IBEX-FinTime: A Harmonized Spanish Macro-Financial Dataset (2005–2025) and Benchmark for Time Series Foundation Models

José Ignacio Meseguer-Gómez[1], Gonzalo Bengoa-Díaz[1], Pablo Reina-Jiménez[1], María Martínez-Ballesteros[1,2], and José C. Riquelme-Santos[1,2]

[1] Department of Computer Languages and Systems, Universidad de Sevilla
Av. Reina Mercedes S/N, 41012 Seville, Spain
jmeseguer@us.es, gonbendia2@alum.us.es, prjimenez@us.es, mariamartinez@us.es, riquelme@us.es
[2] SCORE Lab, I3US Institute, Universidad de Sevilla, Seville, Spain

Abstract. This paper presents IBEX-FinTime, a harmonized dataset integrating financial market variables and macroeconomic indicators related to the Spanish economy within an international context. The dataset spans from January 2005 to December 2025 and contains 7,670 daily observations across 86 variables, including stock market data, international indices, bond yields, exchange rates, commodity prices, and macroeconomic indicators for Spain, the United States, and the Euro Area. The dataset integrates nine main categories of variables within a single harmonized daily structure designed for multivariate financial analysis and forecasting applications. Building unified time-series datasets from heterogeneous sources is challenging due to differences in frequency, reporting conventions, and missing values. To address this, the data have been collected from publicly available sources and processed to ensure consistent formatting and alignment on a common daily time index. Lower-frequency macroeconomic variables are transformed into daily series using forward-fill and backward-fill imputation. IBEX-FinTime is designed as a reusable resource for research in artificial intelligence, econometrics, and time-series forecasting. Its structure enables the direct application of machine learning and statistical models without extensive preprocessing. In addition, we provide an initial benchmark of recent foundation models for time series forecasting, including Chronos-2, Moirai 2.0, TimesFM, and TabPFN-TS, establishing baseline results on a subset of Spanish banking stocks. The benchmark illustrates the suitability of IBEX-FinTime as a reproducible evaluation framework for modern forecasting methods. The dataset and accompanying documentation are publicly available, supporting reproducibility and further research.

Keywords: Time series · Financial forecasting · Foundation models · Data harmonization · IBEX 35 · Spanish economy

1 Introduction

The growing adoption of machine learning methods has transformed a wide range of domains, including finance, healthcare, and economic analysis [11,19]. More recently, foundation models for time series forecasting have emerged as a promising paradigm, leveraging large-scale pretraining to perform forecasting tasks across diverse domains with limited task-specific adaptation. Evaluating the strengths and limitations of these models under realistic forecasting scenarios has therefore become an important research direction in both artificial intelligence and quantitative finance [8,21].

Despite these advances, a substantial part of the data science workflow is still devoted to data collection, cleaning, and harmonization. This issue is especially relevant in financial forecasting tasks, where data are typically obtained from heterogeneous sources with different formats, frequencies, and reporting conventions [11]. As a result, researchers often spend considerable effort integrating and preprocessing time series before applying predictive models. This preparation stage can become a major bottleneck in empirical research and limits the reproducibility of forecasting pipelines.

The IBEX 35 is the main benchmark index of the Spanish stock market and includes the most representative publicly traded companies in terms of liquidity and free float adjusted market capitalization. Although several existing datasets and repositories provide macroeconomic indicators or financial market variables separately, relatively few resources integrate both types of information within a unified and temporally aligned framework suitable for machine learning applications. At the same time, the recent emergence of foundation models for time series forecasting has increased the need for realistic and reproducible benchmark datasets that combine heterogeneous financial and macroeconomic information. Existing benchmark collections are often domain agnostic or lack the macro-financial structure required to evaluate forecasting models in realistic stock market environments.

