26	5.66	1.23	1.61	11.05	16.89
AR(1)-TDD	2.25	5.68	1.21	1.51	11.17	17.04
LLM	**2.13**	**5.06***	1.27	1.67	**9.55****	**14.90***
LLM-TDD	2.14	5.11*	1.25	1.57	9.78**	15.15*
MIDAS	2.18	5.11	1.31	1.65	9.74**	15.07*
MIDAS-TDD	2.21	5.18*	1.31	1.60	9.95**	15.34*
MS						
AR(1)	2.21	6.14	0.96*	1.21**	12.65	18.69
AR(1)-LLM	**1.94**	4.30	1.19	1.54	8.44	12.54
AR(1)-TDD	2.33	5.66	1.14	1.47	12.54	16.99
MS	2.62	6.36	1.37	1.88	13.42	18.91
MS-TDD	2.02	**4.26**	1.37	1.78	**7.60**	**12.13**
MS-MIDAS	3.36	7.77	1.68	2.36	17.79	23.06
MS-MIDAS-TDD	3.03	6.46	1.73	2.31	14.17	18.83

Notes: In the table above we show the Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) for the whole forecasting window and the "calm", and the "turbulent" periods. In bold we highlight the best performing model. Stars indicate that the one-sided DM test (Diebold and Mariano, 2002) is significant at the 10% (*) or 5% (**) significance level.

3.6. Forecasting evaluation

For all our models, we perform one-step ahead forecasts with a recursive forecasting window, i.e. the estimation window increases by one month at each step so that the prediction for each month exploits all the information available at that time. The interval that we use to compare the forecasting performance is the last eight years, which corresponds to half of our overall sample. We divide the forecasting sample between a "calm" and a "turbulent" period. The turbulent period includes the months of economic crisis, which in our sample corresponds to those following the COVID-19 outbreak, i.e. from March to December 2020. The remaining of the forecasting sample is classified as the calm period. We note that this distinction does not serve as an input for the identification of the regimes in the MS models, but is simply used to compare the forecasting performance of our models in two different states of the economy.

As most of the energy (Wang et al., 2022; Lehna et al., 2022) and IPI (Hassani et al., 2009; Bulligan et al., 2010) forecasting literature, we evaluate the performance of our models in terms of Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE), defined as:

$$MAE = \frac{\sum_t |\hat{y}_t - y_t|}{T}, \qquad RMSE = \sqrt{\frac{\sum_t (\hat{y}_t - y_t)^2}{T}},$$

where y is the observed IPI at time $t = 1, \dots, T$, and $\hat{y}_t$ is its value as predicted by our models.

4. Results

In this section we report in detail the results for the Italian IPI, whereas an extension to other countries is presented in the next section. Table 1 summarises the forecasting performance of our models on the whole 2014–2021 testing sample, distinguishing between the calm and turbulent periods.[11] The table shows how electricity demand provides important information for IPI predictions. In the top half of the table

[11] Estimating at different points in time can lead the Markov-switching model to allocate periods to either low- or high-volatility regimes in different ways.

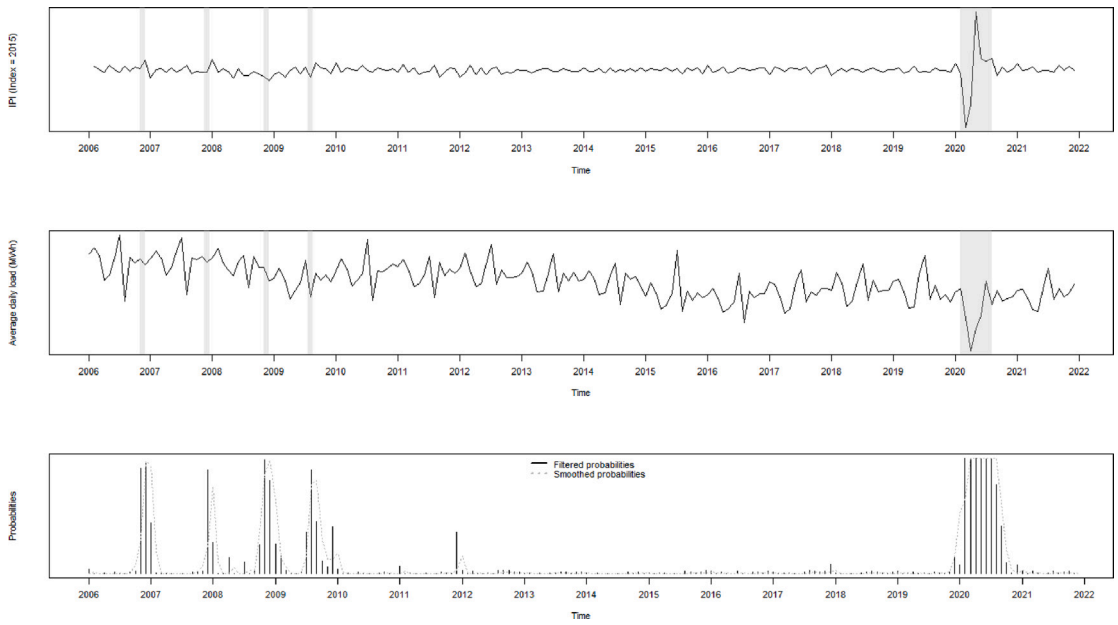

Fig. 3. Markov-switching regimes. Notes: The first two panels show the first differences of IPI and load between 2006 and 2021 (solid line). The grey bands identify the periods of high-volatility, according to the probabilities of being in the high-volatility regime shown in the bottom panel.

