EUROPEAN CENTRAL BANK
EUROSYSTEM

# Working Paper Series

Agostino Consolo, Claudia Foroni,
Claudio Lissona, Christofer Schroeder

## Forecasting the euro area job vacancy rate with earnings calls data

No 3269

**Abstract**

We analyse whether textual information extracted from firms' earnings calls can improve forecasts of the euro area job vacancy rate. Using transcripts from euro area headquartered firms, we construct a monthly indicator of labour demand based on keywords related to labour market pressures and include it into a mixed frequency Bayesian VAR alongside standard hard and soft indicators. A pseudo–real-time evaluation shows that earnings calls provide timely and valuable signals for tracking vacancy dynamics. Among soft indicators, factors limiting production deliver the largest forecasting gains, while real labour-market indicators such as unemployment add little once qualitative signals are included. Forecast improvements are largely driven by information from the manufacturing sector, whose signals prove substantially more informative than those from services, especially when paired with earnings calls. Taken together, our results highlight the usefulness of high-frequency text-based information for improving short-term labour-demand forecasts in the euro area.

**Keywords:** Earnings calls, job vacancy rate, nowcasting, mixed-frequency, sectoral heterogeneity

**JEL classification:** C53, E24, E27

## Non-technical summary

The job vacancy rate is a key indicator of labour demand, providing useful signals on hiring conditions and potential inflationary pressures. However, its usefulness for policy analysis in real time is limited by its short history, frequent revisions, and relatively long publication lag. These constraints make it challenging for policymakers to form timely and reliable assessments of current labour market conditions.

This paper studies whether textual information extracted from firms' earnings calls can improve the monitoring and short-term forecasting of the euro area job vacancy rate. Earnings calls are transcripts from meetings in which firms discuss their recent performance, outlook, and operational challenges, including labour shortages and hiring plans. By extracting and aggregating keywords related to labour market pressures from the transcripts of firms headquartered in the euro area, we construct a monthly indicator of labour demand that is timely, unrevised and available twice per month, well ahead of both official statistics and alternative soft indicators.

We combine this text-based indicator with both hard and soft indicators in a mixed-frequency Bayesian Vector Autoregression (BVAR), which allows us to jointly incorporate quarterly and monthly information within the model. Our analysis focuses on short-term forecasts, which are particularly relevant for policy, and evaluates performance in a pseudo real-time setting that mimics the information available to policymakers at each point in time.

The results show that timely, qualitative information is important for tracking job vacancy dynamics. Among the soft labour market indicators we consider, survey-based measures of the factors limiting firms' production provide the largest forecasting gains, particularly in the post-pandemic period. Our indicator of labour demand derived from earnings calls improves short-term forecasts, especially at very short horizons, by providing additional timely signals on firms' labour demand. Hard indicators, such as the unemployment rate, do not provide a material improvement in forecasting performance above that obtained from the soft indicators.

Our findings also document meaningful differences across sectors. Soft indicators of labour market conditions from the manufacturing sector contain the most informative signals for forecasting the aggregate job vacancy rate. This result is further corroborated by assessing the forecasting performance of sectoral job

vacancy rates. Soft indicators are shown to improve forecasting performance of the job vacancy rate in the manufacturing sector considerably more than forecasting performance of the job vacancy rate in the services sector.

Overall, the findings suggest that high-frequency text-based information from firms' earnings calls serves as a relevant complement to traditional indicators, enhancing real-time monitoring of labour demand in the euro area. More broadly, the paper contributes to the ongoing efforts of research and policy institutions to incorporate alternative data sources into standard and novel econometric frameworks, enriching the information set available for short-term forecasting in real time.

From a policy perspective, timely signals on labour demand are particularly valuable when labour market conditions evolve rapidly and when labour shortages and hiring constraints may suggest wage and price pressures. This is especially relevant for monetary policy, which relies, among other indicators, on timely assessments of labour market slack. By showing that qualitative and text-based indicators provide additional information beyond standard hard data, the paper underscores the benefits of complementing official labour market statistics with high-frequency, forward-looking measures in the short-term analysis of euro area labour market developments, without substituting for established indicators.

# 1 Introduction

Labour demand conditions are a key element for understanding the strength of the business cycle and the degree of labour market tightness, especially in periods when the unemployment rate – a key labour demand indicator – shows little cyclical variation. In this context, a real-time indicator of labour demand can be very informative for the conduct of monetary policy. It complements other business cycle indicators and provides important information on the degree of labour market tightness and matching efficiency (Blanchard and Diamond, 1989; Pissarides, 2011; Consolo and Da Silva, 2019), and can also signal potential inflationary pressures arising from labour market slack, which proved particularly relevant in the post-pandemic period (Benigno and Eggertsson, 2023).

A standard measure of the euro area labour demand is the quarterly job vacancy rate (JVR). However, its usefulness for real-time policy analysis remains limited. The available time series for the euro area is short, subject to sizable revisions, and released with a non-negligible lag – typically around two months after the reference quarter. These limitations restrict policymakers' ability to form timely assessments of current and future labour market conditions. Moreover, many of the standard alternative indicators used to infer on the state of the labour market are subject to similar limitations. This motivates the need for methodological frameworks capable of exploiting alternative sources of information, particularly those available at higher frequency, to improve the monitoring and forecasting of labour demand conditions. Particularly, survey-based indicators, as well as text-based measures, have been shown to contain valuable information for forecasting real activity (Lahiri and Monokroussos, 2013; de Bondt, 2019).

In this paper, we assess whether textual information extracted from firms' earnings calls can help to address the challenges of assessing labour market conditions. In particular, we study whether data from earnings calls can improve nowcasts and short-term forecasts of the euro area job vacancy rate in comparison to approaches which rely solely on standard soft and hard indicators.

Our analysis reveals three key findings. First, the text-based labour demand indicator provides an improvement in the real time forecasting performance of the JVR. This result emerges from our baseline model, which also includes the unemployment rate and the firm survey indicators on "labour as a factor limiting production" in manufacturing, services and construction (henceforth, "FLPs"). FLPs were selected among other soft indicators such as PMI, retail sales and economic sentiment indicators as they are able to capture more effectively the strong mean-reversion of the JVR observed in recent quarters.

Second, we find that hard labour market indicators, including the unemployment rate, do not improve the forecasting performance achieved from only considering relevant soft indicators in the model. This confirms that timely, qualitative information from soft indicators plays an important role in tracking vacancy dynamics. Nevertheless, we retain unemployment in our baseline specification to preserve the information content from the Beveridge curve.

Third, the bulk of the improvement in forecasting performance can be traced to information originating from the manufacturing sector. A simple model that includes only FLP and earnings-calls-derived information from firms in the manufacturing sector outperforms a specification using only services sector information. The model using only manufacturing sector information also performs at least as well, and in some specifications better than a model embedding the full set of FLP information.

Firms' earnings calls have emerged as a promising high-frequency source of economic information. In an earnings call, firm representatives typically first discuss the most recent financial results and the firm's outlook before answering questions from the audience. Earnings calls can, therefore, provide insights into a firm's view on its own performance, as well as the broader macroeconomic environment, including the labour market. In this paper, we follow Dueholm et al. (2024) and apply textual analysis to a large database of earnings calls to construct a monthly indicator of labour demand in the euro area. Our indicator is timely, unrevised, and highly correlated with the euro area job vacancy rate, making it a promising complement to traditional labour market statistics.

To assess the informational content of our indicator of labour demand, we estimate a Mixed Frequency Bayesian VAR (MF-BVAR), which allows us to jointly incorporate monthly and quarterly indicators within a single framework. We consider multiple model specifications combining our indicator derived from earnings calls with official statistics, such as the unemployment rate or real GDP growth. We also consider different sets of soft indicators, including survey-based measures of economic sentiment, FLPs, purchasing managers' indices, and retail turnover indicators, all commonly used to track near-term economic conditions in the euro area.

First, we evaluate the short-term (up to one year ahead) forecasting performance of models that include our labour demand indicator derived from earnings calls, the unemployment rate, and each of the four sets of soft indicators described above. Using a pseudo-real-time exercise starting in 2016Q1, we assess forecast accuracy over the full sample and over a post-COVID subsample beginning in 2021. Across both samples, and especially post-COVID, we find the model including FLPs to outperform the other models.

To assess the relative importance of official statistics versus our text-based indicator, we re-estimate the model including FLPs while excluding either the unemployment rate or the earnings-calls indicator. Removing unemployment does not worsen forecast accuracy, implying that FLPs already contain most of the relevant information. By contrast, omitting our labour demand indicators leads to a modest deterioration of the forecasting performance of the model in the very short-term, with negligible differences at longer horizons. While these accuracy gains are on average modest, the indicator's main strength is its timeliness. Available at monthly frequency but updated biweekly, and free of revisions and publication lags, it offers an early reading of emerging pressures on labour demand well ahead of the official job vacancy statistics, and more timely than other soft indicators.

Finally, to examine whether the informational content of our indicator is sector-specific, we estimate two sub-models using FLPs from manufacturing or services together with alternative high-frequency predictors. We find that our labour demand indicator is particularly informative when combined with manufacturing FLPs, yielding forecast performance comparable to—or at times better than—the model embedding all sectoral FLPs. In contrast, models relying on services FLPs perform worse at all horizons, regardless of the high-frequency predictor used.

Given these results, we then turn to a study of how the performance of our forecasting approach varies by sector. First, we assess forecasts of sector-specific job vacancy rates using our aggregate indicator of labour demand from earnings calls. We find that both a simple model using only our earnings-calls-derived indicator, and a richer model incorporating sectoral FLP information, substantially improve forecasting performance for the job vacancy rate in manufacturing compared with the aggregate job vacancy rate. In contrast, improvements in forecasting performance in the services sector are limited. This provides further evidence that the information captured by earnings calls is particularly relevant for understanding developments in the manufacturing sector. Next, we construct separate indicators of labour demand from earnings calls for each sector, splitting our sample of earnings calls by the originating sector of the associated firms. Although these sector-specific earnings call indicators improve job vacancy rate forecasts relative to services, they do not provide meaningful improvements beyond those obtained with the aggregate earnings call indicator. This suggests that aggregate earnings calls already contain most of the sector relevant information that can be fruitfully exploited for forecasting purposes.

Our paper contributes to several strands of the literature. First, it speaks to a growing body of work forecasting euro area labour market conditions with high-frequency data and mixed-frequency methods. Gogas et al. (2022) forecast euro area unemployment using machine learning and long historical samples; Consolo et al. (2023) develop a mixed-frequency BVAR to forecast employment growth and study the structural drivers of the euro area labour market; and Bańbura et al. (2023) employ dynamic factor and bridge models to nowcast euro area employment growth from monthly indicators. This literature, however, focuses mostly on employment and unemployment, for which long series and timely monthly indicators exist. The job vacancy rate has received far less attention, despite being a direct gauge of labour demand, because its short history, sizeable revisions and publication lag, together with its weaker comovement with other labour market series, make it difficult to forecast with standard hard indicators alone. In this paper, we address this gap by proposing a dedicated modelling strategy in which soft indicators, and particularly earnings calls, play a central role. Relative to the labour market forecasting literature, we are, to the best of our knowledge, the first to formally target the euro area job vacancy rate incorporating text-based information.

Second, our paper contributes to the literature using text-based indicators for economic forecasting, and in particular to the use of earnings calls as a source of timely economic signals. Text-derived indicators have been shown to improve macroeconomic forecasts: Kalamara et al. (2022) find that newspaper-based indicators improve forecasts of several UK macroeconomic variables, including labour market ones; Aprigliano et al. (2023) document forecasting gains from text data for a range of macro-financial variables; and Barbaglia et al. (2024) use text-based sentiment to enhance euro area GDP forecasts. Earnings calls are an especially attractive source in this respect: they are firm-level, forward-looking, available at high frequency and free of revisions, and have been used to construct reliable measures of firms' exposure and sentiment along several dimensions: Brexit-related uncertainty (Hassan et al., 2024), political risk (Hassan et al., 2019), and the diffusion of new technologies (Kalyani et al., 2025). These features make them a natural candidate for tracking labour demand in real time, ahead of both official statistics and survey-based soft indicators. In this spirit, we introduce a novel, high-frequency indicator of euro area labour demand built from earnings calls and show that earnings calls carry a distinct, timely signal that complements both hard and soft indicators, rather than substituting for them.

The rest of the paper is structured as follows. Section 2 describes the data on earnings calls used in the paper and our approach to using these data to construct an indicator of euro area labour demand. Section 3 includes preliminary results obtained with a univariate model to motivate the usefulness of earnings calls. Section 4 outlines the mixed-frequency BVAR employed in the analysis. Section 5 presents results for the aggregate euro area job vacancy rate, while Section 6 presents results looking at sectoral job vacancy rates as well as sectoral earnings calls. Finally, Section 7 concludes. The paper comes with a Supplementary Appendix, including additional results and robustness checks.

## 2 Constructing an indicator of labour demand from earnings calls

In this paper, we apply textual analysis to a large database of transcripts of corporate earnings calls to construct an indicator of euro area labour demand. An earnings call is a conference call that takes place between the management of a publicly listed company and its investors, financial analysts and the press. Earnings calls are typically held quarterly, coinciding with the release of a firm's financial results for a given quarter. During an earnings call, firm representatives typically first discuss the most recent financial results and the firm's outlook before answering questions from the audience. Earnings calls do not follow a standardized structure, nor do they always convey the same set of information across firms or over time. Nevertheless, the soft information contained in an earnings call typically contains insights into (i.) a company's view on its own recent performance and the factors impacting it, (ii.) a company's outlook on its own performance, and (iii.) a company's outlook on the broader macroeconomic environment. While firms discuss a mix of short-, medium-, and long-term factors, uncertainty surrounding the long-term environment makes it likely that the information conveyed in an earnings call is most pertinent over the short run.

We exploit the database of earnings calls and accompanying analytical platform provided by NL Analytics. As of late 2025, the database contains nearly 450,000 earnings calls from companies around the world dating back to 2002. For each earnings call, the database contains its transcript, a time stamp with the date, hour, and minute each earnings call began, as well as basic characteristics about the associated firm, including country of headquarters and detailed sectoral identifiers. For proprietary reasons, NL Analytics does not provide users with the full transcripts of each earnings call in the database. Rather, users can search through the database for the appearance of pre-specified keywords. For each earnings call, users can observe several metrics, including the number of sentences containing at least one keyword from a pre-specified list. The earnings calls transcripts are processed by NL Analytics to be in a consistent format (i.e., using American English spelling) and the database is updated with new earnings calls every two weeks. While earnings calls often take place around the middle of each quarter, the data contain a reasonable number of observations outside of these periods such that the data can be relied on as a high-frequency source of economic information.

