A Work Project, presented as part of the requirements for the Award of a Master's degree in Economics from the Nova School of Business and Economics.

FORECASTING PORTUGUESE GDP GROWTH USING U-MIDAS AND VAR MODELS: A COMPARATIVE ANALYSIS

António Ramalho Lourinho Santana Fernandes

Work project carried out under the supervision of:

Paulo M. M. Rodrigues

20/01/2025

Abstract

This study examines the benefits of using high-frequency data in forecasting Portuguese GDP growth, comparing U-MIDAS and VAR models. U-MIDAS integrates monthly data for flexibility, while VAR uses quarterly data. Using data from 2000Q1 to 2020Q4 for model estimation, forecasts cover 2021Q1 to 2024Q1, a period of COVID-19 recovery and inflationary pressures. Key indicators include industrial production, unemployment rate, PSI-20 index, and exports. The results show that U-MIDAS performs better in short- and long-term forecasts, while VAR excels in the medium term. These findings are statistically compared using methods such as the Diebold-Mariano test, and the Clark-West test.

Keywords: Forecasting, mixed data sampling, high-frequency data, U-MIDAS, VAR, Portuguese GDP, nowcasting

Acknowledgments

I want to express my sincere gratitude to my supervisor, Professor Paulo M. M. Rodrigues, for his constant guidance and thoughtful support throughout this journey. I also extend my deepest thanks to my family and friends for always supporting me during these months.

This work used infrastructure and resources funded by Fundação para a Ciência e a Tecnologia (UID/ECO/00124/2013, UID/ECO/00124/2019 and Social Sciences DataLab, Project 22209), POR Lisboa (LISBOA-01-0145-FEDER-007722 and Social Sciences DataLab, Project 22209) and POR Norte (Social Sciences DataLab, Project 22209).

1. INTRODUCTION

Gross Domestic Product (GDP) is an essential measure of economic performance, serving as a benchmark for assessing growth, stability, and overall economic health. Accurate GDP forecasting is essential for timely decision-making, enabling policymakers to respond proactively to shifts in the economic cycle. However, a key challenge in forecasting GDP is its low frequency, as official GDP figures are typically released every quarter and often with a significant time lag. To address this issue, high-frequency data, such as monthly indicators, has become increasingly valuable in economic forecasting. One of the primary challenges in integrating high-frequency data into GDP forecasts is the unbalanced dataset that arises when using monthly data, for example, to predict a quarterly figure. A straightforward solution is aggregating monthly data to match the quarterly frequency, creating a balanced dataset. However, this approach leads to the loss of valuable information as the nuanced monthly variations that could enhance forecast accuracy are smoothed out.

Utilising high-frequency data for GDP forecasting is advantageous for a small, open economy like Portugal. The Portuguese economy encounters unique challenges, including high levels of public debt, a heavy reliance on exports, and vulnerability to external economic fluctuations, making GDP forecasting essential and complex. High-frequency data provides a helpful solution by delivering timely insights that capture real-time variations in crucial areas, including exports, financial markets, and industrial production. These sectors play a critical role in the Portuguese economy, improving the precision of GDP forecasts and providing a clearer understanding of Portugal's economic dynamics in the face of an uncertain global and regional landscape.

This work project seeks to demonstrate the advantages of high-frequency data in forecasting by comparing distinct methods that operate different data frequencies. Specifically, we employ two forecasting methodologies using a dataset that encompasses Portuguese economic indicators frequently used to forecast GDP growth over short-term (1-step ahead) medium-term (2-step ahead) and longer-term (4-step ahead) horizons. It includes four key explanatory variables – industrial production, unemployment rate, PSI-20 index, and exports – and the target variable, real GDP growth. These indicators were chosen to reflect critical aspects of the Portuguese economy, thereby improving the forecasts' reliability.

The literature offers several methods to study and optimise GDP forecasting with high-frequency indicators. The Mixed Data Sampling (MIDAS) model, introduced by Ghysels, Santa-Clara, and Valkanov (2004), directly integrates monthly or weekly data to predict quarterly GDP. Extensions like Unrestricted MIDAS (U-MIDAS) (Foroni et al. 2015) and Factor-MIDAS refine this approach: U-MIDAS allows for flexible parameter estimation, while Factor-MIDAS simplifies by reducing multiple high-frequency indicators into a few factors (Marcellino and Schumacher 2010).

In contrast, traditional Vector Autoregressive (VAR) models, as Stock and Watson (2001) discussed, typically use only low-frequency data, such as quarterly indicators, to capture relationships over time. The mixed-frequency VAR (MF-VAR), however, extends this approach by integrating both high- and low-frequency data, allowing monthly data to inform quarterly forecasts (Schorfheide and Song 2015).

This study explores the effectiveness of two prominent forecasting models in predicting GDP growth in Portugal. The first model, U-MIDAS, is grounded in the methods proposed by Foroni et al. (2015) and utilises high-frequency monthly data. The methodology employed enables estimation through Ordinary Least Squares (OLS), particularly when the frequency gap – such as that between quarterly and monthly data – is minimal. In contrast to the conventional MIDAS model, U-MIDAS enhances flexibility by omitting the polynomial weighting function. The second model utilised is a VAR, which relies exclusively on quarterly data to estimate its GDP growth forecasts.

We estimate the models' using data from 2000Q1 to 2020Q4 and then compute out-of-sample forecasts from 2021Q1 through 2024Q1. This period is marked by the economic rebound from the COVID-19 crisis and the inflationary pressures from the Russian invasion to Ukraine, adding extra difficulty to the forecasting process and highlighting the value of monthly data for tracking these changes.

To evaluate the models' accuracy, we focus on their out-of-sample performance across the three forecast horizons. Key metrics, including the Root Mean Square Error (RMSE) and Mean Absolute Error (MAE), are used to measure forecast accuracy. Additionally, we employ the Diebold-Mariano (DM) and Clark-West (CW) tests to compare the predictive performance of the models, identifying which one delivers superior out-of-sample forecasts.

This research also incorporates an analysis aimed at "nowcasting" Portuguese quarterly GDP growth using the U-MIDAS model. It investigates whether monthly data can enhance GDP estimates prior to the end of the quarter. Predictions were generated in stages: initially, using data from the first month only, then including data from the second month, and ultimately incorporating data from all three months. The objective was to ascertain whether more accurate predictions could be made earlier in the quarter, providing insights ahead of the official GDP release.

2. LITERATURE REVIEW

Forecasting time series data has been a critical focus in economic research for decades. The 1970s marked significant progress with the development of models like ARMA, laying the foundation for many of the methods used today. However, a significant leap forward came with the introduction of Vector Autoregressive models by Sims (1980), which allowed economists to model multiple economic variables simultaneously. This capability made VAR highly effective for macroeconomic analysis, establishing it as a standard tool for economic forecasting. Recently, another approach known as Mixed Data Sampling has gained reputation. Initially proposed by Ghysels et al. (2004) and further developed by Ghysels et al. (2007), MIDAS allows the sampling of independent and dependent variables at different frequencies without causing parameter proliferation by utilising a distributed lag polynomial. This flexibility has proven particularly useful when higher-frequency variables are expected to improve forecast accuracy.

Ghysels et al. (2004) method was originally designed for financial applications, such as assessing risk-return trade-offs and forecasting volatility. However, recent research has shown its effectiveness in predicting macroeconomic variables, including GDP, inflation, and private consumption. MIDAS has emerged as a valuable tool for these applications, with numerous studies utilising it for GDP growth forecasts. Furthermore, the comparison of MIDAS with low-frequency models has gained significant attention in the literature, especially concerning various European and international economies.

For example, Kuzin et al. (2009) examined the effectiveness of MIDAS and mixed-frequency VAR techniques in nowcasting the Euro Area GDP. They found that both methods complemented each other: MIDAS delivered better results for shorter time horizons, whereas the mixed-frequency VAR was more effective for long-term forecasts. Similarly, Clements and Galvão (2009) used MIDAS to forecast US output growth, focusing on how well leading indicators could predict economic activity. Bhaghoe et al. (2019) also applied a MIDAS model to estimate year-to-year quarterly real GDP growth in Suriname. In line with this, Zuanazzi and Ziegelmann (2014) argue that MIDAS and U-MIDAS deliver better results than ARMA when applied to Brazil's GDP. For inflation forecast, Li et al. (2015) demonstrate that MIDAS surpasses ARIMA in inflation prediction accuracy when daily Google search trends are used as early indicators in China.

In the context of the Portuguese economy, the implementation of the MIDAS approach and its comparison to low-frequency models has been relatively limited. Duarte et al. (2016) used MIDAS to forecast Portuguese private consumption, taking advantage of high-frequency ATM/POS data. Similarly, Dias et al. (2014) applied factor models to forecast Portuguese GDP growth, focusing on pre- and post-crisis periods.

One handy extension of the MIDAS model is the U-MIDAS approach. Foroni et al. (2015) found that U-MIDAS often performs better than the standard MIDAS model when the frequency gap between the variables is not too large, such as when working with quarterly and monthly data. U-MIDAS is based on linear lag polynomials, which makes it easier to estimate using OLS, simplifying the computational process while maintaining accuracy.