we show the forecasting performance, in terms of MAE and RMSE, of our linear models, whereas in the bottom half we present the performance of the Markov-switching specifications. Considering the entire sample, in the first two columns, all the static linear models that include electricity market information significantly outperform the AR(1) benchmark. Columns 3 to 6 show that this superior performance is concentrated during the turbulent months, and that in the calm period all the approaches are not significantly different from each other. The models that combine the AR component with electricity market data (AR(1)-LLM and AR(1)-TDD) do not perform better than the simpler AR(1).

Moving to the bottom half of the table, columns 1–2 show that, depending on the indicator, the model with the best overall forecasting performance is either the AR(1)-LLM or the MS-TDD. However, columns 5–6 show that the forecasting performance of MS-TDD is clearly superior during the turbulent period. In addition, we note that, in general, our most complex models (in terms of parameter numbers) do not seem to perform particularly well. For example, all MS-MIDAS present MAE and RMSE that are higher than those provided by the AR(1). Our interpretation of this result is that our monthly data does not contain enough information to generate precise parameter estimates of such complex specifications. For example, our entire sample includes 192 monthly observations and the MS-MIDAS includes 38 parameters, which appear to be too many. By the same token, the standard version of each model is always outperformed by its TDD counterpart. In other words, assuming the impact of HDD and CDD to be the same and, thereby, reducing the number of parameters constitutes an advantage, at least in our case-study.

In Table 2 we show the MS-TDD estimates over our entire sample (2006–2021). The two regimes identified by this model are indeed clearly distinct, with the load-related parameters are all non-significantly different from zero during the low- volatility regime, while they become significantly different during the high-volatility one. Signs are as expected from Eq. (2), with the parameter of load being positive and the parameter of degree days being negative. Finally, the standard deviation of the error term of the high volatility regime is more than two times the one of the low-volatility one.

We display the two regimes in Fig. 3, where the grey bands represent the periods in which the system is in its high-volatility regime, i.e. when the smoothed probability (displayed in the bottom panel) of that regime is higher than 0.5. The model identifies a few months during the 2008–2009 great financial crisis and the first wave of COVID-19 as high-volatility regimes, reassuring us on its ability to detect the most important periods of crisis that the economic system experienced in the last decade. Analogously to the linear specification, we also test a MS model in which HDD and CDD are constrained to have symmetric effects (MS-TDD), and several combinations with an AR component.

Table 2
MS-TDD estimation results.

	Regime 1 (Low volatility)		Regime 2 (High volatility)	
Intercept	0.03	0.28	2.14	2.93
Load	0.03	0.08	5.24***	0.51
TDD	0.00	0.00	−0.03**	0.01
Spring	−0.13	0.36	4.30	3.75
Summer	−0.48	0.51	−22.35***	4.85
Winter	−0.02	0.35	−7.68**	3.61
August	0.76	1.09	48.44***	6.37
December	0.06	0.59	8.90**	4.04
R^2	0.01		0.89	
Std. deviation	1.55		4.07	
Log-likelihood:	−393.71			

Notes: In the table above, we expressed load in terms of GWh. Stars indicate significance as follows: * 10%, ** 5%, *** 1%.

We further investigate the relative performance of our models through a series of fluctuation tests, proposed by Giacomini and Rossi (2010). This test measures the relative performance of two series of out-of-sample forecasts, computing a statistic analogous to the DM statistic over a rolling window. In this way, the test allows us to monitor how the relative performance of our approaches vary along the forecasting sample. We set the rolling window to 10% of the forecasting sample, that is roughly 10 months. Fig. 4 compares our best model (MS-TDD) against the AR(1) and the MS-AR(1) models. For positive values of the statistic, our model outperforms the benchmarks, and vice versa. The plot shows that the forecasting performance of our approach dramatically improves from the onset of the COVID-19 pandemic and confirms that such improvement remains significant during the entire high-volatility period.