To address these limitations, this work introduces IBEX-FinTime, a harmonized dataset combining daily financial market variables with macroeconomic indicators related to the Spanish economy in an international context. The dataset integrates IBEX 35 stock prices, international market indicators, exchange rates, commodity prices, government bond yields, and macroeconomic variables from Spain, the United States, and the Euro Area into a single daily aligned structure.

Beyond the dataset itself, we also provide an initial benchmark of modern foundation models for time series forecasting. This benchmark illustrates how the proposed dataset can be used as a reproducible evaluation framework for studying the performance of pretrained forecasting models under different financial and macroeconomic conditions.

The remainder of this paper is organized as follows. Section 2 reviews related work on large scale economic and financial datasets. Section 3 describes the construction of the proposed dataset, including data sources, preprocessing, harmonization procedures, and dataset organization. Section 4 presents the benchmark of foundation models. Section 5 discusses potential applications, lim-

itations, and research opportunities enabled by the dataset. Finally, Section 6 concludes the paper and outlines future directions.

2 Related Work

This section reviews two lines of research relevant to this work: the construction of large-scale economic and financial datasets, and the recent emergence of foundation models for time series forecasting.

2.1 Economic and Financial Datasets

The construction of large scale economic datasets from heterogeneous sources has become a central task in empirical macroeconomics and data-driven research. Several contributions have focused on integrating dispersed data into unified and reusable resources, enabling advanced econometric and machine learning applications [9,14].

Several large-scale economic databases have been developed to harmonize heterogeneous data sources and support empirical analysis. The Global Macro Database [18] extends the temporal and cross-country coverage of macroeconomic indicators through the integration of contemporary and historical sources, while the Korea Research Economic Database (KRED) [3] provides a standardized framework for macroeconomic forecasting based on multiple national datasets and aligned with FRED MD [15]. Similarly, the Universal Database for Economic Complexity [20] combines international trade data and applies machine learning techniques to address missing values and inconsistencies. Despite their value for macroeconomic and structural economic research, these datasets focus primarily on annual or monthly indicators and are therefore less suitable for high-frequency financial forecasting and stock market prediction tasks.

Despite these advances, existing datasets typically focus either on macroeconomic indicators or on specific economic domains, often at low or medium frequencies. In contrast, relatively less attention has been paid to the construction of datasets that jointly integrate high-frequency financial market data with macroeconomic variables in a unified framework [4].

The dataset proposed in this work addresses this gap by providing a harmonized collection of daily financial and macroeconomic time series centered on the Spanish economy. Unlike previous datasets, it combines stock market data, including IBEX 35 constituents, with international financial indicators and macroeconomic variables within a single temporally aligned structure. This design enables the direct application of machine learning and statistical models for forecasting tasks, particularly in the context of financial markets.

By bridging macroeconomic information and high-frequency financial data, the proposed dataset is specifically tailored for stock market prediction, offering a reusable benchmark for evaluating forecasting methods and studying macro-financial interactions.

2.2 Foundation Models for Time Series Forecasting

Recent advances in machine learning have led to the emergence of foundation models for time series forecasting [16]. Inspired by the success of large language models, these approaches are pretrained on large collections of heterogeneous time series and subsequently applied to downstream forecasting tasks with little or no task-specific training. Their goal is to learn transferable temporal representations that generalize across datasets and domains.

Among the most influential models in this area are TimesFM [5], a decoder-only transformer trained on large-scale time series corpora, and Chronos [2], which formulates forecasting as a sequence modeling problem through the tokenization of continuous time series. Building upon this paradigm, Chronos-2 [1] extends the framework with richer contextual representations and support for more general forecasting settings.

More recently, Moirai 2.0 [13] has focused on improving forecasting performance through efficient large-scale pretraining, while TabPFN-3 [10] demonstrates how foundation models for tabular data can also be adapted to forecasting tasks. The emergence of these models has increased the need for realistic benchmark datasets capable of evaluating their performance in domain-specific applications such as financial forecasting.