The VAR serves as a natural benchmark for comparing the performance of the U-MIDAS approach. Stock and Watson (2001) showed that VAR models are particularly effective for short- and medium-term forecasting of unemployment, inflation, and interest rates, making them reliable tools in macroeconomic analysis. Additionally, Sdrakas and Viguie (2003) showed that VAR models outperform a single autoregressive model in the short-term forecast for the Euro Area GDP. Given the VAR model's strong performance in forecasting macroeconomic variables and its noted stability, using it as a benchmark will allow for a meaningful comparison with the U-MIDAS's performance, especially in leveraging high-frequency data.

Another critical element of this research is the comparison of different forecasting horizons. Clements and Galvão (2009) explored periods of up to one year (1, 2, and 4 steps ahead), and Foroni et al. (2015) computed their forecasts of US and Euro Area GDP growth to four quarters ahead. We will examine three-time horizons: 1-quarter ahead, 2-quarters ahead, and 4-quarters ahead, similar to the methodology used by Andreou et al. (2010) in forecasting US real GDP growth and by Clements and Galvão (2009). The comparison between time frames will provide valuable insights into how each model handles different periods.

To handle the mixed-frequency data, the methodology outlined by Cottrell and Lucchetti (2017) was followed, which jointly models quarterly and monthly data. Specifically, quarterly GDP growth is explained using monthly variables, and Stata is used to estimate the models and generate the forecasts.

Ghysels et al. (2018) emphasise the critical role of selecting appropriate variables in any time-series forecasting exercise, including MIDAS models. The inclusion of too many variables or those with low predictive power can adversely affect forecast accuracy (Ghysels et al. 2018). A consensus appears to be forming in the literature on this issue. For instance, Marcellino and Schumacher (2010) demonstrate that survey data and industrial production are vital for forecasting German GDP. Similarly, Kuzin et al. (2009) assert that financial indicators and the industrial production index are reliable indicators for forecasting GDP for the Euro Area. Focusing on Brazil, Zuanazzi and Ziegelmann (2014) specify industrial production, financial market data and exports as effective predictors of real GDP among 16 potential indicators when employing MIDAS and U-MIDAS methods. This research incorporates industrial production, exports, and financial market data to forecast GDP growth, which aligns with established and proven practices in the literature.

Nowcasting the Portuguese GDP growth serves as the final focus, demonstrating the practical application of high-frequency data in real-time economic analysis. The literature on nowcasting has expanded significantly, reflecting the growing interest in methods that improve forecast timeliness and accuracy. Giannone et al. (2008) emphasised the value of nowcasting in providing more immediate GDP insights, while Bańbura et al. (2011) showed how different mixed-frequency models could efficiently incorporate monthly data for short-term forecasts.

The remaining sections are organised as follows: the data is discussed in section 3; the econometric methodology is described in section 4 followed by the empirical results and discussion in section 5 and conclusion in section 6.

3. DATA

The datasets for the U-MIDAS and VAR models are structured to meet their specific needs. Both include quarterly real GDP in millions and quarterly GDP growth derived from that real value. The U-MIDAS model uses monthly data for the unemployment rate (%), exports of goods and services (in millions of euros), industrial production (volume), and the PSI-20 stock index closing prices (in euros). In contrast, the VAR model uses the same variables but on a quarterly basis. The quarterly data is sourced directly from its original form rather than calculated from the monthly data to avoid losing information.

The monthly and quarterly data were sourced from Eurostat, *Instituto Nacional de Estatística*, and *Banco de Portugal*. Detailed information about the data can be found in section A of the Appendix.

Both datasets cover the 2000Q1 to 2024Q1 period, with no missing values. This timeline includes major economic events in Portugal and Europe, such as the Global Financial Crisis, the Sovereign Debt Crisis, the COVID-19 pandemic, and the inflation peak caused by the Russian invasion of Ukraine. These events enhance the model's ability to make accurate predictions for the turbulent out-of-sample period from 2021Q1 to 2024Q1.

Four essential data series are used: industrial production, unemployment rate, PSI-20 index, and exports of goods and services. The industrial production series reflects mining, quarrying, manufacturing, electricity, gas, steam, and air conditioning supply activity; it is measured in production volume and is seasonally and calendar-adjusted, with an index base of 100 in 2021. The unemployment rate captures the active Portuguese population aged 16 to 74 years who are not employed, while the PSI-20 index represents the closing values of the stock market on the last trading day of each month (or the last day of the quarter in the VAR dataset). The exports of goods and services show the credit of the goods and services balance in millions of euros. Since this data was only available monthly, the quarter aggregation was done by summing up the corresponding months.

These four series were selected from a larger dataset that initially included four additional variables: the Harmonised Index of Consumer Prices (HICP), the Real Effective Exchange Rate (REER18), the M2 aggregate, and the consumer confidence index. However, only the four chosen variables – industrial production, unemployment rate, PSI-20, and exports of goods and services – provided the highest forecast accuracy. Various data combinations were tested to identify which variables could improve forecast accuracy. Adding other variable or excluding any of the main four variables led to a decline in accuracy. These results align with findings in the literature, which emphasise the strong predictive power of industrial production and exports for GDP forecasts and the importance of financial indicators, represented by the Portuguese stock index PSI-20.

To ensure the data's stationarity, identical transformations were applied in both the VAR and U-MIDAS datasets. The first difference was taken in monthly and quarterly formats for GDP in millions of euros and the unemployment rate. The PSI-20 index was converted into a monthly (or quarterly) rate of change, and for industrial production and the exports of goods and services, the log transformation and first difference were applied to approximate a growth rate. Further details on these transformations are provided in section B of the Appendix.

Constructing the VAR dataset was simple because all variables were available at the same frequency (quarterly). In contrast, additional steps were necessary to handle the mixed frequencies for the MIDAS dataset. For each variable, three monthly columns were created to represent the data. For example, in the 2000Q1 row, the PSI-20 variable was represented by three columns labelled "Month 1" (M1), "Month 2" (M2) and "Month 3" (M3) each capturing the data from the respective months within that quarter.

4. ECONOMETRIC METHODOLOGY

This study employs two distinct models: the Mixed Data Sampling approach, specifically its Unrestricted variant – U-MIDAS, and the Vector Autoregressive model as a benchmark for comparison. The following sections will provide a detailed specification of each model.

4.1 Restricted Mixed Data Sampling (MIDAS)

The original MIDAS regression is the baseline for part of this work, and it was presented by Ghysels et al. (2004) and developed by Ghysels et al. (2007). We want to forecast the Portuguese GDP quarterly growth, denoted by y_t. This variable is available only once between $t-1$ and t. We use monthly variables $x_t^{(m)}$, observed m times in the same period to realise the forecast. The MIDAS model, suggested by Ghysels, Santa-Clara and Valkanov (2004), is specified as:

$$y_{t+h} = \beta_0 + \beta_1 B\left(L^{\frac{1}{m}};\theta\right) x_t^{(m)} + \varepsilon_{t+h} \tag{1}$$

where $B\left(L^{\frac{1}{m}};\theta\right) = \sum_{k=0}^{K} B(k;\theta) L^{\frac{k}{m}}$, $L^{\frac{1}{m}}$ represents a lag operator. The lag coefficient in $B(k;\theta)$, corresponding to the lag operator $L^{\frac{k}{m}}$, is parametrised as a function of a small-dimensional vector of hyperparameters $\theta.$ Here, $B(k;\theta)$ acts as a weighting mechanism for aggregation, normalised so that its components sum to 1, and ε_{t+h} is a standard independent and identically distributed error term (Duarte et al. 2016). When m is large, such as when x is observed daily and y quarterly, equation 1 may involve a significant number of parameters. If the parameters are left unrestricted, it could result in significant parameter proliferation. For this reason, MIDAS is prepared with a parametrisation of the lagged coefficients of $B(k;\theta)$.

The exponential Almon and Beta lag functions (Ghysels et al. 2007) are the most commonly used polynomial functional forms. Because of their nonlinear specifications, MIDAS regressions for both forms must be estimated using nonlinear least squares.

4.2 Unrestricted Mixed Data Sampling (U-MIDAS)

The Restricted MIDAS model, discussed in the preceding section, is efficient when the frequency difference between datasets is significant, as it uses distributed lag functions to prevent parameter proliferation. The Unrestricted MIDAS variant introduced by Foroni et al. (2015) is used because it is particularly beneficial when there is a moderate discrepancy in the sampling frequencies between variables. In this model, Portuguese GDP quarterly growth is forecasted using monthly data. A key benefit of U-MIDAS is that it can be estimated using OLS, simplifying the computation process while maintaining forecasting accuracy. The U-MIDAS regression is expressed as:

$$y_{t+h} = \beta_0 + \emptyset y_t + \sum_{k=0}^{K} \beta_{k+1} L^{\frac{k}{m}} x_t^{(m)} + u_{t+h}$$

$$= \beta_0 + \emptyset y_t + \beta_1 x_t^{(m)} + \beta_2 x_{t-\frac{1}{m}}^{(m)} + \cdots + \beta_{K+1} x_{t-\frac{k}{m}}^{(m)} + u_{t+h} \qquad (2)$$

The low-frequency variable y_t, represented here by Portuguese GDP quarterly growth, is modelled as a function of its own lag and of the lags of $x_{j,t}$, the *j* monthly variable, at time t.