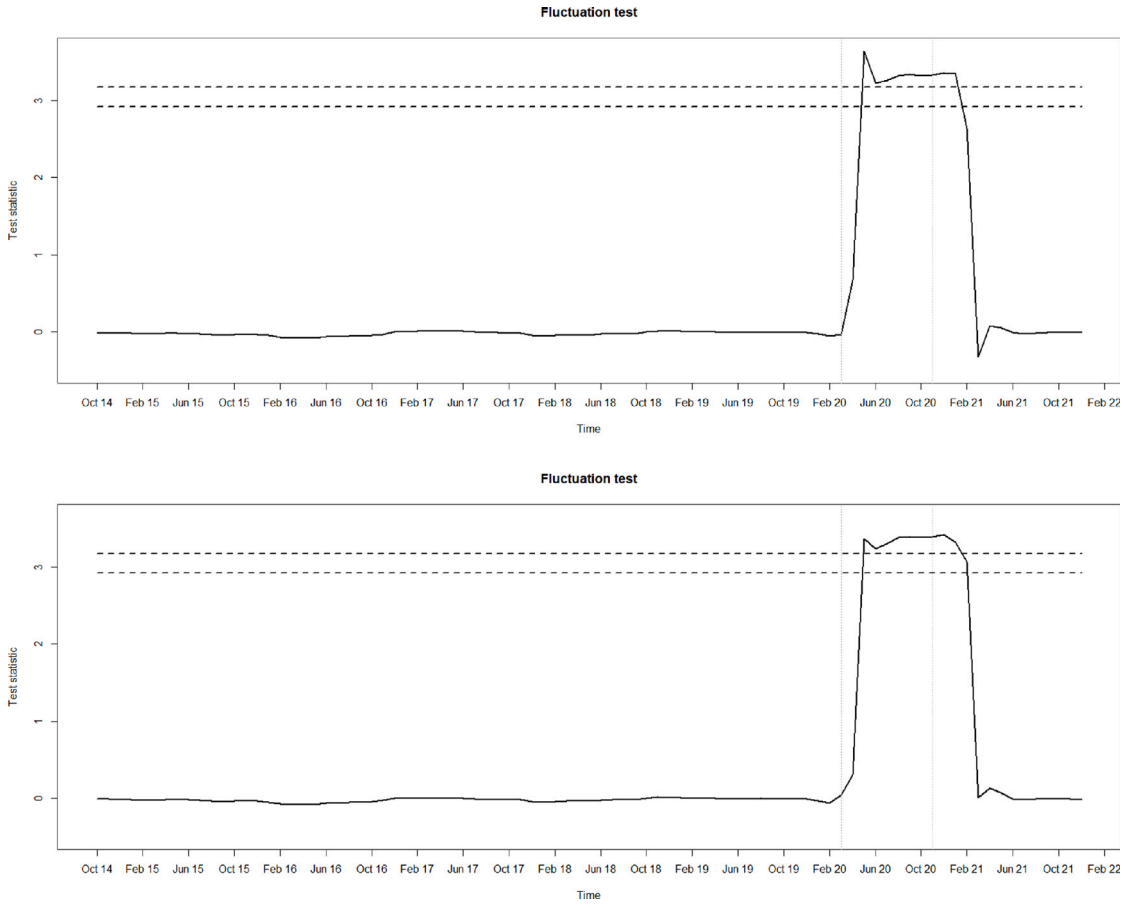

Fig. 4. Fluctuation test — MS-TDD relative to AR(1). Notes: In the figures above, we show the value of the test statistic (solid lines) along our forecasting sample. In the top panel, the benchmark model is the linear AR(1), whereas in the bottom panel it is the MS-AR(1). When the lines lie above the zero, they indicate that our model is performing better than the benchmark, and vice versa. The dotted lines are the asymptotic critical values for the statistic at 5% and 10% significance level, as provided by Giacomini and Rossi (2010).

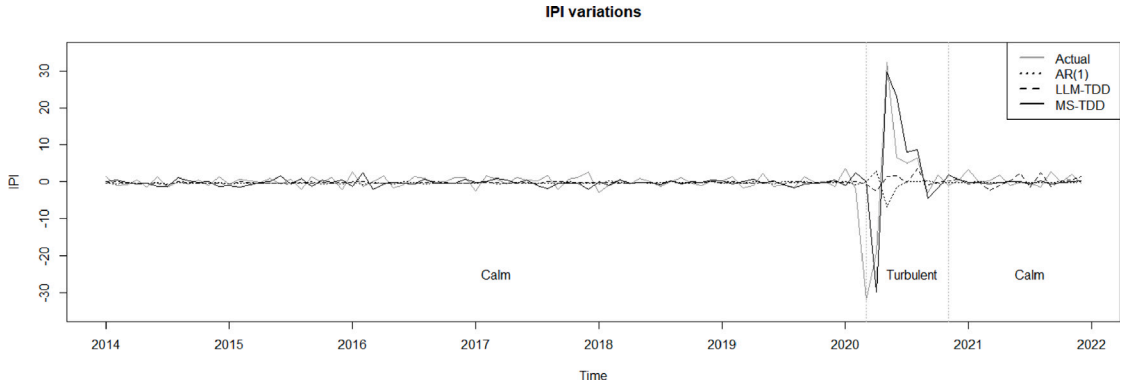

Fig. 5. Actual and forecasted IPI. Notes: Comparison of actual IPI first differences against AR(1), LLM-TDD, and MS-TDD forecasts.