3 Dataset

The IBEX-FinTime dataset comprises 7,670 daily observations spanning from January 2005 to December 2025, covering 86 variables related to financial markets and macroeconomic activity in Spain, the Euro Area, and the United States. The dataset has been designed to provide a unified and temporally aligned resource suitable for time-series analysis and machine learning applications. The following subsections describe the data sources employed, the preprocessing and harmonization procedures, the final structure of the dataset, and its main limitations.

3.1 Data Sources

The dataset is constructed from multiple publicly available financial and statistical sources, combining both automated and manual data acquisition procedures. Table 1 summarizes the main data categories, their corresponding sources, and the methods used to retrieve them.

Financial market data, including stock prices, indices, exchange rates, and commodities, are obtained from Yahoo Finance [22] using Python-based tools. For equity and index series, adjusted closing prices (`Adj Close` in yfinance) are used in order to account for dividends, stock splits, and other corporate actions.

Macroeconomic indicators for Spain and the Euro Area are retrieved from Eurostat [6], while Euribor rates are collected through the European Central Bank API. U.S. macroeconomic variables are sourced from the FRED database [7], and government bond yields are manually downloaded from Investing.com [12] when automated access is not available.

Table 1. Main data sources and acquisition methods.

Data	Source	Acquisition
Stock prices, international stock indices, exchange rates, and commodity prices	Yahoo Finance	Python using `yfinance`
Macroeconomic indicators for Spain and the Euro Area, including GDP, inflation, and unemployment	Eurostat	Python API
Euribor rates with one-month, three-month, and twelve-month maturities	European Central Bank	Public API
U.S. macroeconomic indicators, including GDP, inflation, unemployment, and SOFR	FRED	Manual CSV download
Government bond yields for Spain and the United States	Investing.com	Manual CSV download

3.2 Data Processing and Harmonization

Data processing aims to ensure temporal consistency, structural uniformity, and full integration of heterogeneous time series variables into a single dataset. The construction pipeline combines automated data retrieval and manual collection, followed by a series of preprocessing steps including cleaning, temporal alignment, and imputation.

All variables are stored in CSV format and processed using Python 3.11 with standard data analysis libraries. Data from different sources are first parsed into structured data frames and then merged into a unified dataset indexed by date. To enable joint analysis, all series are aligned to a common daily frequency.

The unified time index is based on calendar days rather than trading days. Consequently, the dataset includes weekends and holidays in order to preserve temporal continuity across heterogeneous financial and macroeconomic variables. Missing observations associated with non-trading days are handled through forward-fill imputation for financial variables [17].

Macroeconomic variables originally reported at monthly or quarterly frequency are transformed into daily series through forward-fill and backward-fill procedures, depending on the reporting convention of each source. In particular, values reported at the beginning of a period are propagated forward, while values reported at the end of a period are propagated backward. This distinction is necessary because different data providers associate the same macroeconomic observation with different calendar dates. For example, some sources register a monthly indicator on the first day of the corresponding month, whereas others associate the same value with the last day of the reporting period.

Financial variables, such as stock prices, are consistently forward-filled to preserve their temporal dynamics. The harmonization strategy is designed pri-

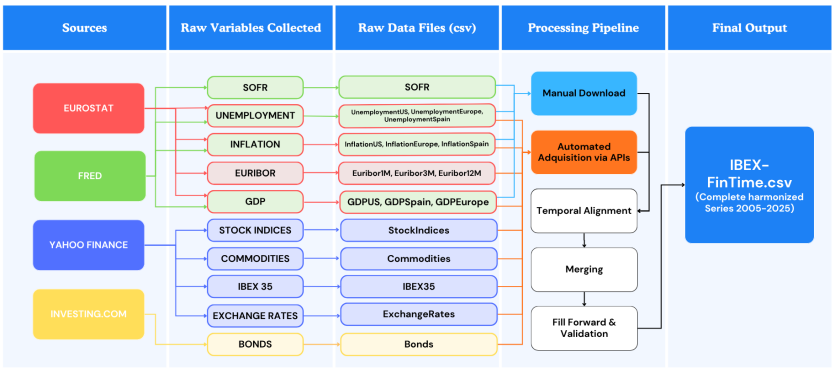


Fig. 1. Workflow of data acquisition, preprocessing, harmonization, and integration into the IBEX-FinTime dataset.

marily for multivariate financial analysis and stock market forecasting rather than strict real-time macroeconomic forecasting evaluation.