4.3 Vector Autoregressive (VAR)

The Vector Autoregressive model is a crucial tool in time series analysis that captures the dynamic relationships between multiple variables. It analyses a vector of time series variables, represented as Y_t, which are influenced by their past values and the past values of other variables in the system over a specified number of lag periods, denoted by p. The VAR model of order p, denoted as VAR (p), can be expressed as:

$$Y_t = A_1 Y_{t-1} + A_2 Y_{t-2} + \cdots + A_p Y_{t-p} + u_t. \quad (3)$$

In equation 3, Y_t is a vector $(n \times 1)$ of endogenous variables at time t. The matrices A_1, A_2,…, A_p represent the coefficients for each lag $i, i = 1, \dots, .p$, capturing the dynamic relationships between the variables and their past values as well as the past values of the other variables. The error term u_t represents shocks to the system and is assumed to be white noise, characterised by a zero mean and constant variance.

Forecasting with a VAR model is straightforward, where the one-step ahead forecast corresponds to the conditional expected value of Y_{t+1}, which in notation is:

$$E\left[\mathrm{Y}_{\mathrm{t+1}} \middle| Y_{t,} Y_{t-1,} \dots\right] = A_0 + A_1 Y_t + \cdots + A_p Y_{t-p+1} \quad (4)$$

The two-step-ahead forecast uses the same process, substituting $t + 1$ into the equation. This iterative method is then extended to the four-step ahead forecast.

5. FORECASTING PORTUGUESE GDP QUARTERLY GROWTH: RESULTS AND DISCUSSION

This section is organised by forecast horizons and details the forecasting outcomes for each model across the three studied periods. We present and analyse the results for each horizon, highlighting key forecast statistics such as Root Mean Square Error and Mean Absolute Error. To ensure robustness in the conclusions, we apply the Diebold-Mariano test (more details in section C of the Appendix) and the Clark-West test (details in section D of the Appendix) to assess whether the performance differences among the models are statistically significant. In addition, section 5 explores the nowcasting results for GDP quarterly growth using monthly variables, focusing on the conclusions and practical advantages of this approach.

The forecasting abilities of the U-MIDAS and VAR models are explored over an out-of-sample period from 2021Q1 to 2024Q1, focusing on three forecast horizons. The 1-step ahead forecasts project GDP growth for the next quarter using all data available up to the previous quarter. For example, the forecast for the first quarter of 2021 (2021Q1) is based on data from 2000Q1 to 2020Q4. The forecast for the second quarter of 2021 (2021Q2) uses information up to 2021Q1. By only including actual data from the previous quarters and not relying on earlier forecasts, the rolling window method reduces the risk of carrying over errors. This approach results in a total of 13 one-quarter ahead predictions.

The 2-step ahead forecasts assess how effectively each model predicts GDP growth two quarters into the future, utilising the rolling window approach. This method generates 12 forecasts, starting with 2021Q2 as the initial point in the out-of-sample period. For each forecast, such as for 2022Q2, data available up to two quarters prior (in this case, through 2021Q4) is used, consistently maintaining a two-quarter gap between the forecast period and the real data utilised.

Lastly, the 4-step ahead forecasts aim to assess the accuracy of each model over a longer time frame of one year. This method generates 10 forecasts, starting from 2021Q4. For these long-term forecasts, data available through 2020Q4 is used to project GDP growth for 2021Q4, with the same process applied to the subsequent quarters. This setup allows for a comprehensive evaluation of how well each model performs across different forecasting horizons, providing valuable insights into their effectiveness in capturing GDP growth over different time frames.

For forecasting and running the models in Stata, different lags were used in all the variables based on the number of steps: one lag for 1-step ahead forecasts, two for 2-step ahead forecasts, and four for 4-step ahead forecasts. Overall, the forecasts show solid predictive power, especially in the 4-step ahead and in the later quarters for the 1-step and 2-step predictions. It is important to note that the out-of-sample period was particularly volatile due to economic disruptions caused by COVID-19 pandemic and the inflationary pressures triggered by the Russian invasion of Ukraine.

5.1 One-Step Ahead Forecast

Table 1 presents the out-of-sample forecasting results for Portuguese GDP growth at the 1-step ahead horizon for both models and Figure 1 provides a clearer insight and visualisation.

***Table 1**: 1-Step Ahead Forecast Results*

Out-of-Sample Period	*Real Values*	*U-MIDAS Results*	*VAR Results*
2021Q1	-2.819	-2.246	2.143
2021Q2	4.576	0.096	0.268
2021Q3	3.593	0.809	0.166
2021Q4	1.864	4.235	1.338
2022Q1	1.665	0.168	-0.093
2022Q2	0.367	2.538	1.080
2022Q3	1.012	-0.534	-0.081
2022Q4	0.488	0.941	-0.201
2023Q1	1.369	0.692	0.128
2023Q2	0.041	-0.881	0.405
2023Q3	-0.122	0.925	0.951
2023Q4	0.833	0.583	0.460
2024Q1	0.643	0.381	0.076

***Figure 1**: 1-Step Out-of-Sample Results*

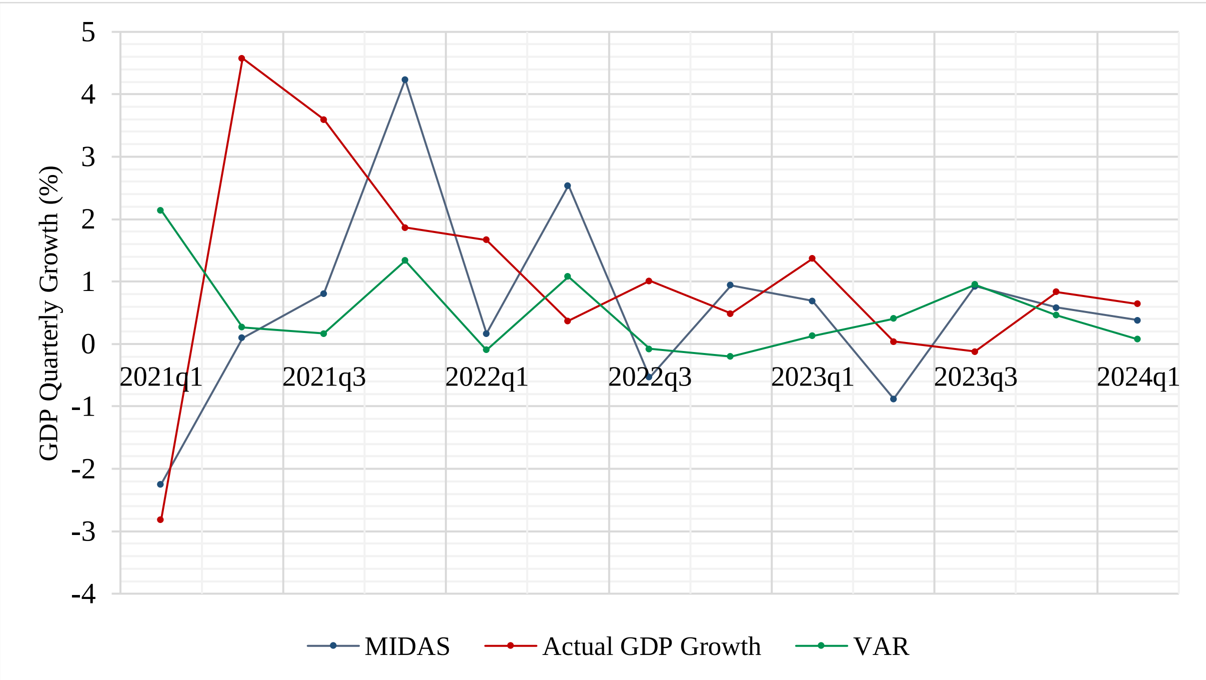

Both models demonstrate limited predictive power during the first quarters. The U-MIDAS model effectively captures the abrupt downturn in 2021Q1 but begins to show improved accuracy only starting from 2022Q4, closely mirroring actual values and performing well in the final two quarters of the out-of-sample period. Conversely, the VAR model exhibits overall stability; although its accuracy is also weaker in the early quarters, it starts to align more closely with actual values by 2022Q2. This initial lack of precision in both models likely reflects the unusual economic fluctuations at the beginning of the out-of-sample period, which were marked by the COVID-19 recovery phase and sudden inflationary pressures. These rapid and unpredictable changes made it difficult for the models to capture short-term dynamics. However, as conditions began to stabilise, both models adapted more effectively, with U-MIDAS showcasing particular strength in capturing monthly variations, thereby contributing to its enhanced performance in the later quarters. The U-MIDAS model has a lower RMSE and MAE than the VAR model, as shown in Table 2.