Finally, Fig. 5 provides a visual comparison amongst the forecasts generated by three selected models: (1) our benchmark model, i.e. the AR(1); (2) the linear model using electricity market data (LLM-TDD) and (3) our best performing model (the MS-TDD). During calm periods, all models generate similar results. On the other hand, the performance of the AR(1) significantly deteriorates during turbulent periods, and in particular during the first half of 2020, when industrial production significantly dropped in response to the COVID-19 pandemic and related lockdown measures. Load-based models appear to perform better, in particular the MS-TDD. At first, this model does not react promptly to the onset of COVID-19, and fails to nowcast the sudden IPI drop in March 2020. However, in the following months it behaves remarkably well, outperforming both the LLM-TDD and the benchmark. Recalling Table 2, by the end of our sample this model has "learned" to distinguish two distinct regimes and to exploit electricity markets' information to nowcast IPI during high-volatility phases. On these grounds, we expect that this model will perform even better in the event of future economic crises.

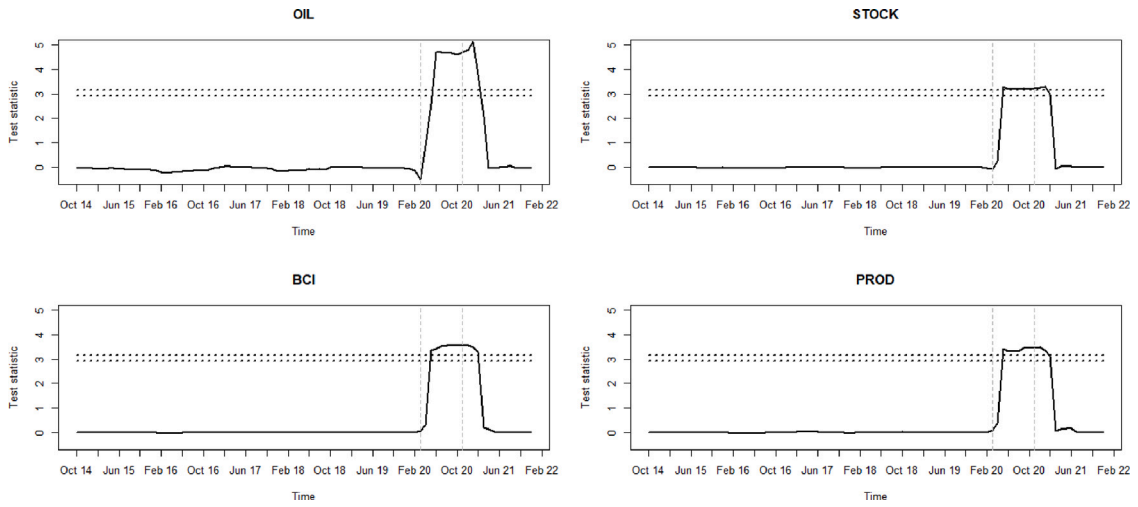

Fig. 6. Actual and forecasted IPI. Notes: The panels above show the fluctuation test statistic (solid line) for the MS models using the alternative predictors, with our MS-TDD model as the benchmark. The solid lines represent the test statistics and are above zero when the MS-TDD model outperforms the alternative one whereas the opposite holds below zero. The dotted lines are the asymptotic critical values for the statistic at 5% and 10% significance level, as provided by Giacomini and Rossi (2010).

Table 3
Forecasting performance of alternative predictors.

	Overall 01/2014–12/2021		Calm 01/2014–02/2020 01/2021–12/2021		Turbulent 03/2020–12/2020	
	MAE	RMSE	MAE	RMSE	MAE	RMSE
Linear						
LOAD	2.13	**5.06**	1.27	1.67	**9.55**	**14.90**
OIL	**2.07**	5.07	1.14	1.39	10.10	15.19*
STOCK	2.09	5.25	1.10	1.36	10.63**	15.77*
BCI	2.18	5.74	1.12	1.37	11.30*	17.33
PROD	2.14	5.31	1.12	1.35	10.96***	15.98*
MS						
LOAD	**2.02**	**4.26**	1.37	1.78	**7.60**	**12.13**
OIL	2.05	5.05*	1.10	1.37	10.25*	15.11*
STOCK	2.33	6.67	1.09	1.37	12.97	20.27
BCI	2.78	8.51	1.14	1.45	16.84*	26.03
PROD	2.55	7.19	1.21	1.56	14.12	21.81

Notes: In the table above, we also show the results of DM tests (Diebold and Mariano, 2002) of alternative predictors. Stars indicate that the model is significantly outperformed by our best load model, as follows: * 10%, ** 5%, *** 1%.

Table 4
Forecasting performance of load models in the three countries.