Missing dates across variables are synchronized to ensure full temporal coverage. The resulting dataset is therefore fully aligned, allowing direct application of statistical, econometric, and machine learning methods without additional preprocessing.

Figure 1 summarizes the complete workflow used to acquire, preprocess, harmonize, and integrate the financial and macroeconomic variables included in IBEX-FinTime. The pipeline illustrates the sequential stages from raw data acquisition to temporal alignment and final dataset construction.

3.3 Dataset Structure and Description

The dataset is distributed as a single compressed repository containing both processed and raw data, together with supporting documentation. Its internal structure is designed to facilitate reproducibility and ease of use. The final version of IBEX-FinTime contains 86 variables covering multiple dimensions of financial markets and macroeconomic activity. The dataset integrates Spanish equity prices, international market indices, exchange rates, commodities, sovereign bond yields, interest rate benchmarks, and macroeconomic indicators for Spain, the United States, and the Euro Area. Table 2 summarizes the main groups of variables included in the dataset.

In addition, the repository includes an exploratory report, which provides descriptive statistics for each variable, and a directory containing the original data files prior to integration, organized by data category. This structure allows

Table 2. Main categories of variables included in the IBEX-FinTime dataset.

Category	Variables
IBEX 35 constituent stocks	ACS, Acerinox, AENA, Amadeus, Acciona, Acciona Energía, ArcelorMittal, BBVA, Bankinter, CaixaBank, Cellnex, Colonial, Endesa, Enagás, Fluidra, Ferrovial, Grifols, IAG, Iberdrola, Indra, Inditex, Logista, Mapfre, Merlin, Naturgy, Puig, Redeia, Repsol, Rovi, Sabadell, Santander, Sacyr, Solaria, Telefónica, Unicaja
International stock indices	CAC 40, FTSE 100, DAX, S&P 500, IBEX 35, NASDAQ 100, Euro Stoxx 50
Exchange rates	EURBRL, EURCHF, EURCNY, EURGBP, EURJPY, EURMXN, EURUSD
Commodity market variables	BrentOil_USD, Cocoa_USD, WTIOil_USD, Cotton_USD, Gold_USD, LeanHog_USD, Copper_USD, Coffee_USD, LiveCattle_USD, NaturalGas_USD, Platinum_USD, Sugar_USD, Silver_USD, NaturalGas_NL_EUR, Corn_USD, Soy_USD, Wheat_USD
Sovereign bond yields	SpainBond3M, SpainBond5Y, SpainBond10Y, USBond3M, USBond5Y, USBond10Y
Interest rate benchmarks	Euribor1M, Euribor3M, Euribor12M, SOFR
Gross Domestic Product (GDP)	GDP_ES, GDP_EA20, GDP_US
Inflation indicators	HICP_SP, HICP_EA20, SPCPI_US
Unemployment indicators	Unemployment_SP, Unemployment_EA20, Unemployment_US

users to trace each variable back to its origin and to reproduce or extend the dataset construction process.

Overall, the dataset provides a consistent and well-documented structure that supports direct use in time-series analysis and data-driven applications.

3.4 Dataset Limitations

Some limitations should be considered when using the dataset. First, real GDP data for Spain and the Euro Area are provided as chain-linked volumes referenced to 2015, while U.S. GDP data are referenced to 2017. This difference may introduce minor scale inconsistencies in cross-country comparisons.