Table 2*: 1-Step Out-of-Sample Comparison*

	RMSE	**MAE**
U-MIDAS	1,88	0,45
VAR	2,21	0,53

These results indicate that U-MIDAS have better forecast accuracy. To confirm this difference statistically, we conducted the Diebold-Mariano test and the Clark-West test. Both tests compare the precision between models (Table 3). The tests' null hypothesis assumes that both forecasts have the same level of accuracy, while the alternative hypothesis suggests that one model (U-MIDAS) is significantly more accurate. We fail to reject the null hypothesis for the 1-step ahead forecasts for Diebold-Mariano. In contrast, the Clark-West test conclude that U-MIDAS is statistically superior to VAR in terms of forecast accuracy.

Table 3: Out-of-Sample Tests for 1-Step Forecasts

Test	p-value	H1	H0	Model 1	Model 2
Diebold-Mariano	0,2413	Model 1 is more accurate than Model 2	Both forecasts have the same accuracy	1-Step Ahead U-MIDAS	1-Step Ahead VAR
Clark-West	0,0006	The alternative model is more accurate	Both forecasts have the same accuracy	1-Step Ahead U-MIDAS (Alt)[1]	1-Step Ahead VAR

5.2 Two-Step Ahead Forecast

A weaker performance in the two-step ahead forecast were demonstrated by both U-MIDAS and VAR. This forecasting horizon challenges the models' ability to account for mid-term fluctuations. While the high-frequency data from U-MIDAS may introduce noise rather than clarity, the low-frequency data from VAR, although more stable, may lack responsiveness. Table 4 presents the out-of-sample results for each model, and Figure 2 compares their predictions, providing a closer look at their relative performance.

Table 4: 2-Step Ahead Forecast Results

Out-of-Sample Period	*Real Values*	*U-MIDAS Results*	*VAR Results*
2021Q2	4.576	2.373	0.003
2021Q3	3.593	2.484	1.290
2021Q4	1.864	-2.695	-1.578
2022Q1	1.665	-3.19	-1.359
2022Q2	0.367	-0.588	-0.161
2022Q3	1.012	-0.429	-0.158
2022Q4	0.488	0.926	0.653
2023Q1	1.369	-0.705	-0.032
2023Q2	0.041	-1.532	-0.535
2023Q3	-0.122	1.082	-0.516
2023Q4	0.833	0.465	0.747
2024Q1	0.643	0.423	0.252

[1] Alt denotes the alternative model and is used in the Clark-West test to specify which model is being evaluated as potentially superior.

***Figure 2**: 2-Step Out-of-Sample Results*

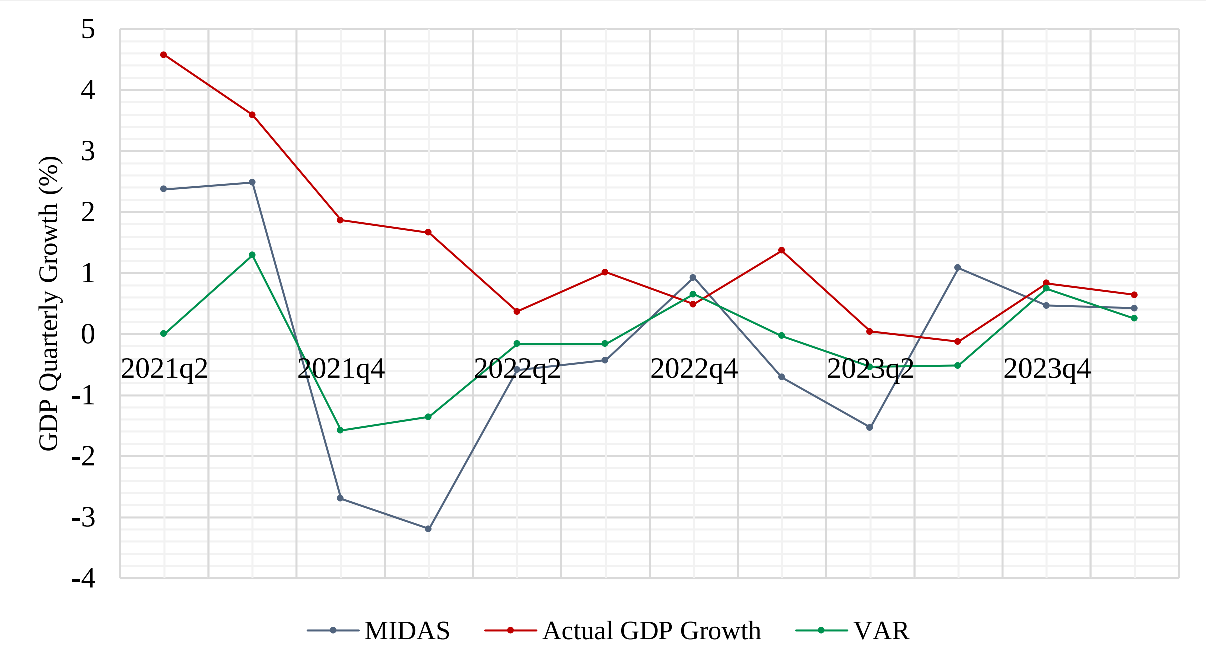

In the 2-step ahead forecasts, the VAR model demonstrates greater predictive power, aligning closely with actual trends from the first quarter of 2022 onward and providing reliable estimates through the first quarter of 2024. This consistency is likely due to the stability of the VAR. In contrast, the U-MIDAS model benefits from using monthly data in its initial forecasts but experiences a decline in accuracy between the fourth quarter of 2021 and the second quarter of 2022, only regaining its predictive strength in the final quarters of the out-of-sample period.

In this case, the VAR model demonstrates lower RMSE compared to U-MIDAS, indicating better overall accuracy, as can be observed in Table 5. The Diebold-Mariano test confirms that this difference in accuracy is statistically significant. Notably, the test's p-value allows us to weakly reject the null hypothesis for equal forecasting accuracy, suggesting that VAR forecasts are statistically more accurate than the U-MIDAS model in this instance. On the other hand, the Clark-West test does not allow the rejection of the null hypothesis, meaning that both forecasts have equivalent accuracy.

***Table 5**: 2-Step Out-of-Sample Comparison*

	RMSE	MAE
U-MIDAS	2,27	1,48
VAR	2,07	1,48

***Table 6**: Out-of-Sample Tests for 2-Step Forecasts*

Test	p-value	H1	H0	Model 1	Model 2
Diebold-Mariano	0,0948	Model 1 is more accurate than Model 2	Both forecasts have the same accuracy	2-Steps Ahead U-MIDAS	2-Steps Ahead VAR
Clark-West	0,4825	The alternative model is more accurate	Both forecasts have the same accuracy	2-Steps Ahead U-MIDAS	2-Step Ahead VAR (Alt)

5.3 Four-Step Ahead Forecast

The 4-step ahead forecast evaluates the ability of both models to predict the GDP growth rate over the long term by requiring estimates for quarters a full year in advance. This forecast horizon spans from the fourth quarter of 2021 (2021Q4) to the first quarter of 2024 (2024Q1). It excludes the highly volatile periods of the first three quarters of 2021 (2021Q1, 2021Q2, and 2021Q3), during which GDP growth experienced significant fluctuations due to the post-COVID-19 recovery. By omitting these turbulent quarters, the forecast period becomes more balanced, offering a clearer perspective on each model's performance. The out-of-sample results for this period are presented in the table below, while Figure 3 provides a side-by-side comparison of the forecasts from each model.

***Table 7**: 4-Step Ahead Forecast Results*

Out-of-Sample Period	*Real Values*	*U-MIDAS Results*	*VAR Results*
2021Q4	1.864	1.163	0.354
2022Q1	1.665	0.09	-0.698
2022Q2	0.367	-0.037	0.352
2022Q3	1.012	1.609	0.706
2022Q4	0.488	-0.866	0.173
2023Q1	1.369	0.903	0.499
2023Q2	0.041	-0.012	-0.185
2023Q3	-0.122	0.399	0.137
2023Q4	0.833	0.167	0.476
2024Q1	0.643	0.591	0.482