	Overall 01/2014–12/2021		Calm 01/2014–02/2020 01/2021–12/2021		Turbulent 03/2020–12/2020	
	MAE	RMSE	MAE	RMSE	MAE	RMSE
Italy						
AR(1)	2.20	5.92	1.06	1.31	11.99	17.94
LLM	2.13	5.06*	1.27	1.67	9.55**	14.90*
TDD	2.14	5.11*	1.25	1.57	9.78**	15.15*
MS	2.62	6.36	1.37	1.88	13.42	18.91
MS-TDD	**2.02**	**4.26**	1.37	1.78	**7.60**	**12.13**
Germany						
AR(1)	1.64	3.33	1.07	1.40	6.48	9.47
LLM	1.60	2.85*	1.17	1.44	5.32*	7.75*
TDD	**1.59**	2.84*	1.17	1.43	5.27*	7.75*
MS	1.64	**2.53***	1.29	1.67	**4.65***	**6.13****
MS-TDD	1.63	2.69*	1.24	1.61	4.98	6.85*
Spain						
AR(1)	1.57	3.91	0.81	1.06	8.10	11.70
LLM	1.45	3.08	0.86	1.09	6.51	9.00
TDD	1.45	3.09	0.85	1.07	6.55	9.04
MS	1.36	2.92	0.84	1.14	5.80	8.41
MS-TDD	**1.29**	**2.63***	0.81	1.07	**5.47**	**7.53***

Notes: In the table above, bold figures indicate that the row model outperforms the AR benchmark. Stars indicate that the alternative predictor is significantly better according to the DM test (Diebold and Mariano, 2002), as follows: * 10%, ** 5%, *** 1%.

5. Extensions

In this section we present the results of three extensions of our analysis. First, we compare electricity consumption to the other nowcasting predictors commonly used in the literature, introduced in Section 3.5. Second, we generalise our finding to two other European countries (Germany and Spain) and, finally, we evaluate alternative specifications for temperature and seasonality effects.

5.1. Comparing against alternative predictors

To compare load against other commonly used nowcasting indicators, we estimate all our models using oil price, stock market index, business confidence indicator and production expectations. Table 3 summarises the results for the best-performing models using these new predictors. The top half of the table illustrates the results for the linear specifications. It shows that, when considering the whole sample, either oil price or load are the best predictor, according to whether MAE or RMSE is used to evaluate the forecasting performance (columns 1–2). However, when focusing on the turbulent period, we observe that load performs significantly better than all other predictors. The bottom half of the table presents the MS version of the models. While the performance of most nowcasting indicators does not seem to change considerably (and in some cases even deteriorates), the one of load significantly improves. This leads to forecasting errors which are significantly smaller than those provided by the other approaches.

Table 5
Forecasting performance of alternative specifications.

	Overall 01/2014–12/2021		Calm 01/2014–02/2020 01/2021–12/2021		Turbulent 03/2020–12/2020	
	MAE	RMSE	MAE	RMSE	MAE	RMSE
AR(1)	2.20	5.20	1.06	1.31	11.99	17.94
Linear						
LLM-T^2	**2.15**	**5.15**	1.25	1.55	**9.92**	**15.30**
LLM-SINCOS	**2.16**	5.28	1.16	1.44	**10.77**	**15.81**
LLM-SINCOS-TDD	**2.15**	5.31	1.16	1.41	**10.71**	**15.94**
MS						
MS-T^2	2.29	4.70	1.28	1.68	**10.96**	**13.70**
MS-SINCOS	2.52	6.00	1.16	1.50	14.25	18.05
MS-SINCOS-TDD	4.35	12.61	1.12	1.39	32.17	38.85

Notes: We here show the Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) for the models estimated as robustness checks, computed for the whole forecasting window and the "calm", and the "turbulent" periods.

Fig. 6 compares the best performing model (MS-TDD) against the same specification using each alternative predictor by means of fluctuation tests. We observe that, during the calm period, all predictors achieve comparable results, yet the MS-TDD outperforms all alternative predictors at the 5% significance level during the entire year of the COVID-19 crisis.

5.2. Extension to other countries

We also implement our analysis in two other European countries, namely Germany and Spain, in order to test the validity of our results in contexts that are different from the Italian one. Germany has the largest industrial production in the European Union (almost two times the one of Italy) and a power system in which coal (38%) and renewables (35%) have similar generation shares (IEA, 2020). On the other hand, Spain has a relatively smaller industrial production (half of the Italian one) and a generation mix which includes renewables (38%) and natural gas (33%) in almost equal parts, but also a relevant contribution from nuclear plants (22%, IEA, 2021).

Table 4 shows that the MS-TDD outperforms the benchmark and the other specifications also for these countries (columns 1–2), with this conclusion driven by the superior performance during the turbulent period (columns 5–6). It is worth noting that the results for these countries are somewhat stronger than those from Italy, with the DM tests being significant when considering both the whole forecasting sample (column 2) and only the turbulent period (column 6).