Second, the U.S. inflation series is seasonally adjusted, whereas the Eurostat-based inflation data for Spain and the Euro Area do not explicitly indicate seasonal adjustment. This discrepancy may affect short-term comparability between regions.

In addition, differences in market calendars and data availability imply that holidays and non-trading days have been handled using forward-fill imputation, which may smooth short-term variations in the series.

Finally, several stock price series contain missing values outside their inclusion periods in the IBEX 35 index, as companies enter and exit the index over time. This should be taken into account when performing time-series analyses or aggregating across assets. These unavailable observations are preserved as NaN values in the final CSV file rather than being imputed or replaced by zeros, allowing users to explicitly handle missing data according to the requirements of their forecasting or machine learning pipeline.

4 Benchmark of Foundation Models

To demonstrate the usefulness of IBEX-FinTime as a benchmarking resource, we conduct an initial evaluation of several recent foundation models for time series

forecasting. The objective of this benchmark is not to provide an exhaustive comparison of forecasting methods, but rather to illustrate how the dataset can be used as a reproducible evaluation framework for modern forecasting models in a financial context.

4.1 Experimental Setup and Models

The benchmark focuses on six banking stocks from the IBEX 35 index: BBVA, Bankinter, CaixaBank, Sabadell, Santander, and Unicaja. These companies were selected because they belong to the same economic sector while exhibiting different market capitalizations, volatility levels, and temporal dynamics.

A univariate forecasting setting is considered, where each stock is modeled independently using only its historical adjusted closing prices. The evaluation period spans from January 2024 to December 2025. For all models, a context window of 64 observations is employed to generate forecasts. All models are evaluated in a zero-shot setting, without fine-tuning, using a rolling-origin scheme with a forecast horizon of $h = 1$ day.

Forecasting accuracy is evaluated using the Mean Absolute Error (MAE) and Symmetric Mean Absolute Percentage Error (SMAPE), defined as

$$MAE = \frac{1}{N} \sum_{i=1}^{N} |y_i - \hat{y}_i|, \tag{1}$$

$$SMAPE = \frac{100}{N} \sum_{i=1}^{N} \frac{2|y_i - \hat{y}_i|}{|y_i| + |\hat{y}_i|}. \tag{2}$$

where y_i denotes the observed value and $\hat{y}_i$ the corresponding prediction.

In addition, the quality of uncertainty estimation is assessed through *Coverage@20*, defined as the proportion of observed values contained within the predicted 20% prediction interval. An ideal model should achieve a Coverage@20 value close to 0.20, indicating that the empirical coverage matches the nominal confidence level.

The benchmark includes four recent foundation models representing different approaches to time series forecasting: Chronos-2 [1], Moirai 2.0 [13], TimesFM [5] and TabPFN-TS [10], the time-series adaptation of the TabPFN-3 tabular foundation model.

4.2 Results

Table 3 reports the forecasting performance of the evaluated models on the selected banking stocks. The results show that all evaluated foundation models achieve competitive forecasting performance on the selected banking stocks. Moirai 2.0 obtains the lowest average MAE and SMAPE, and the lowest MAE in four of the six series. TimesFM attains the lowest MAE for Bankinter, and ties with TabPFN-TS for Sabadell (0.0304), where it also achieves the lowest SMAPE.

Table 3. Forecasting performance of the evaluated foundation models on the banking sector benchmark with C2 (Chronos-2), M2 (Moirai 2.0), TPFN (TabPFN-TS), and TFM (TimesFM).