***Figure 3**: 4-Step Out-of-Sample Results*

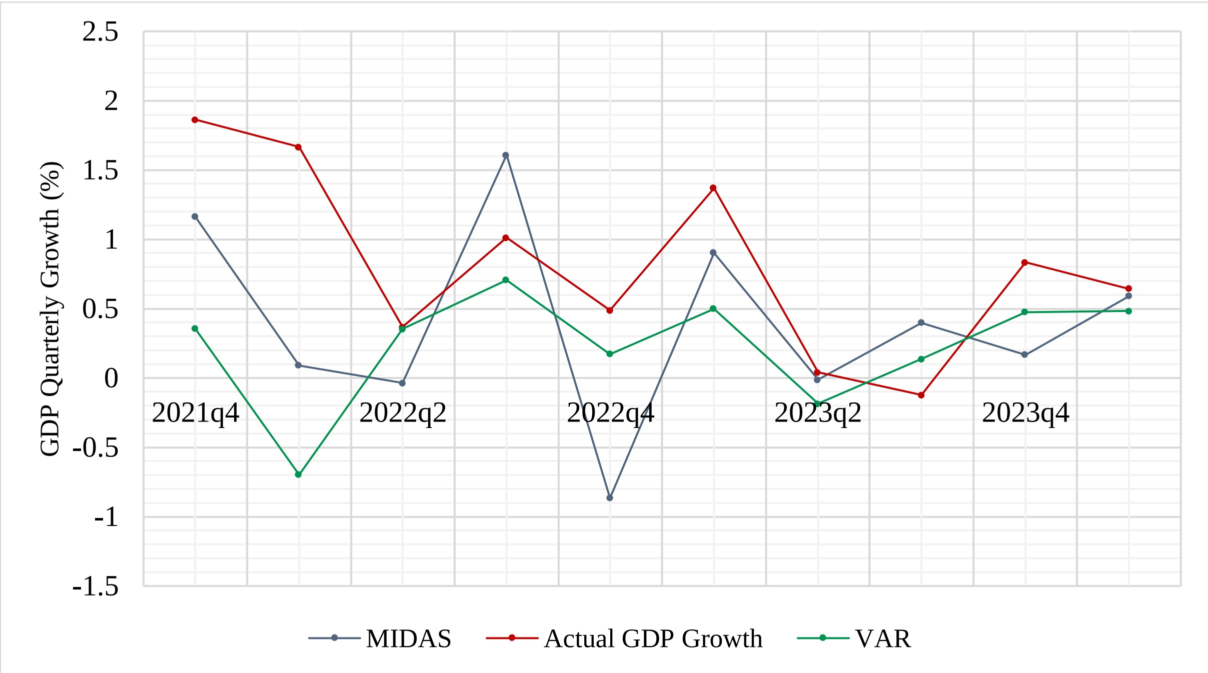

Both models exhibit strong predictive capabilities within this context, effectively tracking actual GDP values. The U-MIDAS model closely mirrors the natural trend, accurately following GDP trajectories, although it occasionally overstates the magnitude of changes. This level of detail in data helps the model respond to subtle trends or turning points that quarterly data may miss, improving the responsiveness and accuracy of long-term forecasts. Conversely, the VAR model slightly underestimates actual values but demonstrates a consistent methodology, making it a reliable option for longer-term forecasting. This balanced out-of-sample evaluation indicates that both models perform significantly better than for other time horizons, underscoring their effectiveness in long-term forecasting, particularly in a more stable economic environment.

Statistical results, shown below, include the RMSE, MAE, Diebold-Mariano and the Clark-West test for both models. The U-MIDAS model has lower RMSE and MAE values compared to the VAR model, indicating greater forecast accuracy. This difference is not statistically significant, as confirmed by the Diebold-Mariano test, which yields a p-value of 0.33, and by the Clark-West test, with a p-value of 0.52. These results do not allow us to reject the null hypothesis that there is no significant difference between the two forecasts, meaning we cannot conclude a statistically significant advantage for U-MIDAS.

Table 8*: 4-Step Out-of-Sample Comparison*

	RMSE	**MAE**
U-MIDAS	0,79	0,42
VAR	0,95	0,59

Table 9*: Out-of-Sample Tests for 4-Step Forecasts*

Test	**p-value**	**H1**	**H0**	**Model 1**	**Model 2**
Diebold-Mariano	0,3373	Model 1 is more accurate than Model 2	Both forecasts have the same accuracy	4-Steps Ahead U-MIDAS	4-Steps Ahead VAR
Clark-West	0,5269	The alternative model is more accurate	Both forecasts have the same accuracy	4-Steps Ahead U-MIDAS (Alt)	4-Step Ahead VAR

5.4 Nowcasting the Portuguese GDP Quarerly Growth using U-MIDAS

A nowcasting analysis of quarterly GDP growth was conducted to leverage the benefits of high-frequency data. This approach uses U-MIDAS and allows us to estimate GDP growth for the upcoming quarter every month, enabling to observe how forecasts develop and determine whether accurate predictions of this low-frequency variable can be made several months in advance. The primary objective is to assess whether incorporating monthly data improves the accuracy of GDP quarterly growth forecasts.

For example, the GDP growth for the first quarter of 2021 (2021Q1) is estimated three times: in January, February, and with the final value calculated in March, in alignment with the U-MIDAS 1-step model estimation discussed in section 5.1.

The findings highlight several advantages of utilising high-frequency data. However, including data from the second and third months of the quarter does not always improve forecast accuracy. Interestingly, estimates based solely on data from the first month of each quarter are closest to the actual quarterly value 46% of the time, compared with 1-step ahead U-MIDAS. This suggests that data from the first month alone provides a strong predictive signal for quarterly GDP growth, effectively capturing key economic dynamics early on. Consequently, additional monthly updates may have a limited impact or could even introduce unnecessary noise.

These results suggest that the economic indicators available in the first month serve as reliable leading indicators for the Portuguese economy, closely aligning with quarterly GDP trends and establishing a solid foundation for nowcasting without further adjustments. From an economic perspective, these findings are significant, as accurately predicting a quarter's GDP in the first month enables more effective economic planning and decision-making.

Table 10 compares quarterly forecasts based on first-month data (for Q1 - January; Q2 - April; Q3 - July; Q4 - October) with estimates derived from data across all three months of the quarter, similar to the U-MIDAS 1-step ahead approach. The comparison highlights the advantages of early high-frequency indicators in enhancing forecasting accuracy.

***Table 10**: Nowcasting and Forecast Results*

Out-of-Sample Period	*Real Values*	*Nowcasting Using M1 Data Results*	*1-Step Ahead U-MIDAS Results*
2021Q1	-2.819	-2,587	-2,246
2021Q2	4.576	-0,178	0,096
2021Q3	3.593	2,247	0,809
2021Q4	1.864	1,848	4,235
2022Q1	1.665	0,479	0,168
2022Q2	0.367	2,197	2,538
2022Q3	1.012	0,799	-0,534
2022Q4	0.488	0,012	0,941
2023Q1	1.369	0,299	0,692
2023Q2	0.041	0,496	-0,881
2023Q3	-0.122	2,286	0,925
2023Q4	0.833	-0,610	0,583
2024Q1	0.643	0,467	0,381

***Figure 4**: Nowcasting Out-of-Sample Results*

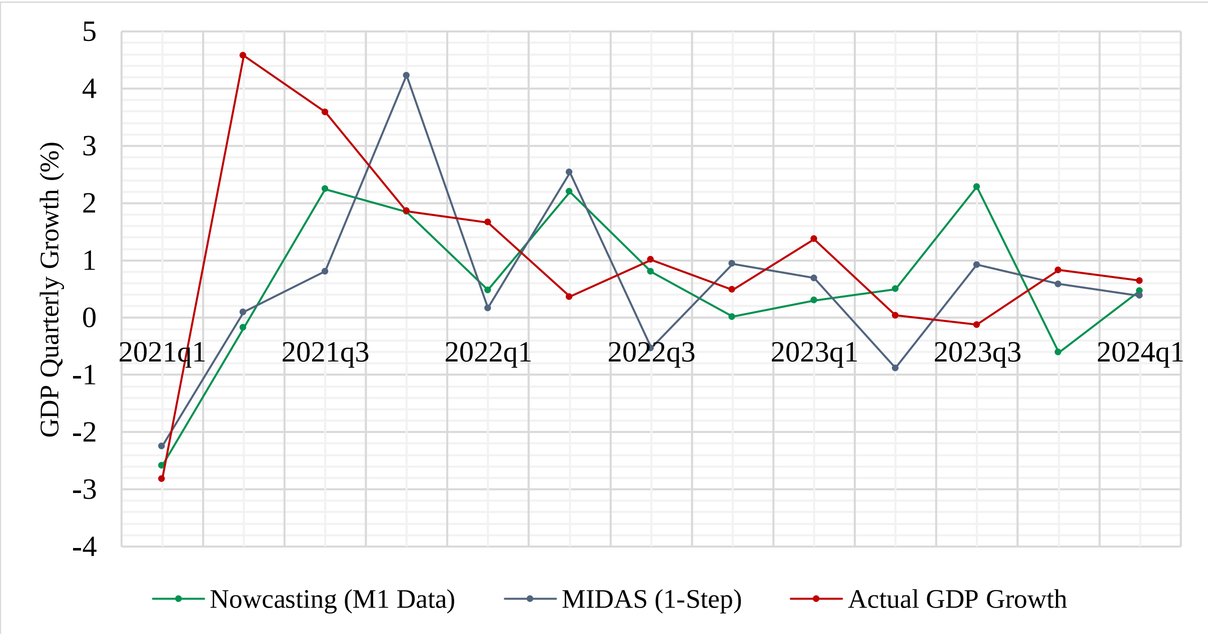

Utilising only M1 data leads to enhanced forecast accuracy, as the model suggests that the initial month's data often captures the most relevant early economic signals for the quarter. Including M2 and M3 data might introduce noise, as subsequent data points can indicate transient fluctuations rather than stable trends. The discrepancy between forecasting for the complete quarter and relying solely on M1 data can be attributed to the data types used in this work project. The PSI-20 index, and exports typically respond quickly to economic conditions, reflecting market changes and international demand. Additionally, indicators such as the unemployment rate and industrial production tend to remain stable throughout the quarter, making the exclusive use of M1 data more effective in capturing this underlying trend.