5.3. Further robustness tests

We also employed a few different specifications in order to further test the robustness of our results. Those models are summarised in Table 5. Regarding the treatment of temperature, we replaced degree days with a quadratic function of the first differences of temperature: LLM-T^2. We also tried to reduce the number of seasonal parameters by replacing the seasonal fixed effects with a sinusoidal function including sine and cosine with period one year (LLM-SINCOS). For all these models we considered both the linear and the MS specifications. As shown in the table, these approaches underperformed in terms of both MAE and RMSE with respect to the corresponding models presented in rest of the manuscript. Further information is available upon request.

6. Conclusions

We developed and compared the performance of a set of models employing electricity consumption, temperature, and seasonal effects to nowcast IPI. Including this information provides predictions which are superior to those of a series of alternative specifications, which include autoregressive models and alternative predictors, especially during the recent COVID-19 pandemic. To the extent that future crises resemble in nature the ones included in our sample (e.g., COVID-19, global financial crisis, sovereign debt crisis), our approach is likely to provide more accurate forecasts with respect to commonly used approaches. These results are of particular relevance to policy-makers and economic agents, since periods of crisis are those characterised by the most economic uncertainty.

Specifically, our comparison points towards the MS model augmented with electricity consumption and total degree days as the best performing model. This specification identifies two regimes, one associated to low-volatility periods and one associated to high-volatility ones. In our sample, the model autonomously identifies this second regime with several months during the global financial crisis in 2008–2009 and during the more recent COVID-19 pandemic. Electricity load information is relevant for predictions only during such volatile states, while it does not contain any significant information during the calmer months. These results confirm that electricity market outcomes provide crucial information for nowcasting macroeconomic variables in times of crisis (Fezzi and Fanghella, 2020, 2021).

Future research could investigate whether the use of more geographically refined data improves the forecasting performance of our model or provides insights into regional heterogeneity. In addition, our single-equation approach could be extended to multivariate settings. In particular, since IPI is one of the main variables included in dynamic macroeconomic models (Clements and Galvão, 2008; Kuzin et al., 2011; Cross et al., 2021), including in such methods an additional equation representing electricity market outcomes seems a promising strategy for boosting their forecasting (and nowcasting) performance. Such systems could be set up as Markov-switching VAR models, possibly with mixed-frequency data (Schorfheide and Song, 2015). Moreover, interactions among both real and energy markets from different countries could be accounted for by estimating a mixed-frequency Bayesian Markov-switching model (Casarin et al., 2018) in order to study how shocks propagate amongst different economies. In general, electricity market outcomes appear to provide easy-to-access, timely and relevant information for improving how we understand and forecast some of the most important macroeconomic indicators.

CRediT authorship contribution statement

Giulio Galdi: Methodology, Software, Formal analysis, Validation, Data curation, Writing, Visualization. **Roberto Casarin:** Writing, Funding Acquisition. **Davide Ferrari:** Writing, Funding acquisition. **Carlo Fezzi:** Conceptualization, Methodology, Software, Formal analysis, Validation, Writing, Supervision, Funding acquisition. **Francesco Ravazzolo:** Conceptualization, Methodology, Software, Writing, Supervision, Funding acquisition.

Funding

This research was supported by the OCEN–COVID19 project, funded by the University of Trento under the "Covid 19" research fund and by High-frequency Economic Indicators and Resilience of Society (HEIRS) founded by the FISR-MIUR.

References

Ahmed, T., Vu, D.H., Muttaqi, K.M., Agalgaonkar, A.P., 2018. Load forecasting under changing climatic conditions for the city of Sydney, Australia. Energy 142, 911–919.

Andreou, E., Ghysels, E., Kourtellos, A., 2013. Should macroeconomic forecasters use daily financial data and how? J. Bus. Econom. Statist. 31 (2), 240–251.

Banbura, M., Giannone, D., Reichlin, L., 2011. Nowcasting. In: Clements, M.P., Hendry, D.F. (Eds.), The Oxford Handbook of Economic Forecasting. Oxford Handbooks in Economics, p. 18.

Barbaglia, L., Frattarolo, L., Onorante, L., Pericoli, F.M., Ratto, M., Tiozzo Pezzoli, L., 2022. Testing big data in a big crisis: Nowcasting under Covid-19. Int. J. Forecast..

Bassetti, F., Casarin, R., Leisen, F., 2014. Beta-product dependent pitman–yor processes for Bayesian inference. J. Econometrics 180 (1), 49–72.

Baumeister, C., Guérin, P., 2021. A comparison of monthly global indicators for forecasting growth. Int. J. Forecast. 37 (3), 1276–1295.

Baumeister, C., Leiva-León, D., Sims, E., 2022. Tracking weekly state-level economic conditions. Rev. Econ. Stat. 1–45.

Behmiri, N.B., Fezzi, C., Ravazzolo, F., 2023. Incorporating air temperature into mid-term electricity load forecasting models using time-series regressions and neural networks. Energy 278, 127831.

Billio, M., Casarin, R., Ravazzolo, F., van Dijk, H.K., 2012. Combination schemes for turning point predictions. Q. Rev. Econ. Finance 52 (4), 402–412.

Bodo, G., Golinelli, R., Parigi, G., 2000. Forecasting industrial production in the euro area. Empir. Econ. 25 (4), 541–561.