We proceed in using the soft information contained in transcripts of earnings calls to construct an indicator of euro area labour demand as follows. First, we restrict our sample of earnings calls to those by companies which were headquartered in the euro area at the time the earnings call took place. This leaves us with a sample of roughly 600 calls per quarter. Selecting companies based off of the country in which their headquarters are located is motivated by the fact that this information is readily available for nearly every earnings call in the data.[1] This restriction may, however, exclude earnings calls from firms which conduct a relevant share of their business in the euro area, but are headquartered outside of it.

Next, we add up the number of sentences in each earnings call containing at least one word from a list of keywords capturing the pressures firms face from unmet labour demand. The keywords are listed in Table 1 and are borrowed from Dueholm et al. (2024). This sum of sentences that we compute for each earnings call does not take into account any further contextual aspects of the text, including negation or the appearance of temporal qualifiers. As such, our indicator does not distinguish between sentences that contain, for example, the text “labor shortages” and “no labor shortages”, or “current labor shortages” and “future labor shortages”. We then collapse our dataset to the monthly frequency, summing up the number of sentences from each earnings call in each month.

[1] As of 2025, the country of headquarters was missing for only 0.19% of all earnings calls in the data.

Table 1: Keywords used to construct an indicator of labour demand

| Keyword | Example |
|---|---|
| Labor costs | "At the same time, rising **labor costs** and the lack of skilled workers are driving companies to automate." |
| Labor cost | "And can you give us also maybe some examples of initiatives you're putting in place to offset the rising **labor cost**?" |
| Personnel costs | "Another measure is the streamlined program which aims to reduce **personnel costs**." |
| Wage inflation | "The increase in current personal expenses was due to general **wage inflation**." |
| Labor shortages | "In addition, the robotics and automation market is growing rapidly, driven by demographic change, **labor shortages** and cost pressure." |
| Labor inflation | "It's still material inflation, but then there is also for sure energy inflation, and last but not least, **labor inflation**." |
| Labor shortage | "What we face right now is really a **labor shortage**." |
| Labor issues | "Then when it comes to volumes, that's a little bit for us impacted by the **labor issues** that we have mentioned." |
| Wage pressure | "Costs are up 1.7% year-on-year at constant FX due to **wage pressure** and well below inflation." |
| Labor expense | "So we truly are expecting this **labor expense** to accelerate in Q3 and Q4, but that's really how it's gated and how it offsets with the PRF." |
| Labor challenges | "And as I previously explained, we will use additional funds to mitigate the ongoing **labor challenges** through the remainder of the year." |
| Salary inflation | "And in terms of profitability, we've been able to maintain the EBIT margin above 5%, despite the context of strong **salary inflation**, which has been higher than any other year in the recent past." |
| Wage pressures | "We continue to see productivity improvements in our in center hemodialysis product line, mitigating somewhat the **wage pressures** that exist in many of our markets." |
| Labor pressures | "Along with EPO and **labor pressures**, insurance costs continue to be an issue for our company." |
| Labor pressure | "So you could read into that that you're seeing some **labor pressure** in service delivery, especially given your GDN mix has been constant now for six or seven quarters." |
| Labor tightness | "We see **labor tightness** as high as it's ever been since 2008." |

NOTES: This table lists the keywords used to construct our indicator of labour demand from earnings calls. The keywords are borrowed from Dueholm et al. (2024) and are selected to capture the pressures firms face from unmet labour demand. Transcripts of earnings calls contained in the data have been transcribed following American English spelling, i.e. the word "labour" is consistently spelled as such in each transcript, regardless of the country of the associated firm's headquarters.

Table 2: Descriptive statistics

| | Total | Industry | Services | Missing |
|---|---|---|---|---|
| Number of earnings calls | 37,935 | 15,897 | 21,233 | 805 |
| Share of total | 1.00 | 0.42 | 0.56 | 0.02 |
| | | | | |
| Share from country of headquarters | | | | |
| - Germany | 0.22 | 0.27 | 0.18 | 0.07 |
| - France | 0.15 | 0.12 | 0.17 | 0.05 |
| - Italy | 0.10 | 0.10 | 0.09 | 0.08 |
| - Spain | 0.08 | 0.05 | 0.10 | 0.06 |
| - Netherlands | 0.10 | 0.09 | 0.10 | 0.06 |
| - Other | 0.35 | 0.37 | 0.36 | 0.68 |

NOTES: This table presents a number of descriptive statistics of the earnings calls data as of late 2025. The first column contains statistics for the total sample of earnings calls. The second and third columns present statistics for firms in the industrial and services sectors respectively. We assign each firm to either the industrial or services sector based on its 10-digit Thomson Reuters Business Classification (TRBC) identifier. Lacking a publicly available key, we make this allocation on the basis of our own judgment. The final column presents statistics for earnings calls in the sample which are recorded with a missing TRBC identifier.

Table 2 presents a number of descriptive statistics of the earnings calls data as of late 2025. The first column contains statistics for the total sample of earnings calls. Of the roughly 38,000 earnings calls in our sample, about two-thirds were from firms headquartered in Germany, France, Italy, Spain, or the Netherlands — the "big-5" countries. The shares of earnings calls from each of the big 5 countries are broadly in line with each country's share of total euro area GDP and population as reflected by, for instance, the Eurosystem capital key. Compared to this benchmark, Germany, France, Italy, and Spain are all slightly underrepresented, while the Netherlands is slightly overrepresented. The remaining columns present breakdowns by sector. We assign each firm to either the industrial or services sector based on its 10-digit Thomson Reuters Business Classification (TRBC) identifier. 42% of all earnings calls were from industrial sector firms while 56% were from firms in the services sector. Given the fact that services accounts for roughly three-quarters of gross value added and employment in the euro area, our sample of firms is slightly underrepresented in the services sector compared with these benchmarks. Only about 2% of all earnings calls are not allocated to a sector due to missing data. The majority of these earnings calls are from firms headquartered outside of the big 5 countries.

Figure 1 plots our indicator aggregated to the quarterly frequency next to the job vacancy rate in the euro area. Although it is somewhat more volatile, our indicator closely tracks the job vacancy rate, with a correlation coefficient of around 0.9. In particular, it captures the steady increase in the job vacancy rate between 2016 and 2020 and its sharp rise in 2021. Both our indicator of labour demand and the job vacancy rate peaked in the first half of 2022 and have been trending downwards ever since, suggesting that labour demand has been cooling.

Figure 1: Indicator of labour demand from earnings calls and the job vacancy rate in the euro area

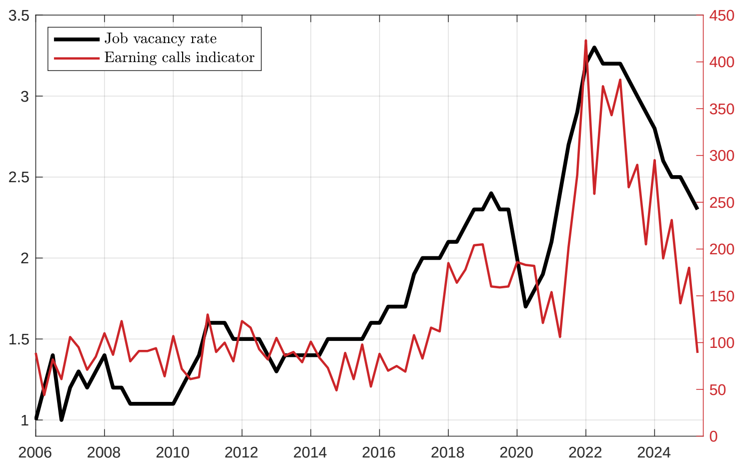


NOTES: This figure plots our monthly indicator of labour demand derived from earnings calls as described in the text (red solid line) aggregated at the quarterly frequency, together with the quarterly job vacancy rate in the euro area (black solid line, y-scale on the left) compiled as by Eurostat.

# 3 Univariate analysis

In this Section, we first estimate simple univariate models to evaluate how useful earnings calls can be for forecasting the euro area job vacancy rate. As a baseline, we consider a Mixed-Data Sampling (MIDAS) model and evaluate its performance in forecasting the euro area job vacancy rate with respect to two workhorse models: an AR(1) model and an ARDL(1,1) model including real year-on-year GDP growth as auxiliary low frequency indicator. MIDAS models provide a natural and flexible framework for exploiting high-frequency information to forecast low-frequency indicators (Ghysels and Marcellino, 2016) and have a well-established track record in macroeconomic forecasting (Clements and Galvão, 2009; Kuzin et al., 2011).

## 3.1 Baseline model

We estimate an AR-MIDAS model of the quarterly euro area job vacancy rate as a function of our monthly earnings-calls-derived indicator of labour demand following Ghysels et al. (2004) and using the notation in

Foroni and Marcellino (2014). The model takes the following form:

$$y_{t_m+h_m} \;=\; \beta_0 + \beta_1 b(L^{\frac{1}{m}};\theta)x^{(m)}_{t_m+w} + \alpha y_{t_m+h_m-3} + \varepsilon_{t_m+h_m} \tag{1}$$

Let $t_q = 1, \ldots, T_q$ denote the time index for the *low-frequency* variables of interest which, in our case, are sampled at the quarterly ($q$) frequency, and $m$ denote the number of times a *high-frequency* variable appears for each $t_q$. Hence, the time-indexing for the high-frequency variable is $t_m = 1, \ldots, T_m$. In our case, $m = 3$ since we focus on monthly data. The low-frequency variable of interest, $y_{t_q}$, can be expressed at the higher frequency by setting $y_{t_m} = y_{t_q}$ $\forall t_m = mt_q$.

$h_m = mh_q$ is the $h_q-$step ahead forecast horizon expressed in high-frequency units and $w$ is the number of monthly observations of $x_{t_m}$ that are available earlier than the lower-frequency variable $y_{t_m}$. $b(L^{\frac{1}{m}};\theta)$ is a real polynomial in the lag operator $L^{\frac{1}{m}}$, such that $b(L^{\frac{1}{m}};\theta) = \sum_{k=1}^{K} c(k,\theta)L^k_m$, where $L^j_m x^{(m)}_{t_m} = x^{(m)}_{t_m-j}$.

The coefficients of $c(k;\theta)$ are parametrized using the *Exponential-Almon* lag polynomial, taking the following form:

$$c(k;\theta) = \frac{\exp\{\theta_1 k + \ldots + \theta_Q k^Q\}}{\sum_{k=1}^{K} \exp\{\theta_1 k + \ldots + \theta_Q k^Q\}} \tag{2}$$

In our application, we follow Ghysels et al. (2005) and adopt a functional form of (2) with two parameters and $K = 6$ lags (i.e., monthly information for two quarters). In practice, this amounts to weighting monthly observations with a flexible *smooth* polynomial while reducing over-parametrization.

### 3.2 Alternative models

To assess the performance of our baseline specification, we also consider two alternative univariate models which are standard benchmarks for forecasting macroeconomic variables (see, e.g., Clements and Galvão, 2009). First, we consider a standard AR(1) process with drift for the job vacancy rate, which takes the following form:

$$y_{t+h} \;=\; \alpha_0 + \alpha_1 y_{t+h-1} + \varepsilon_{t+h} \tag{3}$$

where $t = 1, \ldots, T$ denotes time periods at the quarterly frequency. In this framework, the job vacancy rate only depends on its past and not on any other additional variables.

Second, we also consider a richer workhorse model in the forecasting literature — an ARDL(1,1) model with drift. This model augments the standard AR model by adding lagged variables of auxiliary regressors. Specifically, we consider as an auxiliary, low-frequency variable the year-on-year growth rate of real GDP, denoted as $z_t$. The model takes the following form:

$$y_{t+h} = \alpha_0 + \alpha_1 y_{t+h-1} + \beta_0 z_{t+h-1} + \beta_1 z_{t+h-2} + \varepsilon_{t+h} \quad (4)$$

The choice to include real GDP growth as an auxiliary indicator is motivated by two main reasons. First, real GDP and the job vacancy rate exhibit strong comovement over the business cycle. Second, developments in real GDP are more generally informative of broad labour market conditions. Strong GDP growth is often associated with tight labour market conditions, and vice versa.

Model 4 includes lagged real GDP growth information, accounting for the true data flow of these variables. Indeed, in forecasting the job vacancy rate at time $t + h$ the information available for other real variables is available at most up to time $t + h - 1$. While the forecast can be conditioned on institutional projections of contemporaneous GDP, this becomes increasingly difficult in quasi-real time, where all vintages of such projections are needed.

### 3.3 Results

Here we report the performance of our baseline MIDAS model including the indicator of labour demand derived from earnings calls compared with the two alternative specifications outlined above. When evaluating forecasting performance, we consider two different samples: a *full sample*, where the forecast evaluation starts from 2016:Q1 and a *post-COVID sample* where the forecast evaluation starts in 2021:Q1. In both cases, we carry out a pseudo-real time exercise, where the forecasting performance of the model at time $t+h$ is evaluated with the model estimated up to time $t$.

Table 3 reports the results of this exercise. Our baseline model including information derived from earnings

Table 3: RMSFE: Exponential Almon MIDAS vs AR(1) and ARDL(1,1)

**Full sample**

| | **H=1** | | | **H=2** | | | **H=3** | | | **H=4** | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | M1 | M2 | M3 | M1 | M2 | M3 | M1 | M2 | M3 | M1 | M2 | M3 |
| **AR(1)** | 1.02 | 0.99 | 0.98 | 1.05 | 0.97 | 0.98 | 1.12 | 1.03 | 0.99 | 1.10 | 1.04 | 1.00 |
| **ARDL(1,1)** | 0.77 | 0.74 | 0.74 | 1.04 | 0.96 | 0.98 | 1.16 | 1.08 | 1.02 | 1.17 | 1.10 | 1.06 |

**Post-COVID**

| | **H=1** | | | **H=2** | | | **H=3** | | | **H=4** | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | M1 | M2 | M3 | M1 | M2 | M3 | M1 | M2 | M3 | M1 | M2 | M3 |
| **AR(1)** | 0.88* | 0.79* | 0.79* | 0.80 | 0.73* | 0.73* | 0.73* | 0.70 | 0.70 | 0.69 | 0.67 | 0.67 |
| **ARDL(1,1)** | 0.90 | 0.81 | 0.81 | 0.83 | 0.76 | 0.76 | 0.76 | 0.72 | 0.72 | 0.74 | 0.72 | 0.73 |

NOTES: The table reports the relative mean-squared forecast error (RMSFE) for the MIDAS, compared with the AR(1) model (first row) and the ARDL(1,1) model (second row) with lagged GDP as an auxiliary indicator. On the columns, the different forecast horizons ($H$) and the different ending months for the high-frequency indicators within quarter $T+1$ ($M$). Numbers lower than 1 indicate a better performance of the MIDAS model. Asterisks denote whether the different performance of the two models is statistically significant according to the Diebold and Mariano (2002) test. The upper part of the table shows results for the *full sample*, with the forecast evaluation starting in 2016:Q1, while the bottom part of the table shows results for the *post-COVID*, with the forecast evaluation starting in 2021:Q1.

calls performs at least as well as, and over some horizons better than, the alternative models[2]. Particularly, our baseline model seems to perform better in the post-COVID period, where we observe the strongest improvements in terms of forecasting performance with respect to both benchmark models. Furthermore, the information contained in the last month of each quarter conveys the strongest signal for the job vacancy rate, which may reflect the fact that firms whose earnings calls fall at the end of the quarter have a broader information set compared with firms who hold their earnings calls earlier in the quarter.