***Table 11**: Nowcasting and U-MIDAS (1-Step Ahead) Out-of-Sample Comparison*

	RMSE	**MAE**
U-MIDAS (1-STEP AHEAD)	1,88	0,45
NOWCASTING (M1 DATA)	1,73	0,44

***Table 12**: Clark-West Test for Nowcasting*

CW Statistic	**H1**	**H0**	**Model 1**	**Model 2**
0,0003	Alternative model is more accurate	Both models have the same forecast accuracy	Nowcasting Using only M1 Data (Alt)	1-Step Ahead U-MIDAS

Regarding accuracy, the RMSE and MAE are lower when using only the first month's data than the entire quarter. The Diebold-Mariano test was not conducted in this case because the two models are nested, meaning the error differences are not independent across observations. Instead, the Clark-West test was performed and yielded a p-value close to zero, so we reject the null hypothesis, meaning that the alternative model (Nowcasting) is more accurate in its predictions than the 1-Step ahead U-MIDAS.

The results presented in section 5 are consistent with existing literature. Zuanazzi and Ziegelmann (2014) found that MIDAS and U-MIDAS outperformed low-frequency benchmarks, such as ARMA, in 1-step-ahead GDP forecasts for Brazil. Andreou et al. (2010) demonstrated that MIDAS models enhance real GDP growth forecasts for the US compared to the Autoregressive (AR) model, Factor Autoregressive (FAR) model, and the mean and median forecasts from the Survey of Professional Forecasts (SPF), showing consistent advantages across all horizons. Clements and Galvão (2008) identified significant advantages of MIDAS over autoregressive distributed lag models for short-term forecasts, with diminishing benefits for medium-term horizons. Although the performance improvement of MIDAS may not always align with the levels reported here, its ability to utilise monthly indicators is associated with gains over low-frequency models in forecasting.

Two principal limitations warrant consideration in the context of this analysis. Firstly, the out-of-sample period from 2021Q1 to 2024Q1 is characterised by significant disruptions, including the COVID-19 pandemic and the Russian invasion of Ukraine. These events have introduced substantial complexity into the forecasting of GDP growth during this interval, particularly in the earlier stages, thereby diminishing the model's potential to achieve a high level of accuracy. Secondly, while the selection of leading indicators was rigorous, it remains plausible that some relevant factors influencing GDP growth, such as survey data and retail and consumer behaviour data, may have been excluded. This exclusion can slightly constrain the model's ability to capture the full spectrum of economic dynamics in Portugal.

6. Conclusions

The delay in publishing GDP data in Portugal, compared to monthly economic activity data, has practical implications for forecasting this essential variable, which is crucial for policymakers to anticipate trends and monitor the country's economic health. This delay makes using methods to integrate data of different frequencies necessary to improve GDP forecasts and provide monthly updates (nowcasting). A benchmark model, VAR, relying solely on low-frequency data, was used to demonstrate the benefits of high-frequency data.

Regarding forecasting performance, the U-MIDAS model showed higher accuracy in the 1-step- and 4-step ahead forecasts, as judged by the RMSE and MAE metrics. While U-MIDAS performed better, only the Clark-West test showed a statistically significant advantage for the 1-step ahead. For the 4-step ahead, both the Diebold-Mariano and Clark-West tests indicated no statistically significant differences between U-MIDAS and VAR. On the other hand, in the 2-step ahead forecast, VAR outperformed U-MIDAS, and after testing the significance of this difference, the Diebold-Mariano test confirmed that VAR's results were statistically superior while the Clark-West test concludes that they have an equivalent performance. Similarly, the nowcasting model, using data from only the first month of each quarter, demonstrated better forecast accuracy than the 1-step U-MIDAS. In this case, the difference was statistically significant according to the Clark-West test and we can conclude that the nowcasting performance is more accurate than the 1-step U-MIDAS. Overall, U-MIDAS showed better performance in most scenarios as proved in the literature, and visually, both models appear to track the actual out-of-sample trend closely. When comparing the predictive power of using just the first month's data versus using all three months to forecast quarterly GDP, it became clear that this approach has a positive impact, likely tied to the data type used to predict Portugal's GDP.

Future research could explore the use of higher-frequency data, such as daily and weekly observations, to improve nowcasting models for Portuguese GDP growth. By estimating the GDP growth rate for the next quarter daily, it would be possible to monitor economic activity in real time and detect early signs of potential recessions.

References

Andreou, Elena, Eric Ghysels, and Andros Kourtellos. 2010. "Should macroeconomic forecasters use daily financial data and how?" Working Paper Series 42_10. Rimini Centre for Economic Analysis. https://ideas.repec.org/p/rim/rimwps/42_10.html.

Banerjee, Anindya, Massimiliano Marcellino, and Igor Masten. 2003. "Leading indicators for Euro area inflation and GDP growth." Working Paper No. 235. IGIER (Innocenzo Gasparini Institute for Economic Research). https://doi.org/10.2139/ssrn.396359.

Bańbura, Marta, Domenico Giannone, Michele Modugno, and Lucrezia Reichlin. 2013. "Now-casting and the real-time data flow." Working Paper Series No. 1564. European Central Bank. https://doi.org/10.2139/ssrn.2284274.

Banerjee, Anindya, Massimiliano Marcellino, and Igor Masten. 2005. "Forecasting Macroeconomic variables for the new member states of the European Union." Working Paper Series No. 482. European Central Bank. http://dx.doi.org/10.2139/ssrn.711161.

Bhaghoe, Sailesh, Gavin Ooft, and Philip Hans Franses. 2019. "Estimates of Quarterly GDP Growth using MIDAS regressions." Econometric Institute Research Papers EI2019-29. Erasmus University Rotterdam, Erasmus School of Economics, Econometric Institute. http://hdl.handle.net/1765/118667.

Clements, Michael P., and Ana Beatriz Galvão. 2009. "Forecasting US output growth using leading indicators: An appraisal using midas models." *Journal of Applied Econometrics* 24 (7): 1187-1206. https://doi.org/10.1002/jae.1075.

Cottrell, Allin, and Jack Lucchetti. 2017. "MIDAS in gretl."

Dias, Francisco, Maximiano Pinheiro, and António Rua. 2015. "Forecasting Portuguese GDP with factor models: Pre- and post-crisis evidence." *Economic Modelling* 44: 266-272. https://doi.org/10.1016/j.econmod.2014.10.034.

Duarte, Cláudia, Paulo M.M. Rodrigues, and António Rua. 2016. “A mixed frequency approach to forecast private consumption with ATM/POS data.” *International Journal of Forecasting* 33 (1): 61-75. https://doi.org/10.1016/j.ijforecast.2016.08.003.

Ferreira, Julia. 2020. “Transmuting Unequally Spaced Data: A MIDAS Regression Touch to Forecast Real GDP Growth in Brazil.” Master’s Thesis. Fundação Getulio Vargas São Paulo School of Economics. https://hdl.handle.net/10438/29972.

Foroni, Claudia, and Massimiliano Marcellino. 2013. “A survey of econometric methods for mixed-frequency data.” Working Paper No. 2013|06. Norges Bank Research. https://doi.org/10.2139/ssrn.2268912.

Foroni, Claudia, Massimiliano Marcellino, and Christian Schumacher. 2015. “Unrestricted mixed data sampling (MIDAS): MIDAS regressions with unrestricted lag polynomials”. *Journal of the Royal Statistical Society Series A: Statistics in Society* 178 (1): 57-82. https://doi.org/10.1111/rssa.12043.

Ghysels, Eric. 2016. “Macroeconomics and the Reality of Mixed Frequency Data.” *Journal of Econometrics* 193 (2): 294-314. https://doi.org/10.1016/j.jeconom.2016.04.008.

Ghysels, Eric, and Massimiliano Marcellino. 2018. *Applied Economic Forecasting using Time Series Methods*. 1st ed. New York: Oxford University Press.

Ghysels, Eric, and Hang Qian. 2019. “Estimating MIDAS Regressions via OLS with Polynomial Parameter Profiling.” *Econometrics and Statistics* 9 (C): 1-16. https://doi.org/10.1016/j.ecosta.2018.02.001.

Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov. 2004. “The MIDAS Touch: Mixed Data Sampling Regression Models.” University of California at Los Angeles, Anderson Graduate School of Management. https://ideas.repec.org/p/cdl/anderf/qt9mf223rs.html.

Ghysels, Eric, Arthur Sinko, and Rossen Valkanov. 2007. “MIDAS Regressions: Further Results and New Directions.” Econometric Reviews, Taylor & Francis Journals 26 (1): 53-90. https://doi.org/10.1080/07474930600972467.

Giannone, Domenico, Lucrezia Reichlin, and David Small. 2008. “Nowcasting: The real-time informational content of macroeconomic data.” *Journal of Monetary Economics 55* (4): 665-676. https://doi.org/10.1016/j.jmoneco.2008.05.010.

Jardet, Caroline, and Baptiste Meunier. 2020. “Nowcasting World GDP Growth with High-Frequency Data.” Working Paper No. 788. Banque de France. https://doi.org/10.2139/ssrn.3749139.

Knotek II, Edward, S., and Saeed Zaman. 2019. “Financial nowcasts and their usefulness in macroeconomic forecasting.” *International Journal of Forecasting* 35 (4): 1708-1724. https://doi.org/10.1016/j.ijforecast.2018.10.012.