Bradley, M.D., Jansen, D.W., 2004. Forecasting with a nonlinear dynamic model of stock returns and industrial production. Int. J. Forecast. 20 (2), 321–342.

Bruno, G., Lupi, C., 2004. Forecasting industrial production and the early detection of turning points. Empir. Econ. 29 (3), 647–671.

Bulligan, G., Golinelli, R., Parigi, G., 2010. Forecasting industrial production: the role of information and methods. In: The IFC's Contribution To the 57th ISI Session, Durban, Vol. 33. Bank for International Settlements, pp. 227–235.

Carvalho, V.M., Garcia, J.R., Hansen, S., Ortiz, Á., Rodrigo, T., Rodríguez Mora, J.V., Ruiz, P., 2021. Tracking the COVID-19 crisis with high-resolution transaction data. R. Soc. Open Sci. 8 (8), 210218.

Casarin, R., Foroni, C., Marcellino, M., Ravazzolo, F., 2018. Uncertainty through the lenses of a mixed-frequency Bayesian panel Markov-switching model. Ann. Appl. Stat. 12 (4), 2559–2586.

Chang, Y., Kim, C.S., Miller, J.I., Park, J.Y., Park, S., 2016. A new approach to modeling the effects of temperature fluctuations on monthly electricity demand. Energy Econ. 60, 206–216.

Chetty, R., Friedman, J.N., Hendren, N., Stepner, M., Team, T.O.I., 2020. How did COVID-19 and Stabilization Policies Affect Spending and Employment? A New Real-Time Economic Tracker Based on Private Sector Data. National Bureau of Economic Research Cambridge, MA.

Chiu, C.-W.J., Mumtaz, H., Pinter, G., 2017. Forecasting with VAR models: Fat tails and stochastic volatility. Int. J. Forecast. 33 (4), 1124–1143.

Clements, M.P., Galvão, A.B., 2008. Macroeconomic forecasting with mixed-frequency data. J. Bus. Econom. Statist. 26, 546–554.

Clements, M.P., Krolzig, H.M., 1998. A comparison of the forecast performances of Markov-switching and threshold autoregressive models of US GNP. Econom. J. 1, C47–C75.

Cleveland, W.S., 1981. LOWESS: A program for smoothing scatterplots by robust locally weighted regression. Amer. Statist. 35 (1), 54.

Costantini, M., 2013. Forecasting the industrial production using alternative factor models and business survey data. J. Appl. Stat. 40 (10), 2275–2289.

Cross, J.L., Hou, C., Nguyen, B.H., 2021. On the China factor in the world oil market: A regime switching approach. Energy Econ. 95, 105119.

Dendramis, Y., Kapetanios, G., Marcellino, M., 2020. A similarity-based approach for macroeconomic forecasting. J. R. Statist. Soc.: Ser. A (Statist. Soc.) 183 (3), 801–827.

Dickey, D.A., Fuller, W.A., 1979. Distribution of the estimators for autoregressive time series with a unit root. J. Amer. Statist. Assoc. 74 (366a), 427–431.

Diebold, F.X., Mariano, R.S., 2002. Comparing predictive accuracy. J. Bus. Econom. Statist. 20 (1), 134–144.

Durmaz, T., Pommeret, A., Tastan, H., 2020. Estimation of residential electricity demand in Hong Kong under electricity charge subsidies. Energy Econ. 88, 104742.

Fezzi, C., Bunn, D., 2010. Structural analysis of electricity demand and supply interactions. Oxford Bull. Econom. Statist. 72 (6), 827–856.

Fezzi, C., Fanghella, V., 2020. Real-time estimation of the short-run impact of COVID-19 on economic activity using electricity market data. Environ. Resour. Econom. 76 (4), 885–900.

Fezzi, C., Fanghella, V., 2021. Tracking GDP in real-time using electricity market data: Insights from the first wave of COVID-19 across Europe. Eur. Econ. Rev. 139, 103907.

Foroni, C., Marcellino, M., Schumacher, C., 2015. U-MIDAS: MIDAS regressions with unrestricted lag polynomials. J. R. Statist. Soc.: Ser. A (Statist. Soc.) 29 (1), 57–82.

Foroni, C., Marcellino, M., Stevanovic, D., 2020. Forecasting the COVID-19 recession and recovery: Lessons from the financial crisis. Int. J. Forecast..

Forsythe, E., Kahn, L.B., Lange, F., Wiczer, D., 2020. Labor demand in the time of COVID-19: Evidence from vacancy postings and UI claims. J. Public Econom. 189, 104238.

Franses, P.H., Van Dijk, D., 2005. The forecasting performance of various models for seasonality and nonlinearity for quarterly industrial production. Int. J. Forecast. 21 (1), 87–102.

Giacomini, R., Rossi, B., 2010. Forecast comparisons in unstable environments. J. Appl. Econometrics 25 (4), 595–620.