Figure 2 compares the forecast profile of the AR(1) and ARDL(1,1) models (upper panels) with that of the MIDAS model including our indicator of labour demand from earnings calls as a high-frequency, auxiliary indicator (bottom panel). On average, information from earnings calls seems to be informative about the *direction* of the job vacancy rate, displaying a forecast profile which is more flexible in capturing up- and downturns with respect to alternative models. This is particularly evident in the post-COVID period, where the forecast obtained with the MIDAS model more closely tracks the actual dynamics of the job vacancy rate at all horizons.

The preliminary evidence offered in this section suggests that the information contained in earnings calls can be timely and informative for predicting the job vacancy rate in the euro area. Specifically, we find that earnings calls data (i) are particularly valuable for nowcasting, as they contain information not captured by hard indicators; (ii) perform well during periods when hard indicators, such as GDP, are highly volatile (e.g.,

[2]These results are obtained using a 9-month moving average of our labour demand indicator of earnings calls as auxiliary monthly indicator. Results using the raw indicator and alternative transformations are available upon request

Figure 2: Pseudo-real-time forecast: MIDAS vs alternative models

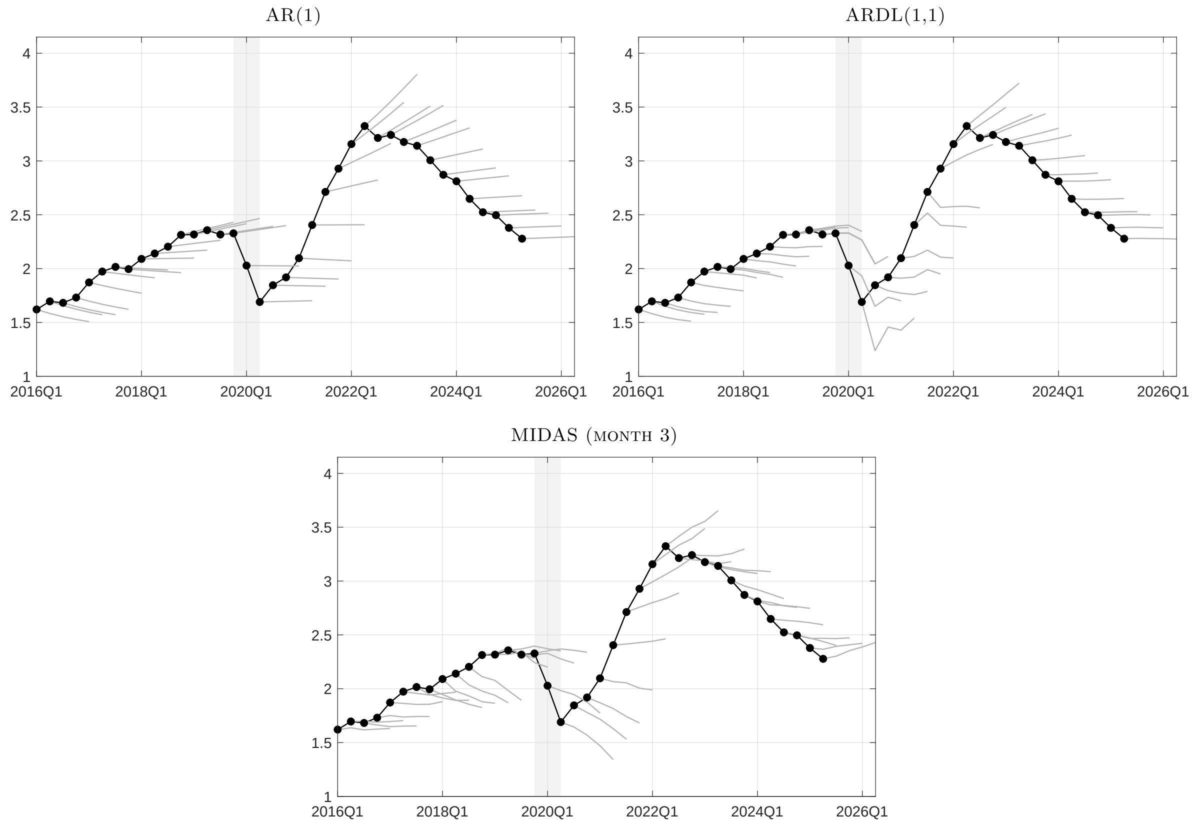


Notes: The figure plot the job vacancy rate from 2016:Q1 to 2025:Q2 (black solid line) along with the recursive forecasts up to 4-quarters ahead. Each quarter is denoted with a black dot, and each thin gray line stemming from each quarter is the forecast with parameters estimated up to that quarter. Estimates for the MIDAS are obtained with the Exponential MIDAS model using information for the last month in each quarter for a 9-month moving average of earnings calls data.

the COVID period); and (iii) serve as effective substitutes for alternative high-frequency hard indicators, which are often released with a lag. These findings motivate a deeper analysis considering more complex forecasting approaches which we turn to in the next section.

# 4 Multivariate Analysis

We now turn to a multivariate framework to examine the role of our indicator of labour demand derived from earnings calls in conjunction with multiple variables, while explicitly modelling their dynamic interdependence.

### 4.1 Mixed-frequency Bayesian VAR

To account for the mixed-frequency nature of our data, with high-frequency indicators recorded at the monthly level and the JVR, among others, recorded at the quarterly level, we estimate a Mixed-Frequency Bayesian VAR (MF-BVAR). While multivariate extensions of the MIDAS framework offer an alternative to handle mixed-frequency data in a multivariate setting (Ghysels, 2016), the MF-BVAR provides a parsimonious representation that is well suited to our short sample. Building on the notation of Schorfheide and Song (2015), the dynamics at the monthly frequency are assumed to follow a VAR($p$) model of the form:

$$\mathbf{x}_t \;=\; \boldsymbol{\phi}_c + \boldsymbol{\Phi}(L)\mathbf{x}_{t-1} + \mathbf{u}_t; \qquad \mathbf{u}_t \sim \text{i.i.d.} N(0, \boldsymbol{\Sigma}) \tag{5}$$

where the $n-$dimensional vector $\mathbf{x}_t$ can be decomposed in $\mathbf{x}_{m,t}$, being the set of *observed* monthly variables, and $\mathbf{x}_{q,t}$ being the set of unobserved monthly realizations of quarterly variables.

Model (5) can be cast in state-space form by defining the state vector $\mathbf{z}_t = (\mathbf{x}'_{t-1}, \ldots, \mathbf{x}'_{t-p})'$, and defining $\boldsymbol{\Phi} = (\boldsymbol{\Phi}_1, \ldots, \boldsymbol{\Phi}_p, \boldsymbol{\phi}_c)'$. The state-space form of the model writes as:

$$\mathbf{y}_t \;=\; \mathbf{M}_t \boldsymbol{\Lambda}_z \mathbf{z}_t \tag{6a}$$

$$\mathbf{z}_t \;=\; \mathbf{F}_1(\boldsymbol{\Phi}) + \mathbf{F}_c(\boldsymbol{\Phi}) + \boldsymbol{\nu}_t \tag{6b}$$

where $\mathbf{M}_t$ is a weighting matrix accounting for the presence of missing values in the data. Specifically, for the quarterly variables in $\mathbf{x}_{q,t}$, $\mathbf{M}_t = \mathbf{M}_{q,t}$ is an identity matrix for each final month in the quarter, and empty otherwise.[3] In contrast, for the monthly variables, $M_{m,t}$ is always an identity matrix, except for periods with missing observations (due to, for instance, data releases). $\boldsymbol{\Lambda}_z = (\boldsymbol{\Lambda}_{mz}, \boldsymbol{\Lambda}_{qz})'$ is a weighting matrix. For quarterly variables, this matrix weights the monthly observations to produce the corresponding quarterly dynamics. Specifically, quarterly flows are defined as the simple average of the three underlying monthly values, that is:

$$\boldsymbol{\Lambda}_{qz}\mathbf{z}_t \;=\; \frac{1}{3}(\mathbf{x}_{q,t} + \mathbf{x}_{q,t-1} + \mathbf{x}_{q,t-2}) \tag{7}$$

[3] Notice that, in this framework, the dimension of the state-space changes according to the reference month within the quarter.

while no weighting is applied to monthly variables, so that $\mathbf{\Lambda}_{mz}$ is the identity matrix. Finally, the state equation (6b) is the companion form of (5).

## 4.2 Model specification and estimation

In this exercise, we consider alternative model specifications to evaluate the informational content of our indicator of labour demand derived from earnings calls in forecasting the euro area job vacancy rate.

To derive a baseline specification, we include in the MF-BVAR model our novel earnings calls indicator, a hard indicator for the labour market – the unemployment rate – and a set of alternative soft indicators. This baseline specification naturally extends the univariate analysis presented in Section 3, and allows us to assess whether information extracted from earnings calls can provide a reliable signal for nowcasting the JVR while adding relevant information not captured by traditional variables. We focus on timely soft indicators suitable for high-frequency, real-time analysis and therefore particularly useful for monetary policy. To keep the model tractable and flexible, we embed these indicators in groups and check their forecasting performance. Specifically, we consider four models. All of them include the earnings calls indicator and the euro area unemployment rate and:

1. *Model 1* includes survey-based Economic Sentiment indicators (ESIs) from the business surveys run by the European Commission. In particular, we include the indicators on general confidence, industry and construction confidence, and industry production expectations over the next three months.

2. *Model 2* includes survey-based indicators of "labour as a factor limiting production" (FLPs) also from the business surveys run by the European Commission. Specifically, indicators of labour shortages as a FLP in manufacturing, services (both available at the quarterly frequency), and in construction (available monthly). These indicators are strongly correlated with the JVR and are commonly used as proxies for it. Furthermore, these indicators are also relatively timely, as the information for the current quarter are released with the monthly survey in the first month of the quarter.

3. *Model 3* includes the composite employment PMI index, as well as the sector-specific ones for employment in manufacturing, services, construction, and the PMI indicator of total output.

4. *Model 4* includes indicators of retail turnover for food items, non-food items, fuel, and total retail.

Soft indicators, particularly ESIs and PMIs, have proven to be a reliable input in forecasting real activity (Giannone et al., 2009; Chudik et al., 2016; de Bondt and Saiz, 2024). Furthermore, other indicators, such as those capturing FLPs, are direct proxies of labour shortages, embedding firms' expectations on labour market outcomes (see, e.g., Arce et al., 2023). This begs the question of whether they convey meaningful signals on labour market conditions, and how useful they are for forecasting the job vacancy rate.

To assess the performance of these models, we conduct a pseudo–real time forecasting exercise. Starting from 2016:Q1, we estimate each model and generate forecasts of the job vacancy rate up to four quarters (twelve months) ahead. We then expand the estimation sample by one month, re-estimate the model with the newly available data, and produce another set of forecasts up to four quarters ahead. We repeat this procedure recursively until the end of the sample, thereby obtaining monthly forecast vintages (three per quarter), which we use to evaluate the relative performance of the models described above.

We consider two evaluation samples: a *full sample*, covering the period 2016:Q1-2025:Q2[4] and a *post-COVID* sample, covering the period 2021:Q1-2025:Q2.[5] When using the full sample, we exclude forecasts corresponding to 2020:Q2 to avoid over-representing the sharp and atypical drop observed in many indicators during that period. All models are estimated with $p = 4$ lags of the indicators, corroborated by standard information criteria and by maximizing the marginal data density (Sims and Zha, 1998), and for each forecast vintage we retain 2,000 draws from the Gibbs sampler. For estimation, we exploit a standard Minnesota prior with standard hyperparameter values commonly used in the literature (see, e.g., Canova and Ferroni, 2021).

# 5 Forecasting the euro area job vacancy rate

In this section, we present the results obtained in forecasting the euro area job vacancy rate using the alternative suites of models introduced in Section 4.2.

[4]For the *full sample*, the mean squared forecast error (MSFE) for each model is computed excluding the months in 2020:Q2 to prevent extreme COVID outliers from distorting the evaluation

[5]Results are virtually identical when starting in 2021:Q3, i.e., after excluding the period affected by the second COVID wave.

## 5.1 Model comparison

We consider four different sets of soft indicators to inform our MF-BVAR, with the aim of assessing which, if any, provide the most informative signals for forecasting the euro area job vacancy rate. To compare the relative performance of the models, we estimate each specification in quasi-real time with an expanding window starting from 2016:Q1 ($T_0$) and ending in 2025:Q2 ($T$), using the latest vintage and increasing the estimation window by one month. This yields a total of $V = 114$ vintages from which we compute forecast errors and the MSFE for each model. Specifically, let $x^{\mathrm{JVR}}_{t+h}$ denote the quarterly $h$-step-ahead job vacancy rate, for $h = 1, \ldots, H$, and let $\widehat{x}^{\mathrm{JVR},m}_{t+h}$ be the corresponding forecast obtained with the MF-BVAR estimated using data from sub-model $m = 1, \ldots, M$. Then, for each period $t = T_0, \ldots, T-H$, we define the $h$-step-ahead forecast error for model $m$ as:

$$\widehat{\nu}^{m}_{t+h} = x^{\mathrm{JVR}}_{t+h} - \widehat{x}^{\mathrm{JVR}}_{t+h} \tag{8}$$

and the corresponding MSFE as:

$$\mathrm{MSFE}^{m}(h) = \frac{1}{T_q - T_0 - h} \sum_{t=T_0}^{T-h} (\widehat{\nu}^{m}_{t+h})^2 \tag{9}$$

Figure 3 compares the MSFEs across the four sub-models used in our analysis with the full sample. The horizons considered stretch from the nowcast ($H = 1$) to a one-year-ahead forecast ($H = 4$). The results indicate that no single model dominates across all horizons. While ESI and PMI indicators perform well in the nowcast, their accuracy deteriorates as the forecast horizon lengthens, with models including FLP and retail turnover indicators yielding more reliable forecasts at medium-term horizons.

This pattern changes markedly in the post-COVID period. As shown in Figure 4, in the post-COVID sample the model augmented with FLP indicators outperforms all alternatives at every horizon.[6] Detailed results for this exercise are included in Appendix B.1.

These results naturally raise the question of why the model augmented with FLP indicators outperforms

[6]These results hold also with alternative measures of real activity, e.g. real GDP. There are two reasons for which we include the unemployment rate rather than GDP as a measure of real activity in the model. First, the unemployment rate is observed at monthly frequencies, hence more timely. Second, as shown in Appendix B.3, although models with real GDP growth seem to be particularly useful in the post-COVID period, their performance considerably worsens in the pre-COVID period and in an extended forecast sample starting from 2016:M1, which makes the unemployment rate a more reliable hard indicator for our analysis.