	MAE				SMAPE				Coverage@20			
Stock	C2	M2	TPFN	TFM	C2	M2	TPFN	TFM	C2	M2	TPFN	TFM
BBVA	0.1603	**0.1573**	0.1677	0.1645	1.4208	**1.3952**	1.4936	1.4581	0.1778	**0.2008**	0.2084	0.2103
Bankinter	0.1073	0.1031	0.1120	**0.1030**	1.2289	**1.1808**	1.2824	1.1875	0.1472	**0.1874**	0.2199	0.1740
CaixaBank	0.0787	**0.0774**	0.0817	0.0784	1.3131	**1.2931**	1.3832	1.3040	0.1702	**0.1989**	0.2180	0.1931
Sabadell	0.0309	0.0307	**0.0304**	**0.0304**	1.4802	1.4702	1.4695	**1.4519**	**0.1702**	0.1644	0.2543	0.1530
Santander	0.0795	**0.0773**	0.0809	0.0807	1.4115	**1.3735**	1.4512	1.4334	0.1702	**0.1989**	0.2218	0.1549
Unicaja	0.0192	**0.0191**	0.0205	0.0194	1.3118	**1.2935**	1.4084	1.3231	0.1702	**0.2161**	0.2371	0.1721
Average	0.0793	**0.0775**	0.0822	0.0794	1.3611	**1.3344**	1.4147	1.3597	0.1676	**0.1944**	0.2266	0.1762

Overall, the differences among models remain relatively small, suggesting that all evaluated approaches are capable of capturing the short-term dynamics of Spanish banking stocks. These findings highlight the suitability of IBEX-FinTime as a benchmark dataset for evaluating modern forecasting models under a common and reproducible experimental protocol.

5 Discussion

The IBEX-FinTime dataset provides an integrated view of Spanish financial market dynamics within a broader macroeconomic and international context. Unlike datasets focused exclusively on macroeconomic indicators or stock market data, IBEX-FinTime combines IBEX 35 constituents, international stock indices, exchange rates, commodities, bond yields, interest rate benchmarks, and macroeconomic indicators into a single daily aligned structure. This design reduces the preprocessing burden commonly required before applying statistical, econometric, or machine learning models to financial forecasting tasks.

Figure 2 illustrates the temporal evolution of the six banking stocks used in the benchmark: BBVA, Bankinter, CaixaBank, Sabadell, Santander, and Unicaja. These series exhibit heterogeneous behaviors, including periods of sustained growth, varying levels of volatility, and the impact of major economic events such as the 2008 global financial crisis, the European sovereign debt crisis, the COVID-19 shock, and the subsequent interest rate tightening cycle. The diversity of temporal patterns observed in these assets highlights the suitability of the dataset for evaluating forecasting models under different market conditions and motivates their selection for the foundation model benchmark presented in Section 4.

The benchmark results further illustrate the usefulness of IBEX-FinTime as a forecasting benchmark. Despite being evaluated under a simple univariate setting restricted to banking stocks, all four foundation models achieved competitive performance across the six series considered. The relatively small differences in MAE suggest that the selected stocks exhibit predictable temporal patterns that can be effectively captured by modern pretrained forecasting architectures.

Among the evaluated models, Moirai 2.0 obtained the lowest average MAE and SMAPE, and the lowest MAE in most of the considered series. Nevertheless,

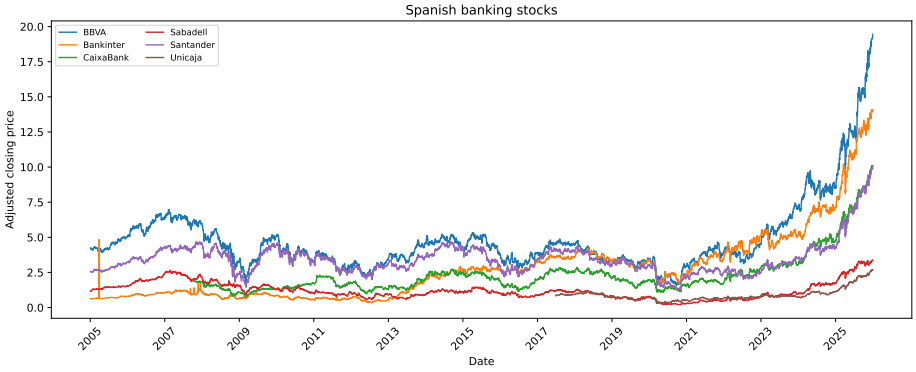


Fig. 2. Time series trends of the six Spanish banking stocks used in the benchmark.

no single model consistently dominated all assets. These findings indicate that different foundation model architectures may capture complementary aspects of financial time series dynamics.