Golinelli, R., Parigi, G., 2007. The use of monthly indicators to forecast quarterly GDP in the short run: an application to the G7 countries. J. Forecast. 77–94.

Goolsbee, A., Syverson, C., 2021. Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020. J. Public Econom. 193, 104311.

Guérin, P., Marcellino, M., 2013. Markov-switching MIDAS models. J. Bus. Econom. Statist. 31 (1), 45–56.

Günay, M., 2018. Forecasting industrial production and inflation in Turkey with factor models. Cent. Bank Rev. 18 (4), 149–161.

Hamilton, J., 1989. A new approach to the economic analysis of nonstationary time series and the business cycle. Econometrica 57 (2), 357–384.

Hassani, H., Heravi, S., Zhigljavsky, A., 2009. Forecasting European industrial production with singular spectrum analysis. Int. J. Forecast. 25 (1), 103–118.

Hassani, H., Rua, A., Silva, E.S., Thomakos, D., 2019. Monthly forecasting of GDP with mixed-frequency multivariate singular spectrum analysis. Int. J. Forecast. 35 (4), 1263–1272.

Heij, C., van Dijk, D., Groenen, P.J., 2011. Real-time macroeconomic forecasting with leading indicators: scan empirical comparison. Int. J. Forecast. 27 (2), 466–481.

Heravi, S., Osborn, D.R., Birchenhall, C., 2004. Linear versus neural network forecasts for European industrial production series. Int. J. Forecast. 20 (3), 435–446.

IEA, 2020. Germany 2020 Energy Policy Review. Technical Report, International Energy Agency.

IEA, 2021. Spain 2021 Energy Policy Review. Technical Report, International Energy Agency.

IEA, 2023. Italy 2023 Energy Policy Review. Technical Report, International Energy Agency.

Kong, E., Prinz, D., 2020. Disentangling policy effects using proxy data: Which shutdown policies affected unemployment during the COVID-19 pandemic? J. Public Econom. 189, 104257.

Kuzin, V., Marcellino, M., Schumacher, C., 2011. MIDAS vs. mixed-frequency VAR: Nowcasting GDP in the euro area. Int. J. Forecast. 27 (2), 529–542.

Lehna, M., Scheller, F., Herwartz, H., 2022. Forecasting day-ahead electricity prices: A comparison of time series and neural network models taking external regressors into account. Energy Econ. 106, 105742.

Lemmens, A., Croux, C., Dekimpe, M.G., 2005. On the predictive content of production surveys: A pan-European study. Int. J. Forecast. 21 (2), 363–375.

Lewis, D., Mertens, K., Stock, J.H., 2020. US Economic Activity During the Early Weeks of the SARS-Cov-2 Outbreak. Technical Report, National Bureau of Economic Research.

Maravall, A., Gómez, V., Caporello, G., et al., 2015. Statistical and econometrics software: TRAMO and SEATS. Statist. Econometr. Softw..

Marchetti, D.J., Parigi, G., 2000. Energy consumption, survey data and the prediction of industrial production in Italy: A comparison and combination of different models. J. Forecast. 19 (5), 419–440.

Martínez-García, E., Grossman, V., Mack, A., 2015. A contribution to the chronology of turning points in global economic activity (1980–2012). J. Macroeconom. 46, 170–185.

Møller, N.F., 2017. Energy demand, substitution and environmental taxation: An econometric analysis of eight subsectors of the Danish economy. Energy Econ. 61, 97–109.

Onorante, L., Raftery, A.E., 2016. Dynamic model averaging in large model spaces using dynamic Occam's window. Eur. Econ. Rev. 81, 2–14.

Perlin, M., 2015. MS_Regress-the Matlab package for Markov regime switching models. Available at SSRN 1714016.

R Core Team, 2021. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria, URL: https://www.R-project.org/.

Ravazzolo, F., Vespignani, J., 2020. World steel production: A new monthly indicator of global real economic activity. Can. J. Econom./Rev. Can. d'Économ. 53 (2), 743–766.

Sanchez-Espigares, J.A., Lopez-Moreno, A., 2021. MSwM: Fitting Markov switching models. R package version 1.5. URL: https://CRAN.R-project.org/package=MSwM.
Schorfheide, F., Song, D., 2015. Real-time forecasting with a mixed-frequency VAR. J. Bus. Econom. Statist. 33 (3), 366–380.
Schreiber, S., Soldatenkova, N., 2016. Anticipating business-cycle turning points in real time using density forecasts from a VAR. J. Macroeconom. 47, 166–187.
Sheridan, A., Andersen, A.L., Hansen, E.T., Johannesen, N., 2020. Social distancing laws cause only small losses of economic activity during the COVID-19 pandemic in Scandinavia. Proc. Natl. Acad. Sci. 117 (34), 20468–20473.
Wang, L., Wu, J., Cao, Y., Hong, Y., 2022. Forecasting renewable energy stock volatility using short and long-term Markov switching GARCH-MIDAS models: Either, neither or both? Energy Econ. 111, 106056.