Figure 3: MSFE across models, full sample

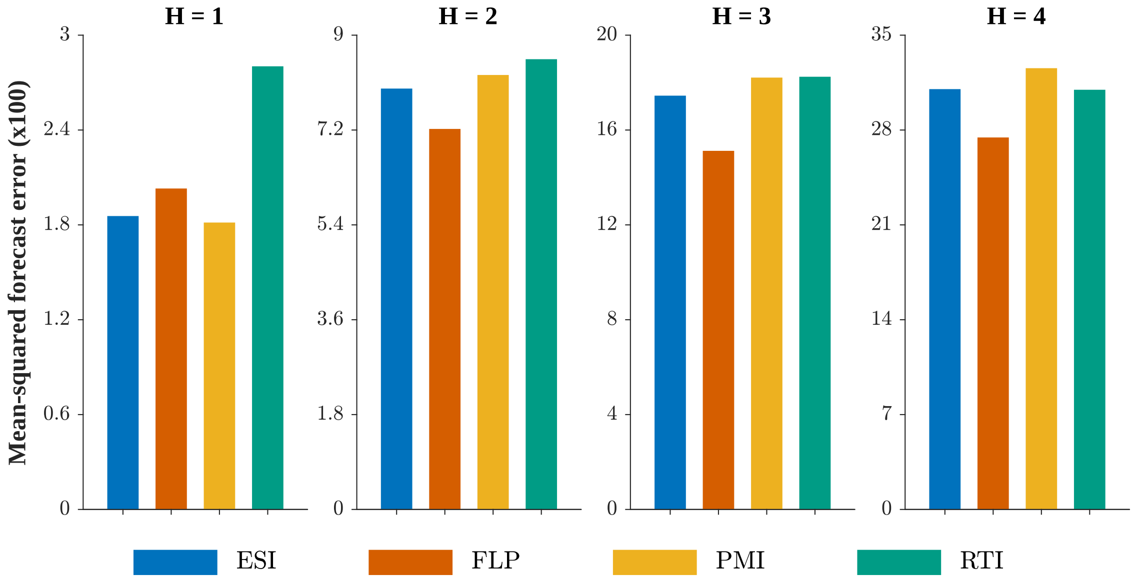


NOTES: The figure shows the MSFE (y-axis) of alternative models (bars) for the *full sample*, scaled by 100 for better interpretability. Each coloured bar corresponds to one of the models outlined above. Each subplot reports results for different forecast horizons, from 1 quarter ($H = 1$) to 1 year ($H = 4$).

Figure 4: MSFE across models, post-COVID period

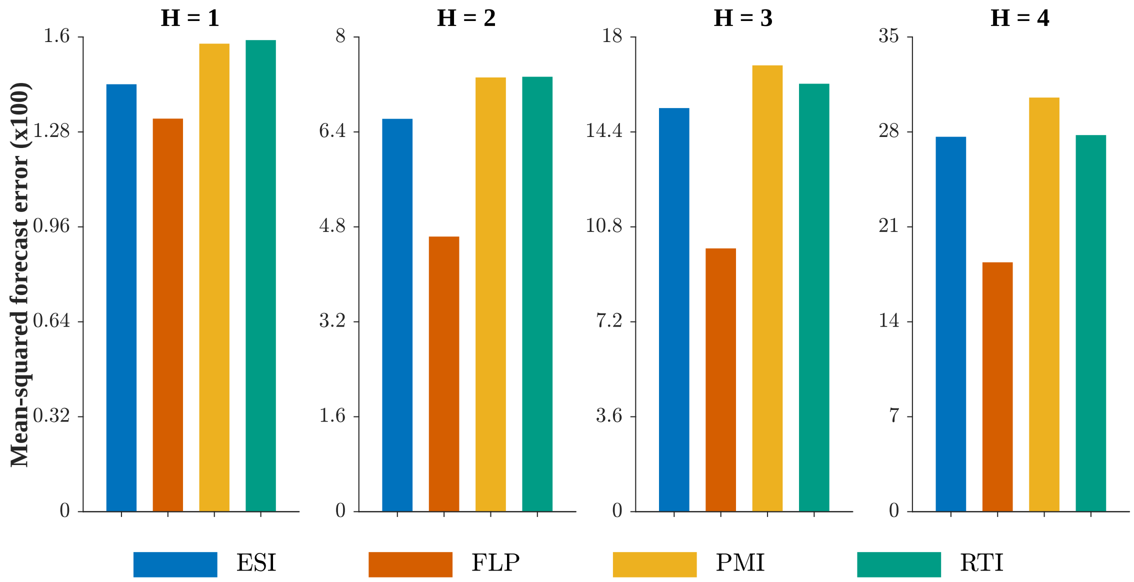


NOTES: The figure shows the MSFE (y-axis) of alternative models (bars) for the *post-COVID sample*, scaled by 100. Each coloured bar corresponds to one of the models outlined above, while the gray bar is a benchmark model only including earnings calls as an auxiliary indicator. Each subplot reports results for different forecast horizons, from 1 quarter ($H = 1$) to 1 year ($H = 4$).

models based on alternative soft indicators. To investigate this, Figure 5 shows the forecast profiles of the first- and second-best performing models over the full sample. Each gray line corresponds to the sequence of forecasts, up to one year ahead, generated using the most recent quarter available in each real-time vintage.

Figure 5: Pseudo-real-time forecast from alternative models

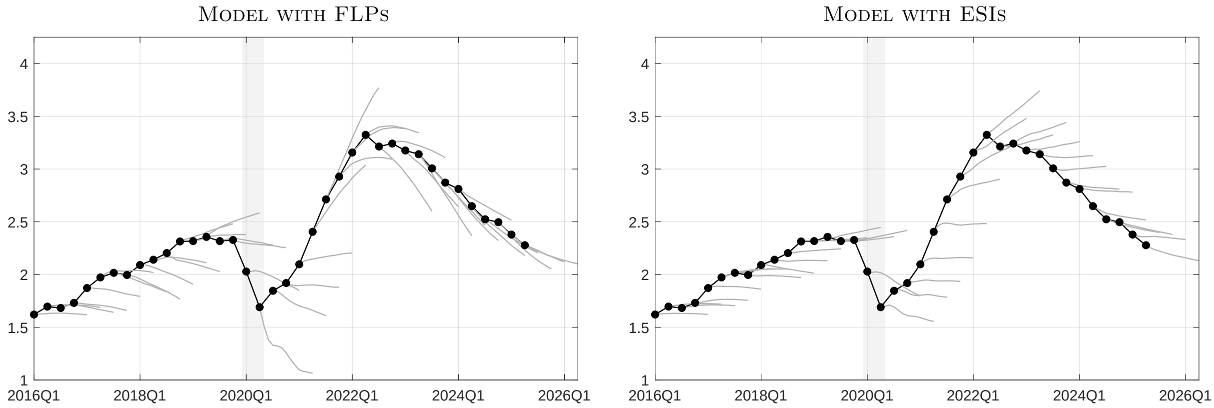


Notes: The figure plots the job vacancy rate from 2016:Q1 to 2025:Q2 (black solid line) along with the forecasts, up to 12-months ahead, obtained with different submodels, starting 2016:M1. Each quarter is denoted with a black dot, and each thin gray line stemming from each quarter is the forecast with parameters estimated up to that quarter.

The figure shows that the model including FLP indicators tracks job vacancy rate dynamics much more closely, especially in the post-COVID period. By contrast, the model with ESIs produces an almost flat forecast profile throughout the sample. This behaviour leads to a reasonably good fit in the pre-COVID period, when job vacancy dynamics were relatively muted, but results in systematic underperformance thereafter, particularly in the most recent years, where the model tends to overshoot at all horizons.

Given these results, we take the model with FLP indicators as the auxiliary indicators and use it to develop some testable implications. Specifically, two questions are of interest:

1. Do real indicators and earnings calls contain meaningful information for forecasting the euro area job vacancy rate *on top of* what is already embedded in soft indicators?

2. Is there any meaningful *sectoral heterogeneity* in the information captured by FLP indicators, and how does this relate to our indicator of labour demand derived from earnings-calls?

To address these questions, we unpack the model augmented with FLP indicators and isolate the drivers of its forecasting performance. In light of the superior performance observed in the *post-COVID sample*, we restrict our forecast evaluation to this period.

Figure 6: MSFE across models: the role of real variables and earnings calls

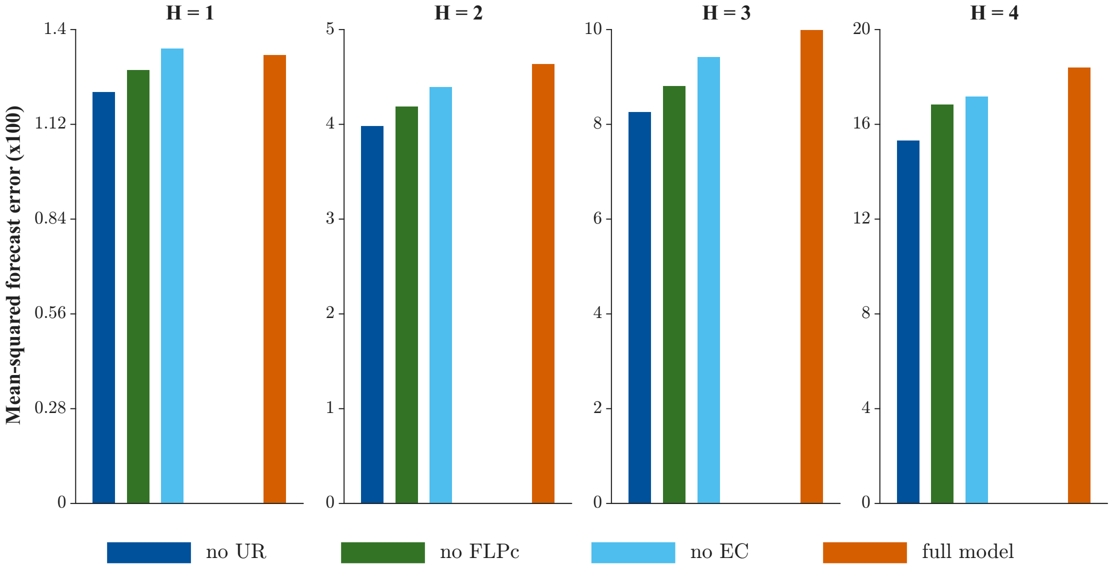


NOTES: The figure shows the MSFE (y-axis) of alternative models (bars) for the *post-COVID sample*, scaled by 100. The red bar corresponds to the model including FLPs, unemployment and earnings calls, the blue bar to the model excluding unemployment, the dark green bar to the model excluding the construction FLP and the turquoise bar to the model excluding earnings calls. Each subplot reports results for different forecast horizons, from 1 quarter ($H = 1$) to 1 year ($H = 4$).

## 5.2 The informational content of real activity and earnings calls

To assess the role of real indicators and earnings calls within our framework, we begin with what we label the *full model*, which includes FLP indicators, the unemployment rate (UR), and our earnings-calls-derived indicator (EC). We then consider two alternative sub-models, obtained by excluding either real variables or the earnings-calls-derived indicator. Comparing the forecasting performance of these restricted specifications with that of the full model allows us to gauge the contribution of each variable to the final forecast.

Figure 6 reports the results of this exercise. The blue bar shows the performance of a model that excludes real activity and therefore relies solely on soft indicators. Relative to the full model (red bar), it is clear that the unemployment rate does not add additional predictive content beyond what is already embedded in the soft indicators. Indeed, the model without real variables slightly outperforms the full model.

Similarly, the turquoise bar in Figure 6 shows the forecasting performance of a model that excludes the earnings-calls-derived indicator. Compared with the full model, earnings calls appear to provide a relevant, though modest signal in the nowcast ($H = 1$), with their contribution fading over longer horizons. This result is intuitive: earnings calls are expected to be informative primarily in the very short run, consistent with the fact that the indicator captures firms' *current* assessment of labour-demand conditions. Moreover, this result is not matched when omitting the other soft monthly indicator included in the model, namely construction FLP (dark green bars). Although excluding this indicator still yields a better performance than a specification without unemployment, it does not match the results observed when removing earnings calls. This suggests that earnings calls convey distinct and complementary information within the model, with respect to other high-frequency indicators.

From this analysis, it is evident that real indicators do not enhance the model's performance and may even deteriorate it. By contrast, our indicator of labour demand from earnings calls yields a modest improvement of around 3% in the very short-run, with no material contribution at longer horizons. We stress, however, that the practical value of the indicator should not be judged solely by the size of this gain. Being timely and revision-free, earnings calls deliver this signal precisely in the very short run, when neither the official job vacancy rate nor many other real and soft indicators are yet available. To more precisely evaluate the role of our earnings-calls-derived indicator, we further investigate whether these indicators provide meaningful information once we disentangle the contribution from sector-specific FLP indicators. Since sectors differ in both dynamics and informational content, looking at this heterogeneity provides us with a preliminary assessment of whether earnings calls can help to forecast the *aggregate* euro area job vacancy rate through specific sectoral channels.

### 5.3 Comparison with alternative labour demand indicators

To further assess the robustness of our results and compare the informational content of our labour market indicator based on earnings calls against alternative high-frequency indicators, we re-estimate our baseline MF-BVAR, including FLPs and unemployment, replacing our earnings-calls-based indicator with a labour-demand measure extracted from Indeed job postings.[7] Since Indeed data are available only from 2018:M1 onward, we also re-estimate our baseline model on the same shorter sample to ensure that differences in performance are not simply driven by the length of the sample.