Koenig, Evan, F., Sheila Dolmas, and Jeremy Piger. 2003. “The Use and Abuse of Real-Time Data in Economic Forecasting”. *The Review of Economics and Statistics* 85 (3): 618–628. https://doi.org/10.1162/003465303322369768.

Kuzin, Vladimir, Massimiliano Marcellino, and Christian Schumacher. 2009. “MIDAS VS. MIXED-FREQUENCY VAR: NOWCASTING GDP IN THE EURO AREA.” Discussion Paper No. 7445. Centre for Economic Policy Research (CEPR). https://doi.org/10.2139/ssrn.2785336.

Sdrakas, Nikolaos, and Cedric Viguie. 2003. “VAR modelling of the euro area GDP on the basis of principal component analysis.” Luxembourg: Office for Official Publications of the European Communities: 1-19.

Stock, James. H., and Mark W. Watson. 2001. “Vector Autoregressions.” *Journal of Economic perspectives* 15 (4): 101-115. https://doi.org/10.1257/jep.15.4.101.

Zuanazzi, Pedro, T., and Flávio A. Ziegelmann. 2014. “Previsões para o crescimento do PIB trimestral brasileiro com séries financeiras e econômicas mensais: Uma aplicação de MIDAS.” *Economia Aplicada* 18 (2): 295-318. https://doi.org/10.1590/1413-8050/ea515.

Appendix

A. Data Description

1. Real Gross Domestic Product (GDP) in millions of Euros

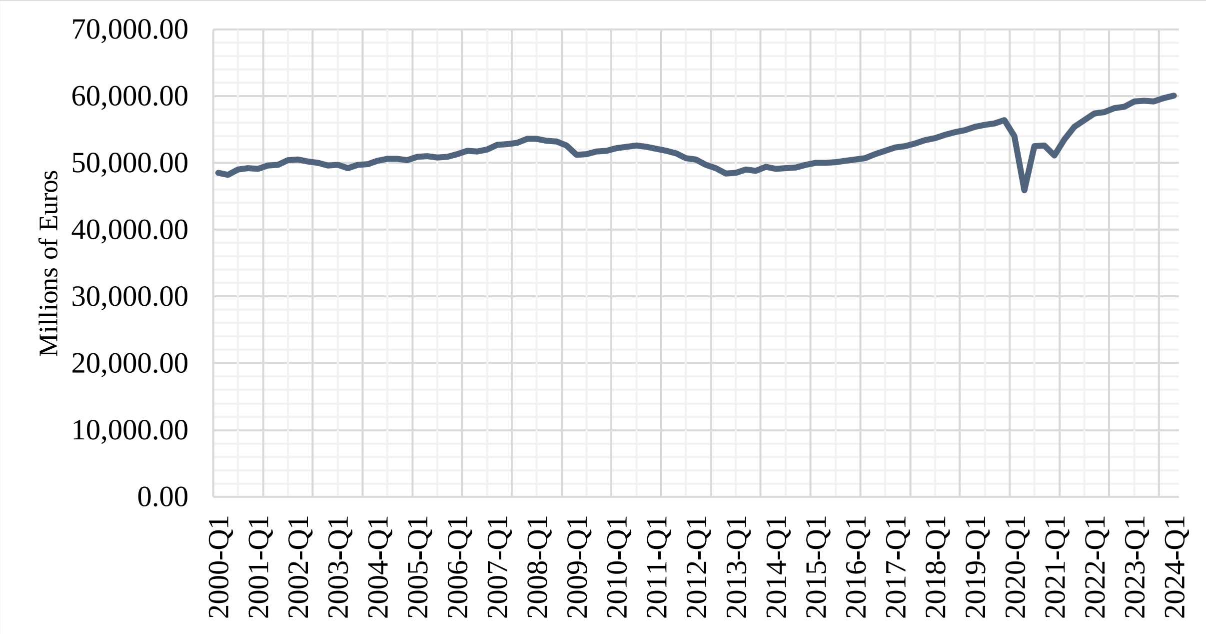

This series was retrieved from *Instituto Nacional de Estatística* (INE) and it is not calendar or seasonally adjusted.

2. Unemployment rate (active population between 16 and 74 years old) (%)

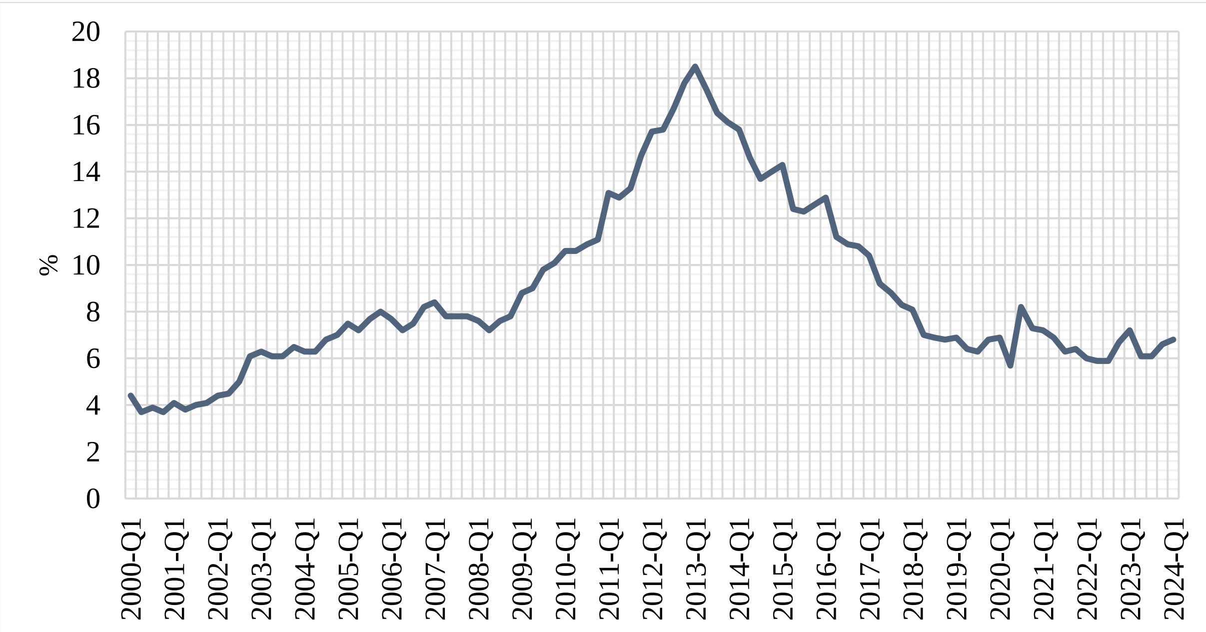

This series was retrieved from *Instituto Nacional de Estatística* (INE). The data from January 2000 to January 2011 were obtained using weights calibrated based on the Monthly Estimates of Resident Population, calculated explicitly for the Employment Survey according to the definitive results of the 2011 Census. The values from February 2011 onward were obtained using weights calibrated based on the Monthly Estimates of Resident Population, calculated explicitly for the Employment Survey according to the definitive results of the 2021 Census.

3. Exports of goods and services in millions of Euros

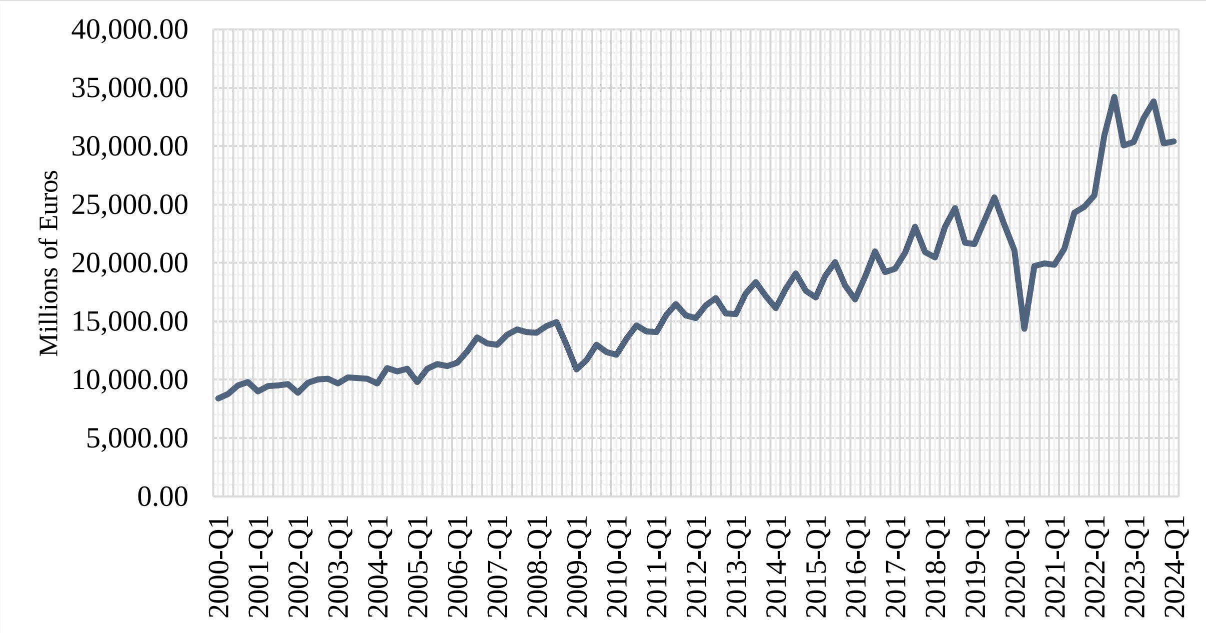

This series was retrieved from BPstat (Bank of Portugal Statistics) and it is not calendar or seasonally adjusted.