Figure 7 reports the results of this comparison by presenting out-of-sample forecasts across different phases of the business cycle. The evidence shows that our baseline specification, which includes information from earnings calls, consistently outperforms the alternative model based on Indeed data. In particular, our model

[7] (see, e.g., Adrjan and Lydon, 2020, for an application of Indeed data on the impact of COVID on job postings)

Figure 7: Model performance over time: earnings calls vs. Indeed

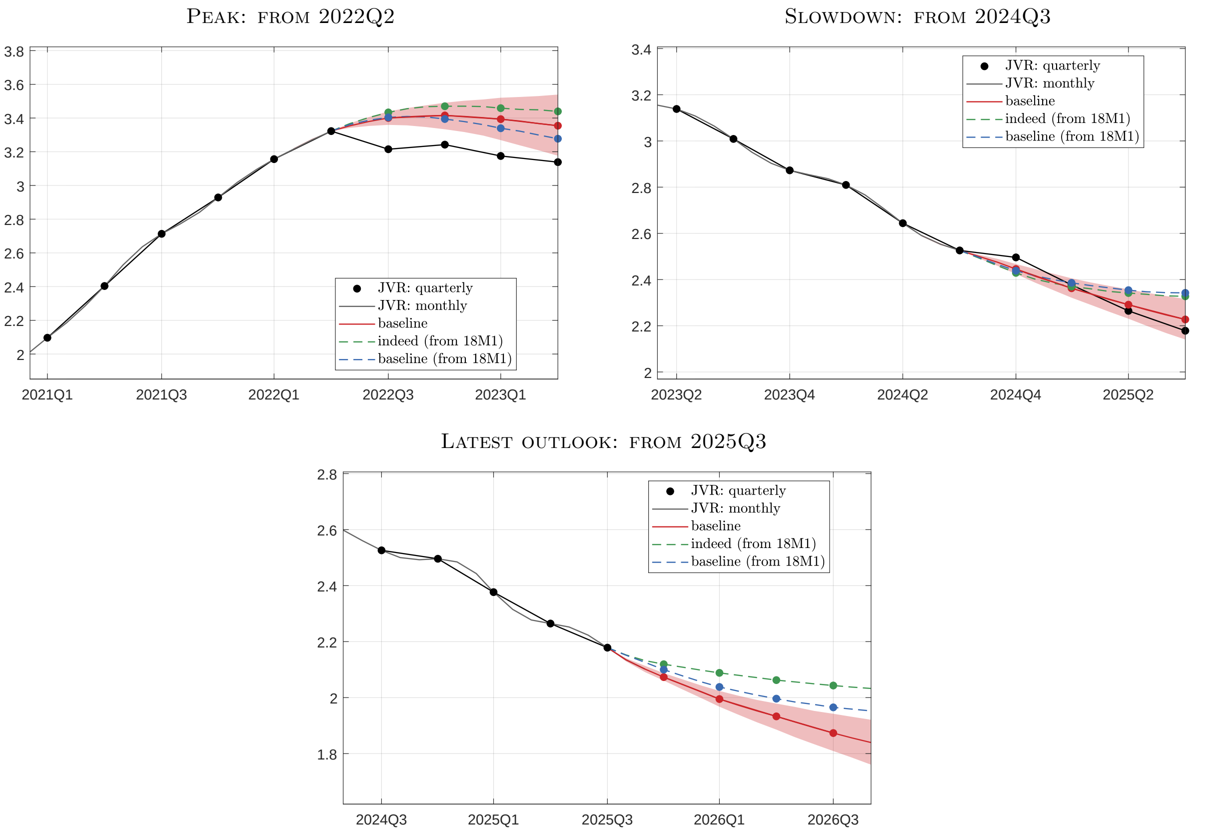


Notes: Each subfigure plots, over different periods of time and business cycle phases, the forecasting performance of different models. Dots represent quarterly values, while lines represent monthly values. Black dots correspond to the observed quarterly JVR while gray lines are the unobserved monthly JVR estimated within the MF-BVAR. Our baseline model (red solid line/dots) is compared with an alternative model using a labour demand indicator extracted from Indeed data, available from 2018:M1 (green dashed lines/dots), together with the same model as the baseline but estimated only from 2018:M1 onwards, to match the Indeed sample (blue dashed lines/dots).

tracks more closely both the turning point after the 2022Q2 peak in the JVR (upper-left subplot) and the pace of the decline thereafter (upper-right subplot). Importantly, this conclusion remains unchanged when restricting estimation to the shorter sample starting in 2018:M1. Taken together, these results indicate that earnings calls embed a reliable and timely signal of labour demand conditions, providing a useful complement to alternative measures while also benefiting from a longer available sample. We nonetheless stress that, given the short sample over which Indeed data are available, our attempt to disentangle the role of the indicator from that of the sample remains tentative.

Finally, based on the latest outlook (bottom panel), forecasts point to a continued cooling of the EA labour market, with the job vacancy rate expected to decline further over the next quarters.

# 6 Sectoral heterogeneity

## 6.1 The role of manufacturing and services within FLPs

The labour market often exhibits distinctly different dynamics between manufacturing and services, reflecting sector-specific productivity trends, wage structures, and cyclical sensitivities (Ngai and Pissarides, 2007; Autor and Dorn, 2013). We next ask whether exploiting this heterogeneity provides additional information for forecasting the euro area job vacancy rate. In particular, we aim to identify whether one sector contributes more strongly to the model's forecasting performance and whether earnings-calls data are more informative at the sectoral rather than aggregate level.

Figure 8: MSFE across models: sectoral decomposition

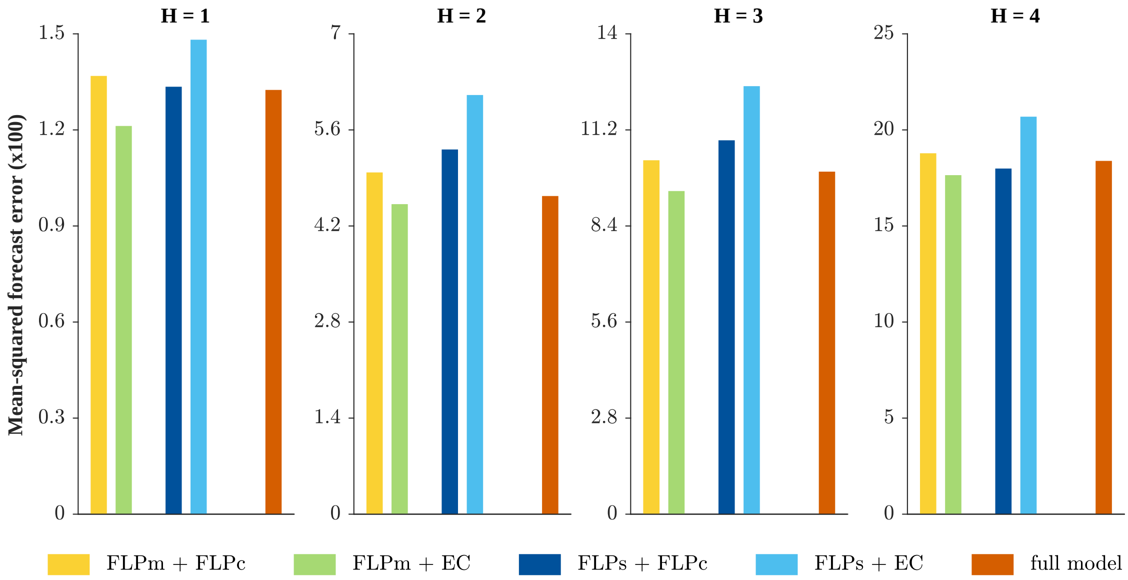


NOTES: The figure shows the MSFE (y-axis) of alternative models (bars) for the *post-COVID sample*, scaled by 100. The red bar corresponds to the model including FLPs, unemployment and earnings calls. The yellow and green bars correspond to a model including FLPs in manufacturing, together with either FLPs in construction (yellow bars) or earnings calls (green bars). Similarly, the blue and turquoise bars correspond to a model including FLPs in services, together with either FLPs in construction (blue bars) or earnings calls (turquoise bars). Each subplot reports results for different forecast horizons, from 1 quarter ($H = 1$) to 1 year ($H = 4$).

Figure 8 compares the full model with two simplified specifications that include either the manufacturing FLP indicator ($m$) or the services FLP indicator ($s$). To evaluate the contribution of earnings calls to forecasting performance, we consider, for each sector, two versions of the model: (i) one combining quarterly FLP indicators for manufacturing or services with the construction FLP indicator (the only FLP indicator available at the monthly frequency) and (ii) the same specification but replacing the construction FLP indicator with our monthly earnings-calls-derived indicator.

We obtain two key results. First, irrespective of the monthly indicator used, models embedding information from manufacturing systematically outperform those based on information from the services sector. This indicates that current developments in the manufacturing sector are more informative for forecasting the euro area job vacancy rate than those in services. Second, earnings calls provide a clear forecasting gain when paired with manufacturing (green bars), with the model including manufacturing information and our earnings-calls-derived indicator even outperforming the full specification. By contrast, information from earnings calls offers little added information when paired with indicators from the services sector (turquoise bars), where the forecasting performance deteriorates. As shown in Appendix C.1, these results are robust when we include the pre-COVID period within our forecast evaluation sample, although the forecast gains from earnings calls are smaller.

These findings motivate a deeper investigation of the sectoral decomposition of the job vacancy rate, both to assess whether information from earnings calls can enhance sector-specific forecasts and to determine whether they contain intrinsically sectoral information that can be meaningfully exploited.

## 6.2 Forecasting sectoral job vacancy rates

Given the evidence of sectoral heterogeneity in forecasting the aggregate job vacancy rate, we now take a step further and focus directly on the sectoral job vacancy rates, with particular attention to manufacturing and services.

Figure 9: Sectoral job vacancy rate and FLP indicators

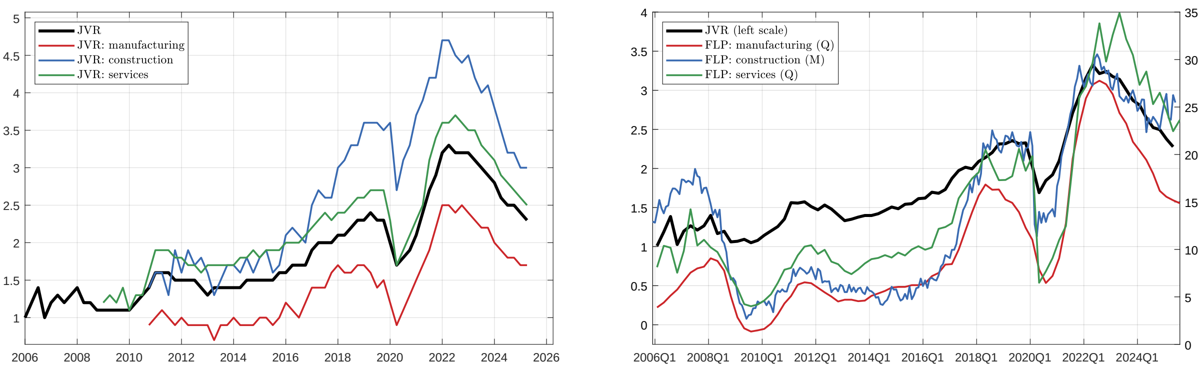


NOTES: The left panel of the figure plots the aggregate euro area job vacancy rate (black thick line) together with its decomposition by sector (coloured solid lines). The right panel of the figure plots the aggregate euro area job vacancy rate (black thick line) together with factors limiting production by sector (coloured solid lines, right scale).

Figure 9 displays the sectoral job vacancy rates (left panel) and the corresponding FLP indicators (right panel) for manufacturing, construction, and services. On the one hand, sectoral job vacancy rates largely mirror the dynamics of the aggregate job vacancy rate, with the notable exception of construction, which begins to rise earlier than the other sectors in the pre-COVID period. On the other hand, the sectoral FLP indicators, which constitute a key input in our forecasting model, exhibit markedly different patterns, especially during the COVID period: services experience a sharper surge and a more muted subsequent slowdown relative to manufacturing, while construction displays a persistent level that remains pronounced despite the easing observed in other sectors.

To assess how these sectoral differences translate into forecasting performance, we repeat the exercise in Section 5, now using sectoral job vacancy rates as the target variables and employing the corresponding sectoral FLP indicators, together with our indicator of labour demand from earnings calls as predictors.

Figure 10: MSFE across models: sectoral job vacancy rates

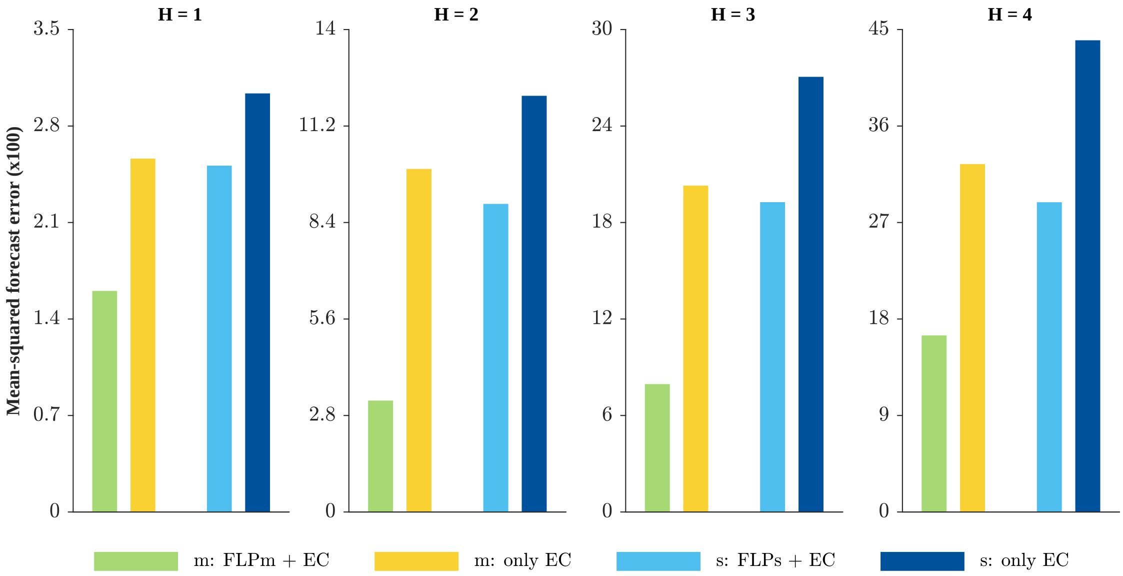


NOTES: The figure shows the MSFE (y-axis) of alternative models (bars) for the *post-COVID sample*, scaled by 100. The yellow and green bars correspond to a model targeting the job vacancy rate in manufacturing ($m$) and services ($s$) using only earnings calls data. Similarly, the blue and turquoise bars correspond to a model targeting the sectoral job vacancy rates using both earnings calls and sectoral FLPs. Each subplot reports results for different forecast horizons, from 1 quarter ($H = 1$) to 1 year ($H = 4$).

Figure 10 compares the forecasting performance of models targeting the manufacturing ($m$) and services ($s$) job vacancy rates, using either only the earnings-calls-derived indicator (yellow and green bars) or both the earnings-calls-derived indicator and the corresponding sectoral FLP indicators (blue and turquoise bars). As anticipated from previous results, models that combine the earnings-calls-derived indicator with sectoral

FLP indicators systematically outperform those relying solely on earnings calls information in both sectors.

More importantly, whether we consider the baseline model with only earnings calls or the richer specification including FLP indicators, the forecasting performance for the manufacturing job vacancy rate is considerably superior, by roughly 36 to 64% depending on the forecast horizon. This highlights two key insights: (i) information from earnings calls maps more closely into developments in manufacturing than in services; and (ii) sectoral FLP indicators carry substantially more predictive information for manufacturing than for services. As shown in Appendix C.3, these results hold also when we employ a longer forecast evaluation sample starting from 2016:M1.[8]

Taken together, these results point to meaningful sectoral heterogeneity: manufacturing dynamics are considerably easier to forecast than those in services, for which both soft and text-based indicators provide noisier signals. One possible explanation is the higher degree of heterogeneity within the services sector, which encompasses a broader range of occupations than manufacturing, each characterized by distinct dynamics (Forsells et al., 2019).

### 6.3 From aggregate to sectoral earnings calls

As a final exercise, we assess whether sectoral indicators of labour demand from earnings calls provide additional information beyond what is already embedded in the aggregate indicator. To this end, we construct separate indicators from our earnings calls data, exploiting the industry classification of the firms associated with each earnings call. Using the original TRBC[9] taxonomy, we map each firm to its corresponding NACE sector. This procedure yields a sectoral match for roughly 92% of the selected earnings calls, with the remaining 8% left unclassified due to missing data.

Figure 11 shows the sectoral decomposition of the earnings-calls-derived indicator. Clear differences are visible across sectors, particularly during the COVID period. The indicator of labour demand from earnings calls in the manufacturing sector rises with a slight delay relative to services, while services revert to their pre-COVID levels more quickly in the subsequent recovery. This raises the question of whether such differ-

[8]The underlying mean-squared forecast errors for both the aggregate and the sectoral job vacancy rates are reported in Table C5. We do not report Diebold-Mariano tests across sectoral job vacancy rates (manufacturing versus services), as the target series differ and the corresponding forecasts are therefore not directly comparable.

[9]Thomson Reuters Business Classification.

Figure 11: Sectoral earnings calls

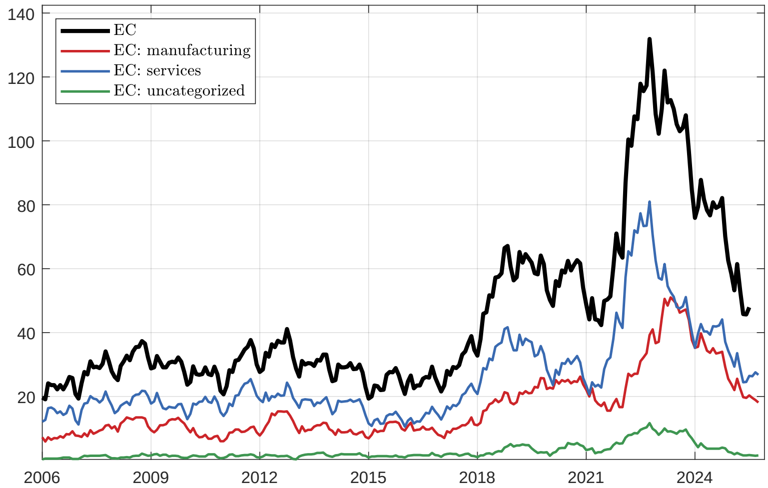


NOTES: The figure plots the aggregate earnings calls indicator (black solid line) together with its decomposition into manufacturing (red solid line), services (blue solid line) and uncategorized earnings calls (green solid line), i.e. data from earnings calls which cannot be attributed either to manufacturing or services.

ences translate into meaningful forecast improvements, or whether they simply capture idiosyncratic sectoral fluctuations already embedded in the aggregate indicator.

Figure 12: MSFE across models: sectoral earnings calls

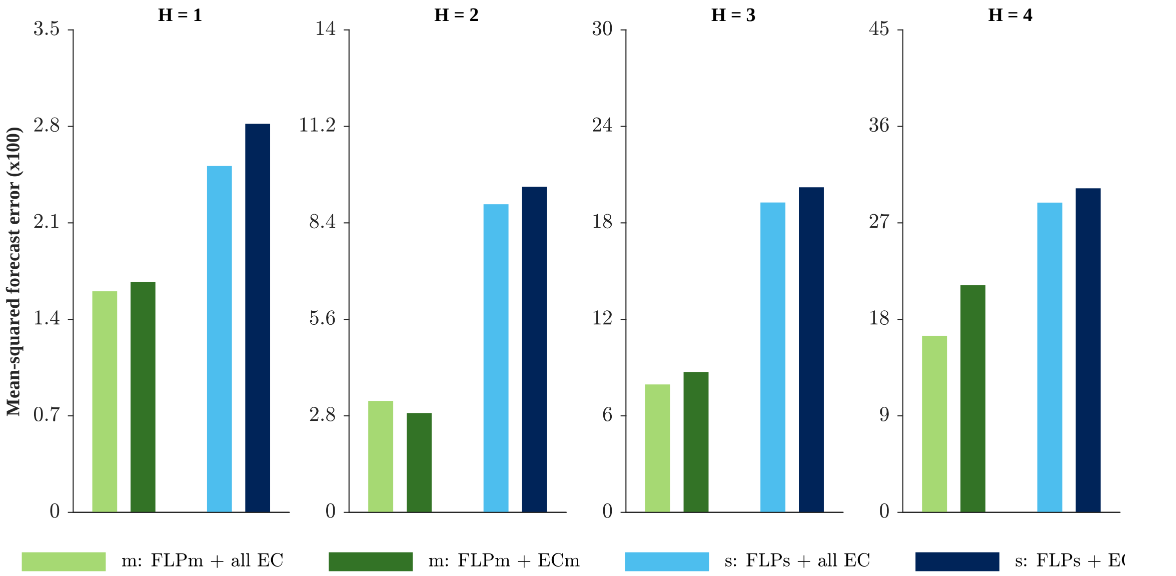


NOTES: The figure shows the MSFE (y-axis) of alternative models (bars) for the *post-COVID sample*, scaled by 100. The bright and dark green bars correspond to models targeting the job vacancy rate in manufacturing ($m$) using sectoral FLPs with aggregate and sectoral earnings calls data, respectively. Similarly, the blue and dark blue bars correspond to models targeting the job vacancy rates in services ($s$) using sectoral FLPs with aggregate and sectoral earnings calls data, respectively. Each subplot reports results for different forecast horizons, from 1 quarter ($H = 1$) to 1 year ($H = 4$).

To address this, Figure 12 compares the forecasting performance of models that include sectoral FLP indicators together with either aggregate or sectoral earnings-calls-derived indicators. The results show that aggregate earnings calls already incorporate all relevant information for forecasting, with no additional gains achieved by further disaggregating the indicator by sector. Appendix C.2 reports results obtained with an extended forecast sample starting from 2016:M1, showing only a marginal improvement in the forecast of the JVR in the services sector.

# 7 Conclusions

This paper evaluates the usefulness of textual information extracted from firms' earnings calls for forecasting the euro area job vacancy rate. Motivated by the fact that the euro area job vacancy rate is released with a lag, available over a limited time horizon and revised, we construct a novel monthly indicator of labour demand based on key specific references on labour market from the earnings call transcripts. We incorporate this textual indicator within a Mixed Frequency Bayesian VAR framework, together with traditional hard and soft indicators, to assess its contribution in terms of forecast performance.

Our analysis leads to three main findings. First, textual information extracted from firms' earnings calls meaningfully improves the real-time forecasting performance of the euro area job vacancy rate with respect to alternative monthly indicators, particularly in the most recent period, where labour demand has been rapidly declining after the post-pandemic surge. Among the high-frequency indicators included in the model, our indicator of labour demand derived from earnings calls provides the largest contribution to forecast accuracy. To obtain this result, we leverage information contained in alternative survey-based indicators, where FLPs were selected as the best to complement our indicator, according to their forecasting performance.

Second, real labour market indicators such as the unemployment rate do not materially improve forecast accuracy once the relevant soft indicators are included. This highlights the key role of timely, forward-looking qualitative information, rather than slower-moving hard data, in tracking labour demand conditions in the short- and medium-run.

Third, within our model, the bulk of the forecasting performance derives from information related to the manufacturing sector. Models combining earnings calls with sector specific FLP indicators perform at least as well as, and often better than, models using the full set of indicators. Furthermore, earnings calls data prove to be better suited in forecasting the sectoral job vacancy rate in manufacturing, rather than services. Finally, although more granular, sector-specific earnings calls do not provide meaningful gains beyond those achieved with aggregate earnings calls, indicating that the aggregate measure already captures the key information relevant for forecasting.

Taken together, these findings highlight the value of incorporating high-frequency text-based information into forecasting frameworks for euro area labour market conditions. Earnings calls provide timely signals that complement existing indicators, enhancing the ability to monitor shifts in labour demand and assess cyclical conditions in real time. This is ultimately relevant for the conduct of monetary policy in the euro area. As the availability and quality of text-based economic data continue to expand, our results suggest that such indicators will play an increasingly important role in supporting short-term assessment and policy analysis.

# References

Adrjan, P. and R. Lydon (2020): "Covid-19 and the global labour market: Impact on job postings," *Economic Letter*, 3, 1–7.

Aprigliano, V., S. Emiliozzi, G. Guaitoli, A. Luciani, J. Marcucci, and L. Monteforte (2023): "The power of text-based indicators in forecasting Italian economic activity," *International Journal of Forecasting*, 39, 791–808.

Arce, O., A. Consolo, A. Dias da Silva, and M. Mohr (2023): "More jobs but fewer working hours," *ECB Blog*, 7.

Autor, D. H. and D. Dorn (2013): "The growth of low-skill service jobs and the polarization of the US labor market," *American economic review*, 103, 1553–1597.

Bańbura, M., I. Belousova, K. Bodnár, and M. B. Tóth (2023): *Nowcasting employment in the euro area*, 2815, ECB Working Paper.

Barbaglia, L., S. Consoli, and S. Manzan (2024): "Forecasting GDP in Europe with textual data," *Journal of Applied Econometrics*, 39, 338–355.

Benigno, P. and G. B. Eggertsson (2023): "It's baaack: The surge in inflation in the 2020s and the return of the non-linear phillips curve," Tech. rep., National Bureau of Economic Research.

Blanchard, O. J. and P. Diamond (1989): "The Beveridge Curve," *Brookings Papers on Economic Activity*, 1989, 1–76.

Canova, F. and F. Ferroni (2021): "A hitchhiker's guide to empirical macro models," Working Paper 2021-15, Federal Reserve Bank of Chicago.

Chudik, A., V. Grossman, and M. H. Pesaran (2016): "A multi-country approach to forecasting output growth using PMIs," *Journal of Econometrics*, 192, 349–365.

Clements, M. P. and A. B. Galvão (2009): "Forecasting US output growth using leading indicators: An appraisal using MIDAS models," *Journal of Applied Econometrics*, 24, 1187–1206.

Consolo, A. and A. D. Da Silva (2019): "The euro area labour market through the lens of the Beveridge curve," *Economic Bulletin Articles*, 4.

Consolo, A., C. Foroni, and C. Martínez Hernández (2023): "A mixed frequency BVAR for the euro area labour market," *Oxford Bulletin of Economics and Statistics*, 85, 1048–1082.

de Bondt, G. and L. Saiz (2024): "Is the PMI a reliable indicator for nowcasting euro area real GDP?" *Economic Bulletin Boxes*, 1.

de Bondt, G. J. (2019): "A PMI-based real GDP tracker for the euro area," *Journal of Business Cycle Research*, 15, 147–170.

Diebold, F. X. and R. S. Mariano (2002): "Comparing predictive accuracy," *Journal of Business & economic statistics*, 20, 134–144.

Dueholm, M., A. Kalyani, and S. Ozkan (2024): "Can Earnings Calls Be Used to Gauge Labor Market Tightness?" Tech. rep., Federal Reserve Bank of St. Louis.

Foroni, C. and M. Marcellino (2014): "Mixed-frequency structural models: Identification, estimation, and policy analysis," *Journal of Applied Econometrics*, 29, 1118–1144.

Forsells, M., N. Kennedy, and L. M. Timm (2019): "Developments in the services sector and its relationship with manufacturing," *Economic Bulletin Boxes*, 7.

Ghysels, E. (2016): "Macroeconomics and the reality of mixed frequency data," *Journal of Econometrics*, 193, 294–314.

Ghysels, E. and M. Marcellino (2016): "The econometric analysis of mixed frequency data sampling," *Journal of Econometrics*, 193, 291–293.

Ghysels, E., P. Santa-Clara, and R. Valkanov (2004): "The MIDAS touch: Mixed data sampling regression models," *CIRANO Working Papers*.

——— (2005): "There is a risk-return trade-off after all," *Journal of financial economics*, 76, 509–548.

Giannone, D., L. Reichlin, and S. Simonelli (2009): "Nowcasting euro area economic activity in real time: the role of confidence indicators," *National Institute Economic Review*, 210, 90–97.

Gogas, P., T. Papadimitriou, and E. Sofianos (2022): "Forecasting unemployment in the euro area with machine learning," *Journal of Forecasting*, 41, 551–566.

Hassan, T. A., S. Hollander, L. V. Lent, and A. Tahoun (2024): "The global impact of Brexit uncertainty," *The Journal of Finance*, 79, 413–458.

Hassan, T. A., S. Hollander, L. Van Lent, and A. Tahoun (2019): "Firm-level political risk: Measurement and effects," *The quarterly journal of economics*, 134, 2135–2202.

Kalamara, E., A. Turrell, C. Redl, G. Kapetanios, and S. Kapadia (2022): "Making text count: economic forecasting using newspaper text," *Journal of Applied Econometrics*, 37, 896–919.

Kalyani, A., N. Bloom, M. Carvalho, T. Hassan, J. Lerner, and A. Tahoun (2025): "The diffusion of new technologies," *The Quarterly Journal of Economics*, 140, 1299–1365.

Kuzin, V., M. Marcellino, and C. Schumacher (2011): "MIDAS vs. mixed-frequency VAR: Nowcasting GDP in the euro area," *International Journal of Forecasting*, 27, 529–542.

Lahiri, K. and G. Monokroussos (2013): "Nowcasting US GDP: The role of ISM business surveys," *International Journal of Forecasting*, 29, 644–658.

Ngai, L. R. and C. A. Pissarides (2007): "Structural change in a multisector model of growth," *American economic review*, 97, 429–443.

Pissarides, C. A. (2011): "Equilibrium in the labor market with search frictions," *American Economic Review*, 101, 1092–1105.

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

Sims, C. A. and T. Zha (1998): "Bayesian methods for dynamic multivariate models," *International Economic Review*, 39, 949–968.