4. Industrial production (volume) (2021=100)

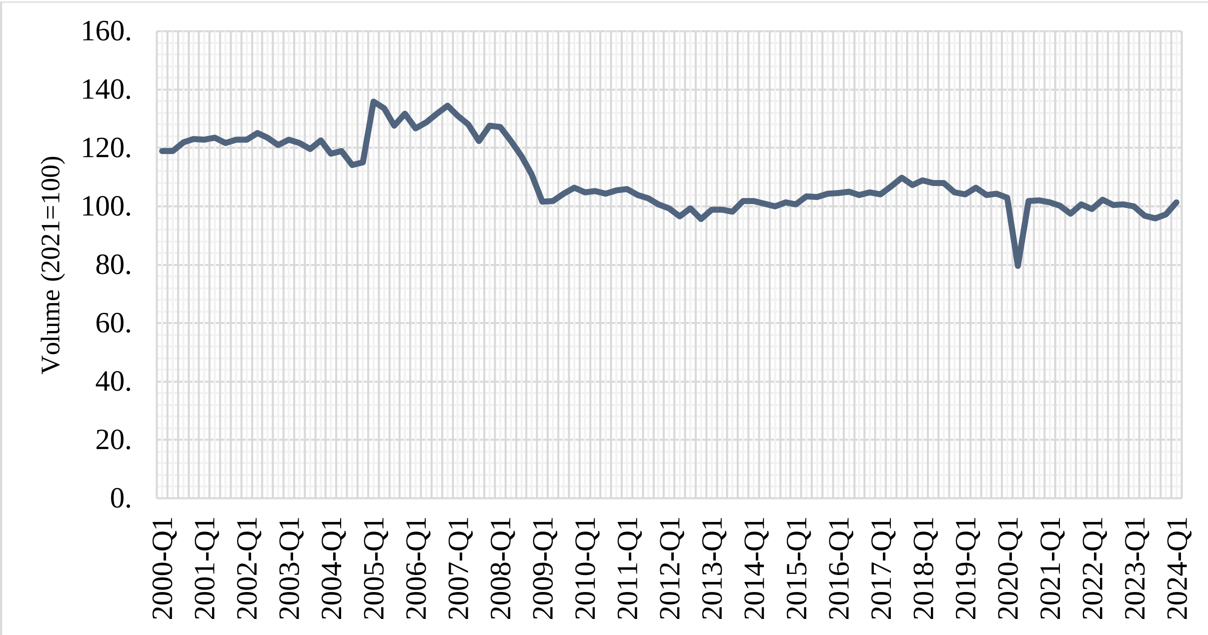

This series, retrieved from Eurostat, includes activities such as quarrying, manufacturing, and the supply of electricity, gas, steam, and air conditioning. It has been seasonally, and calendar adjusted.

5. PSI-20 stock index closing prices in Euros

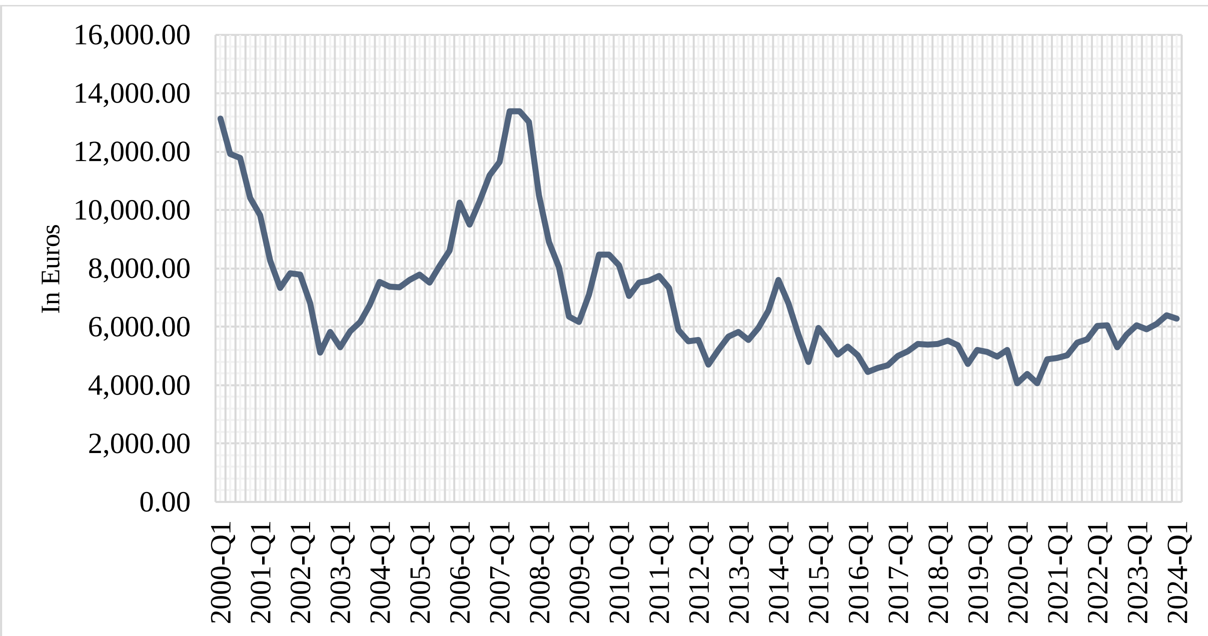

This series was retrieved from BPstat (Bank of Portugal Statistics) and represents the Portuguese Stock Index, comprising the largest companies listed on the Lisbon Stock Exchange.

B. Data Transformation

Various transformations were applied to ensure the series' stationarity. Trends were removed by differencing the variables, while logarithmic transformations stabilised variances. These adjustments were necessary to meet the fundamental assumptions required for robust time series analysis.

Variable	Transformation
Real GDP (Basis 2021)	First Difference: $Y_t - Y_{t-1}$
Unemployment Rate (%)	First Difference: $X_t - X_{t-1}$
Exports of Goods and Services	First Difference and Logarithm: $\log(X_t) - \log(X_{t-1})$
Industrial Production (2021=100)	First Difference and Logarithm: $\log(X_t) - \log(X_{t-1})$
PSI-20 Stock Index	Growth Rate: $\frac{X_t - X_{t-1}}{X_{t-1}}$

C. Diebold-Mariano Test

The Diebold-Mariano (DM) test is employed to evaluate the predictive accuracy of two forecasting models. This test is particularly valuable for assessing whether one model consistently demonstrates superior performance over the other in out-of-sample forecasts. The hypotheses for the test are:

- H_0: No difference in the accuracy of VAR and U-MIDAS.
- H_1: The predictive accuracy of the two models differ significantly, with one model performing better (we assume the one with lower RMSE).

The DM test can be defined as follows:

1. Let e_i and r_i be the residuals for the two forecasts:

$$e_i = y_i - f_i \qquad r_i = y_i - g_i$$

2. The test calculates d_i, referred to as the loss differential, by using the Mean Squared Error (MSE) statistic:

$$d_i = e_i^2 - r_i^2$$

3. Then, we define:

$$\bar{d} = \frac{1}{n}\sum_{i=1}^{n} d_i$$

4. For $n > k \geq 1$, we define:

$$\gamma_k = \frac{1}{n}\sum_{i=k+1}^{n} (d_i - \bar{d})(d_{i-k} - \bar{d})$$

γ_k is the autocovariance at lag k.

5. For $h \geq 1$, we define the Diebold-Mariano statistic as follows:

$$DM = \frac{\bar{d}}{\sqrt{\frac{[\gamma_0 + 2\sum_{k=1}^{h-1}\gamma_k]}{n}}}$$

D. Clark-West Test

The Clark and West (2007) test is an adjusted version of the Diebold and Mariano (1995) statistic, which also follows a standard normal distribution. The hypotheses for the test are:

- H_0: The alternative model does not improve forecast accuracy; both models perform equally well.
- H_1: The alternative model improves forecast accuracy compared to the simpler benchmark model.

The CW test can be defined as follows:

$$CW = \frac{\overline{\hat{s}}\sqrt{p}}{\sqrt{Avar(\hat{s})}}$$

$$\hat{s} = \left(y_{t+1} - \hat{y}_{m2,t+1}\right)^2 - \left(y_{t+1} - \hat{y}_{m1,t+1}\right)^2 + \left(\hat{y}_{m2,t+1} - \hat{y}_{m1,t+1}\right)^2$$

Where:

- $\hat{s}$: adjusted loss differential function.
- $\sqrt{Avar(\hat{s})}$: asymptotic variance of the adjusted loss differential function.