# A Data description

## A.1 Earnings calls

Table A1: Descriptive statistics

| | Total | Industry | Services | Missing |
|---|---|---|---|---|
| Number of earnings calls | 37,935.00 | 15,897.00 | 21,233.00 | 805.00 |
| Share of total | 1.00 | 0.42 | 0.56 | 0.02 |
| Share from country of headquarters | | | | |
| -Germany | 0.22 | 0.27 | 0.18 | 0.07 |
| -France | 0.15 | 0.12 | 0.17 | 0.05 |
| -Italy | 0.10 | 0.10 | 0.09 | 0.08 |
| -Spain | 0.08 | 0.05 | 0.10 | 0.06 |
| -Netherlands | 0.10 | 0.09 | 0.10 | 0.06 |
| Share from economic sector | | | | |
| -Basic Materials | 0.10 | 0.22 | 0.00 | |
| -Consumer Cyclicals | 0.13 | 0.14 | 0.12 | |
| -Consumer Non-Cyclicals | 0.05 | 0.05 | 0.05 | |
| -Energy | 0.06 | 0.08 | 0.05 | |
| -Financials | 0.15 | 0.00 | 0.27 | |
| -Healthcare | 0.10 | 0.13 | 0.07 | |
| -Industrials | 0.18 | 0.17 | 0.19 | |
| -Real Estate | 0.04 | 0.00 | 0.07 | |
| -Technology | 0.15 | 0.10 | 0.19 | |
| -Utilities | 0.05 | 0.10 | 0.00 | |
| Share from business sector | | | | |
| -Applied Resources | 0.02 | 0.04 | 0.00 | |
| -Automobiles & Auto Parts | 0.03 | 0.07 | 0.00 | |
| -Banking & Investment Services | 0.11 | 0.00 | 0.19 | |
| -Chemicals | 0.04 | 0.09 | 0.00 | |
| -Consumer Goods Conglomerates | 0.00 | 0.00 | 0.00 | |
| -Cyclical Consumer Products | 0.03 | 0.07 | 0.00 | |
| -Cyclical Consumer Services | 0.05 | 0.00 | 0.09 | |
| -Energy (Fossil Fuels) | 0.06 | 0.07 | 0.05 | |
| -Financial Technology (Fintech) & Infrastructure | 0.00 | 0.00 | 0.00 | |
| -Food & Beverages | 0.02 | 0.04 | 0.01 | |
| -Food & Drug Retailing | 0.02 | 0.00 | 0.03 | |
| -Healthcare Services & Equipment | 0.04 | 0.05 | 0.03 | |
| -Industrial & Commercial Services | 0.06 | 0.00 | 0.11 | |
| -Industrial Goods | 0.07 | 0.17 | 0.00 | |
| -Insurance | 0.04 | 0.00 | 0.07 | |
| -Investment Holding Companies | 0.01 | 0.00 | 0.01 | |
| -Mineral Resources | 0.04 | 0.09 | 0.00 | |
| -Personal & Household Products & Services | 0.00 | 0.01 | 0.00 | |
| -Pharmaceuticals & Medical Research | 0.06 | 0.08 | 0.04 | |
| -Real Estate | 0.04 | 0.00 | 0.07 | |
| -Renewable Energy | 0.00 | 0.01 | 0.00 | |
| -Retailers | 0.02 | 0.00 | 0.04 | |
| -Software & IT Services | 0.05 | 0.00 | 0.10 | |
| -Technology Equipment | 0.04 | 0.10 | 0.00 | |
| -Telecommunications Services | 0.05 | 0.00 | 0.09 | |
| -Transportation | 0.04 | 0.00 | 0.08 | |
| -Utilities | 0.05 | 0.10 | 0.00 | |

Notes:

## A.2 Additional variables

Figure A1: Job vacancy rate and sentiment/survey indicators

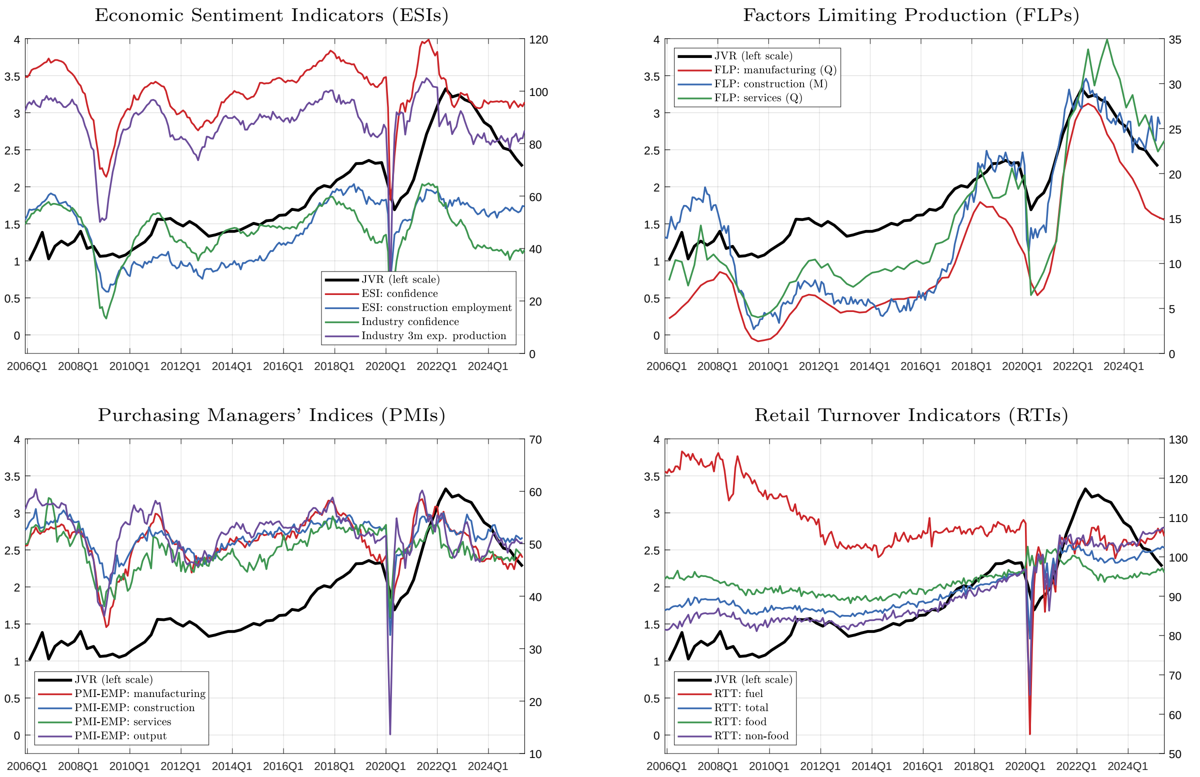


**Notes**: The figure plots the job vacancy rate from 2006:Q1 to 2025:Q2 (black thick line, left scale) together with the additional quarterly and/or monthly indicators used in each submodel (coloured thin lines, right scale).

# B Forecasting the euro area job vacancy rate: additional results

## B.1 Model comparison

Table B2: Forecasting performance: FLPs vs. other indicators

| | H=1 | H=2 | H=3 | H=4 |
|---|---|---|---|---|
| | *Post-COVID sample* | | | |
| **ESI** | 0.920 | 0.700* | 0.652* | 0.665* |
| **PMI** | 0.840 | 0.634** | 0.590* | 0.602* |
| **RTI** | 0.833 | 0.633** | 0.615* | 0.662* |
| | *Full sample* | | | |
| **ESI** | 1.138 | 0.962 | 0.944 | 0.968 |
| **PMI** | 1.167 | 0.940 | 0.910 | 0.929 |
| **RTI** | 0.766 | 0.910 | 0.919 | 0.989 |

NOTES: The table reports the relative mean-squared forecast error, over different horizons (columns), for our baseline MF-BVAR model with FLPs with respect to the same model including alternative sets of soft indicators (rows). Results are reported for both the *post-COVID sample* (upper panel) and for the *full sample* (bottom panel). ***, **, * denote significance at 1%, 5%, 10% level for the Diebold and Mariano (2002) test.

Table B3: MSFE by month within quarter

| | H=1 | | | H=2 | | | H=3 | | | H=4 | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | *Post-COVID sample* | | | | | | | | | | | |
| | M1 | M2 | M3 | M1 | M2 | M3 | M1 | M2 | M3 | M1 | M2 | M3 |
| **ESI** | 0.014 | 0.012 | 0.017 | 0.069 | 0.062 | 0.068 | 0.159 | 0.150 | 0.149 | 0.291 | 0.274 | 0.263 |
| **FLP** | 0.015 | 0.014 | 0.011 | 0.055 | 0.049 | 0.035 | 0.120 | 0.111 | 0.066 | 0.218 | 0.201 | 0.129 |
| **PMI** | 0.016 | 0.013 | 0.019 | 0.078 | 0.065 | 0.076 | 0.182 | 0.162 | 0.164 | 0.325 | 0.301 | 0.289 |
| **RTI** | 0.018 | 0.012 | 0.018 | 0.085 | 0.062 | 0.073 | 0.184 | 0.153 | 0.149 | 0.311 | 0.270 | 0.251 |
| | *Full sample* | | | | | | | | | | | |
| | M1 | M2 | M3 | M1 | M2 | M3 | M1 | M2 | M3 | M1 | M2 | M3 |
| **ESI** | 0.012 | 0.014 | 0.017 | 0.055 | 0.057 | 0.061 | 0.116 | 0.123 | 0.131 | 0.201 | 0.223 | 0.226 |
| **FLP** | 0.012 | 0.020 | 0.015 | 0.051 | 0.064 | 0.051 | 0.112 | 0.126 | 0.111 | 0.199 | 0.227 | 0.202 |
| **PMI** | 0.012 | 0.014 | 0.015 | 0.058 | 0.057 | 0.062 | 0.125 | 0.126 | 0.133 | 0.216 | 0.229 | 0.232 |
| **RTI** | 0.014 | 0.013 | 0.036 | 0.059 | 0.056 | 0.067 | 0.122 | 0.122 | 0.137 | 0.203 | 0.217 | 0.216 |

NOTES: The table reports the mean-squared forecast error, over different horizons (columns), for all the submodels estimated in the paper (rows). The mean-squared forecast error is reported for each month (sub-columns) in which the latest high-frequency observation is available for the forecast. Results are reported for both the *post-COVID sample* (upper panel) and for the *full sample* (bottom panel).

## B.2 The role of real activity and earnings calls: extended forecast sample

Figure B2: MSFE across models: the role of real activity and earnings calls, full sample

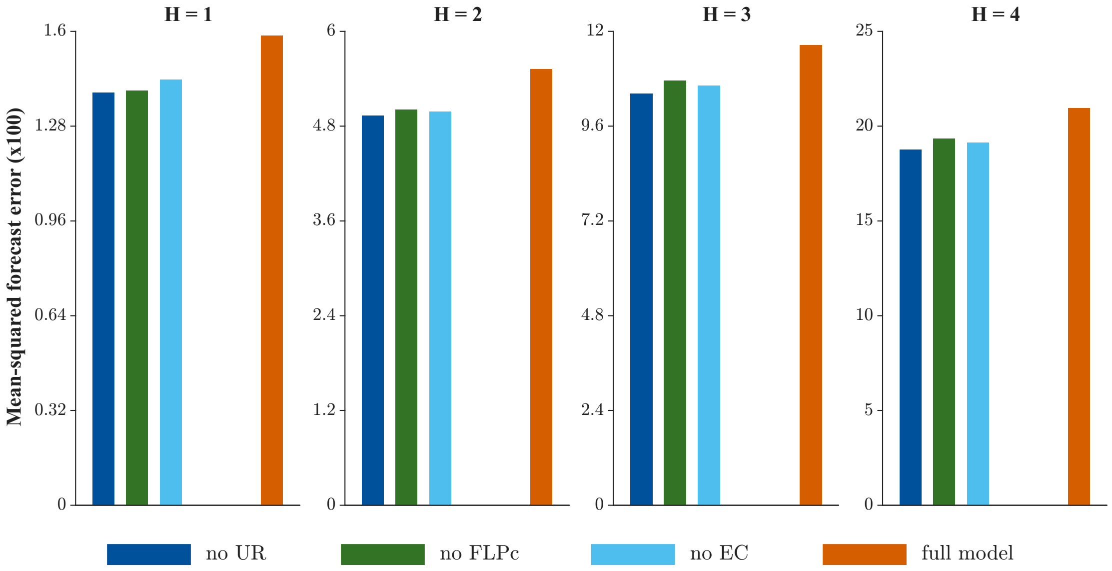


NOTES: The figure shows the MSFE (y-axis) of alternative models (bars) for the *full sample*, scaled by 100 for better interpretability. Each coloured bar corresponds to one of the models outlined above, while the gray bar is a benchmark model only including earnings calls as an auxiliary indicator. Each subplot reports results for different forecast horizons, from 1 quarter ($H = 1$) to 1 year ($H = 4$). The forecast evaluation is carried out on an evaluation sample starting from 2016:M1.

Table B4: Forecasting performance: the role of real activity and earnings calls

| | H=1 | H=2 | H=3 | H=4 |
|---|---|---|---|---|
| | *Post-COVID sample* | | | |
| **no UR** | 1.090* | 1.165 | 1.210 | 1.202 |
| **no FLPc** | 1.035 | 1.107 | 1.134 | 1.090 |
| **no EC** | 0.986 | 1.055 | 1.060 | 1.071 |
| **GDP** | 1.122 | 1.320 | 1.261 | 1.243 |
| | *Full sample* | | | |
| **no UR** | 1.137 | 1.120 | 1.117 | 1.117 |
| **no FLPc** | 1.132 | 1.101 | 1.083 | 1.084 |
| **no EC** | 1.103 | 1.108* | 1.096 | 1.096 |
| **GDP** | 1.014 | 0.866 | 0.764 | 0.729 |

NOTES: The table reports the relative mean-squared forecast error, over different horizons (columns), for our baseline MF-BVAR model with FLPs with respect to four different specifications: (i) excluding UR, (ii) excluding construction FLPs, (iii) excluding EC and (iv) including GDP in place of UR. Results are reported for both the *post-COVID sample* (upper panel) and for the *full sample* (bottom panel). ***, **, * denote significance at 1%, 5%, 10% level for the Diebold and Mariano (2002) test.

## B.3 GDP as alternative real indicator

Figure B3: MSFE across models: GDP vs unemployment

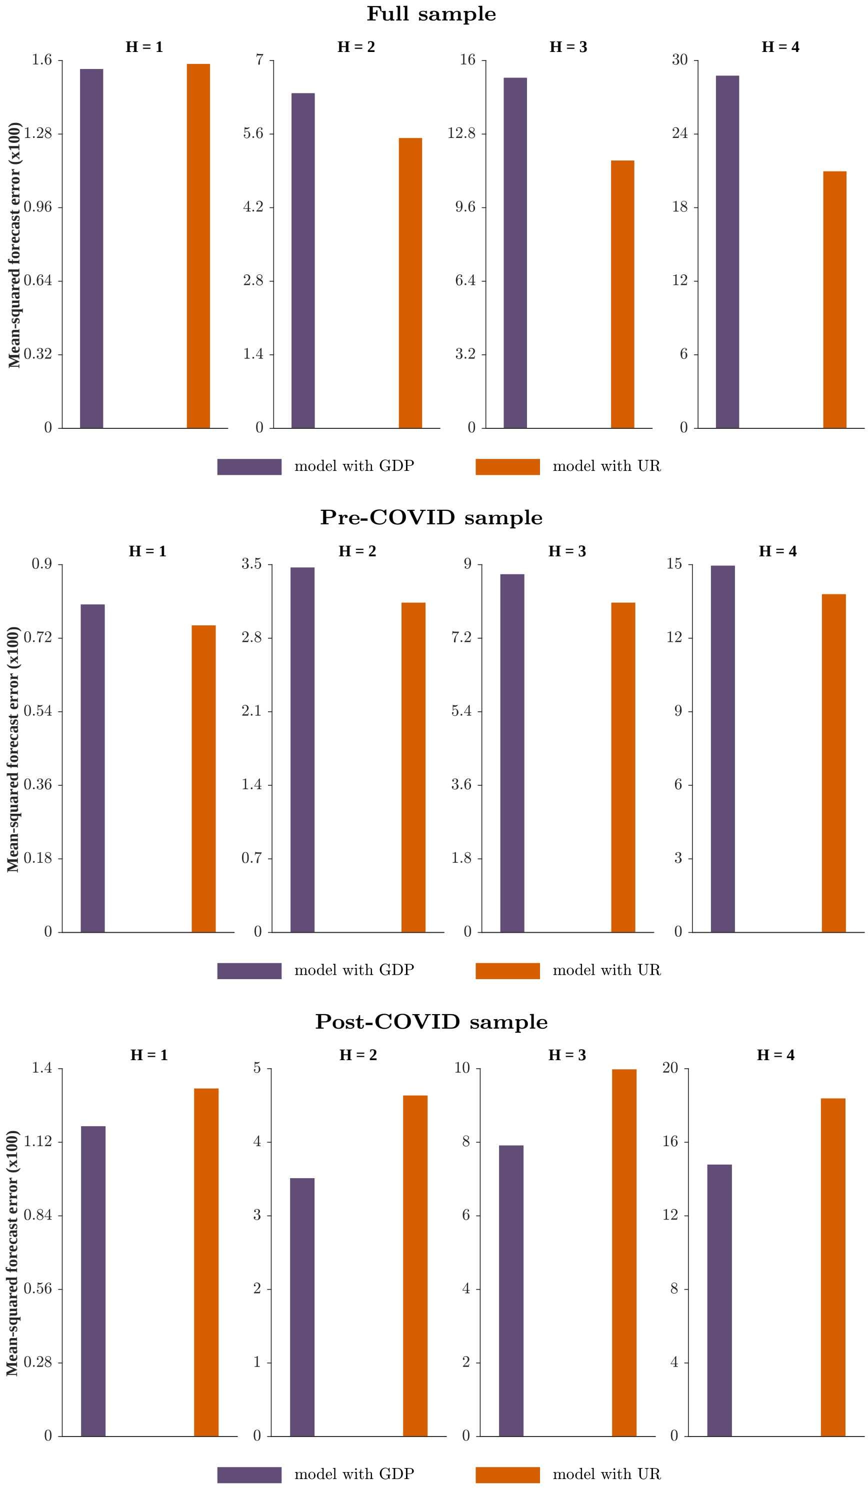


NOTES: The figure shows the MSFE (y-axis) of our benchmark model, including unemployment rate (red bar), against the one of a model including GDP (dark purple bar). The MSFE is scaled by 100 for better interpretability. Each subplot reports results for different forecast horizons, from 1 quarter ($H = 1$) to 1 year ($H = 4$). The forecast evaluation is carried on three evaluation samples: (i) a sample starting from 2016:M1, (ii) a sample starting in 2016:M1 and ending in 2019:M12 and (iii) a sample starting in 2021:M1.

# C Sectoral heterogeneity: additional results

## C.1 Decomposition of FLPs: extended forecast sample

Figure C4: MSFE across models: sectoral decomposition of FLPs, full sample

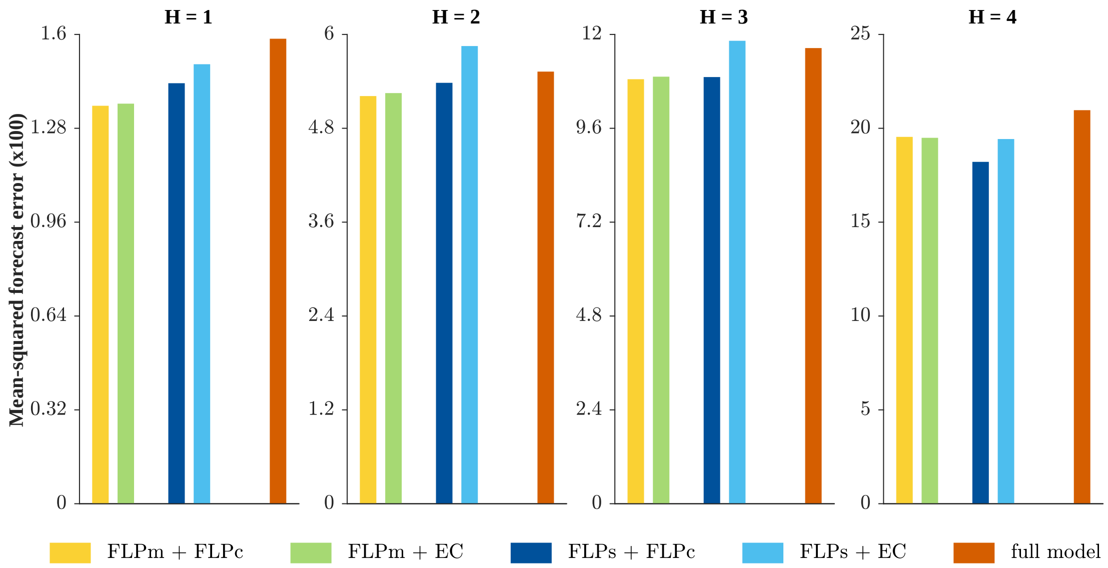


NOTES: The figure shows the MSFE (y-axis) of alternative models (bars) for the *full sample*, scaled by 100 for better interpretability. Each coloured bar corresponds to one of the models outlined above, while the gray bar is a benchmark model only including earnings calls as an auxiliary indicator. Each subplot reports results for different forecast horizons, from 1 quarter ($H = 1$) to 1 year ($H = 4$). The forecast evaluation is carried out on an evaluation sample starting from 2016:M1.

## C.2 Sectoral job vacancy rate: extended forecast sample

Figure C5: MSFE across models: sectoral job vacancy rate, full sample

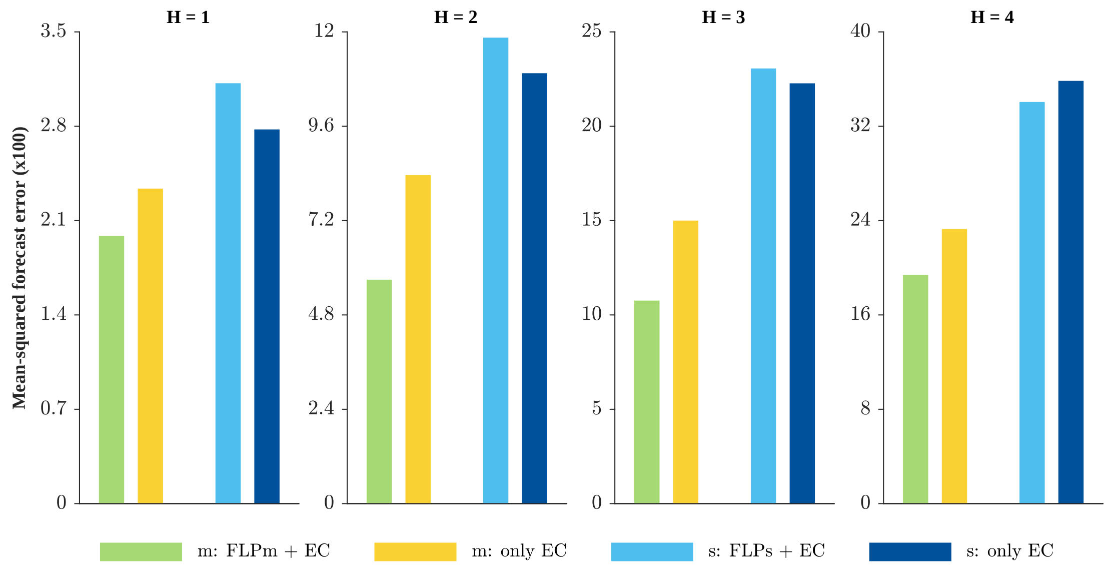


NOTES: The figure shows the MSFE (y-axis) of alternative models (bars) for the *full sample*, scaled by 100. The yellow and green bars correspond to a model targeting the job vacancy rate in manufacturing ($m$) and services ($s$) using only earnings calls data. Similarly, the blue and turquoise bars correspond to a model targeting the sectoral job vacancy rates using both earnings calls and sectoral FLPs. Each subplot reports results for different forecast horizons, from 1 quarter ($H = 1$) to 1 year ($H = 4$).

## C.3 Sectoral earnings calls: extended forecast sample

Figure C6: MSFE across models: sectoral earnings calls, full sample

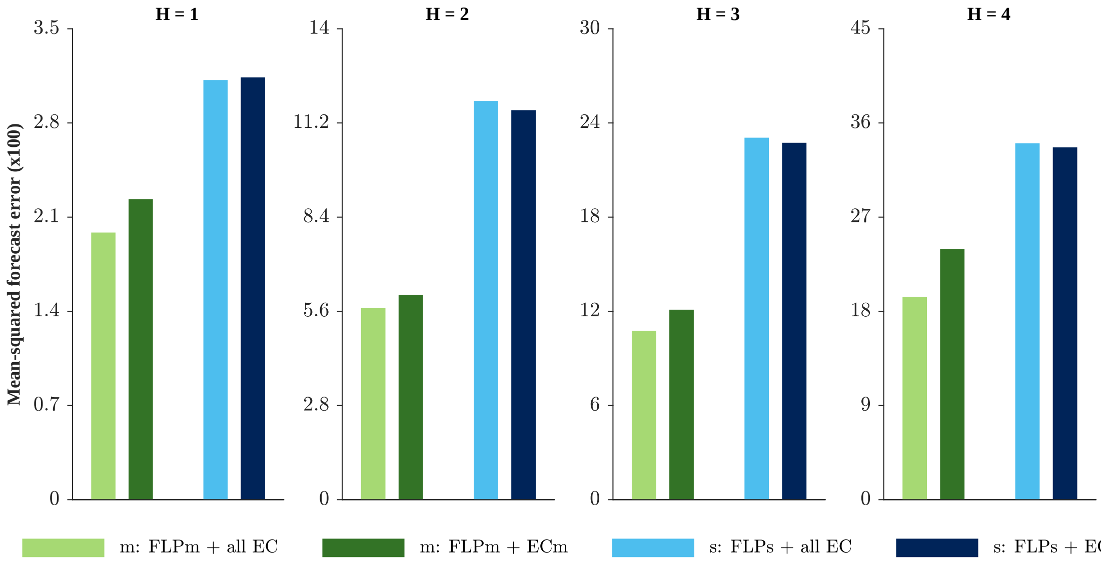


NOTES: The figure shows the MSFE (y-axis) of alternative models (bars) for the *full sample*, scaled by 100. The yellow and green bars correspond to a model targeting the job vacancy rate in manufacturing ($m$) using sectoral FLPs with aggregate and sectoral earnings calls data, respectively. Similarly, the blue and turquoise bars correspond to a model targeting the job vacancy rates in services ($s$) using sectoral FLPs with aggregate and sectoral earnings calls data, respectively. Each subplot reports results for different forecast horizons, from 1 quarter ($H = 1$) to 1 year ($H = 4$).

Table C5: Forecasting performance: sectoral and aggregate JVR

| **Panel A: Aggregate JVR with sectoral FLPs** | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | *Post-COVID* | | | | *Full sample* | | | |
| Model | $h = 1$ | $h = 2$ | $h = 3$ | $h = 4$ | $h = 1$ | $h = 2$ | $h = 3$ | $h = 4$ |
| FLPm + FLPc | 0.014 | 0.050 | 0.103 | 0.188 | 0.014 | 0.052 | 0.109 | 0.195 |
| FLPm + EC | 0.012 | 0.045 | 0.094 | 0.176 | 0.014 | 0.053 | 0.109 | 0.195 |
| FLPs + FLPc | 0.013 | 0.053 | 0.109 | 0.180 | 0.014 | 0.054 | 0.109 | 0.182 |
| FLPs + EC | 0.015 | 0.061 | 0.125 | 0.207 | 0.015 | 0.059 | 0.118 | 0.194 |
| **Panel B: Sectoral JVR** | | | | | | | | |
| | *Manufacturing* | | | | *Services* | | | |
| Model | $h = 1$ | $h = 2$ | $h = 3$ | $h = 4$ | $h = 1$ | $h = 2$ | $h = 3$ | $h = 4$ |
| | | | | | *Post-COVID sample* | | | |
| only EC | 0.026 | 0.100 | 0.203 | 0.325 | 0.030 | 0.121 | 0.271 | 0.440 |
| FLP + all EC | 0.016 | 0.032 | 0.080 | 0.165 | 0.025 | 0.089 | 0.193 | 0.289 |
| FLP + sect. EC | 0.017 | 0.029 | 0.087 | 0.212 | 0.028 | 0.094 | 0.202 | 0.302 |
| | | | | | *Full sample* | | | |
| only EC | 0.023 | 0.084 | 0.150 | 0.233 | 0.028 | 0.109 | 0.223 | 0.358 |
| FLP + all EC | 0.020 | 0.057 | 0.108 | 0.194 | 0.031 | 0.119 | 0.231 | 0.341 |
| FLP + sect. EC | 0.022* | 0.061 | 0.121 | 0.240 | 0.031 | 0.116 | 0.227 | 0.337 |

NOTES: The table reports the mean-squared forecast error for different models (rows) over different horizons (columns) for both the aggregate job vacancy rate (top panel) and sectoral job vacancy rates in manufacturing and services (bottom panel). Stars in Panel A denote Diebold-Mariano significance versus the full model (FLP, UR and EC); in Panel B (FLP + sect. EC row) versus FLP + all EC within the same sector. Comparisons across sectors (manufacturing versus services) have no DM test as the target series differ. ***, **, * denote significance at the 1%, 5%, 10% level for the Diebold and Mariano (2002) test.

**Acknowledgements**

We thank the participants of the Eurosystem's Working Group on Forecasting for valuable comments and feedback.
This paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB. Any remaining errors are our own.

**Agostino Consolo**
European Central Bank, Frankfurt am Main, Germany; email: agostino.consolo@ecb.europa.eu

**Claudia Foroni**
European Central Bank, Frankfurt am Main, Germany; email: claudia.foroni@ecb.europa.eu

**Claudio Lissona (corresponding author)**
European Central Bank, Frankfurt am Main, Germany; email: claudio.lissona@gmail.com

**Christofer Schroeder**
European Central Bank, Frankfurt am Main, Germany; email: christofer.schroeder@ecb.europa.eu



Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.ecb.europa.eu



PDF ISBN 978-92-899-7999-3 ISSN 1725-2806 doi:10.2866/9269320 QB-01-26-197-EN-N