From a benchmarking perspective, the results demonstrate that IBEX-FinTime provides a realistic and challenging environment for evaluating forecasting models. The dataset contains heterogeneous market conditions spanning multiple economic regimes, while simultaneously offering sufficient structure for meaningful model comparisons. This makes the dataset suitable not only for benchmarking foundation models, but also for future comparisons involving statistical methods, deep learning architectures, and hybrid forecasting approaches.

From a research perspective, IBEX-FinTime can be reused in several directions. First, it enables the development and benchmarking of forecasting models for Spanish stock market variables using both domestic and international explanatory factors [4,21]. Second, it supports studies on macro-financial interactions, allowing researchers to analyze how Spanish market behavior relates to unemployment, inflation, GDP, interest rates, commodities, and global equity indices. Third, its transparent structure makes it possible to replicate, extend, or adapt the dataset to other markets, facilitating comparative studies and reproducible experimentation.

The inclusion of both raw and processed data also increases the value of the dataset for applied research. Users can directly employ the harmonized daily dataset for modeling, while also accessing the original files to inspect sources, modify preprocessing decisions, or extend the repository with additional variables. Therefore, IBEX-FinTime is not only a dataset for a specific forecasting task, but also a reusable resource for artificial intelligence, econometrics, and time series research focused on financial markets.

6 Conclusions and Future Work

This paper presented IBEX-FinTime, a harmonized dataset integrating financial market variables and macroeconomic indicators related to the Spanish economy within an international context. The dataset combines daily stock market data, international indices, exchange rates, commodity prices, bond yields, and macroeconomic indicators into a single temporally aligned structure designed for forecasting and data-driven analysis. By addressing differences in frequency, reporting conventions, and missing values through a transparent harmonization pipeline, IBEX-FinTime facilitates the direct application of statistical, econometric, and machine learning methods. Furthermore, the benchmark of foundation models demonstrates the suitability of the dataset as a reproducible evaluation framework for modern forecasting approaches. The temporal coverage from 2005 to 2025 includes multiple market regimes, increasing the value of the dataset for studying financial forecasting under diverse economic conditions. Future work will focus on extending the benchmark to multivariate forecasting settings and additional asset classes, enabling a more comprehensive evaluation of foundation models. Other directions include incorporating alternative data sources such as financial news and sentiment indicators, as well as automating data acquisition procedures to support continuous updates of the dataset.

Data Availability

The IBEX-FinTime dataset, including the processed CSV file, raw data sources, preprocessing scripts, and supplementary documentation, is publicly available on GitHub[3].

Acknowledgments. This research has been supported by grant PID2023-146037OB-C21 funded by MICIU/AEI/10.13039/501100011033. The authors would like to thank the Spanish Ministry of Universities for the FPU predoctoral fellowship (FPU23/03390).

References

1. Ansari, A.F., Shchur, O., Küken, J., Auer, A., Han, B., Mercado, P., et al.: Chronos-2: From univariate to universal forecasting (2025), https://arxiv.org/abs/2510.15821
2. Ansari, A.F., Stella, L., Turkmen, C., Zhang, X., Mercado, P., Shen, H., et al.: Chronos: Learning the language of time series (2024), https://arxiv.org/abs/2403.07815
3. Baek, C., Moon, S.H., Lee, S.: KRED: Korea Research Economic Database for Data-Rich Policy Analysis. arXiv preprint arXiv:2509.16115 (2025)
4. Consoli, S., Negri, M., Tebbifakhr, A., Tosetti, E., Turchi, M.: Forecasting the IBEX-35 Stock Index Using Deep Learning and News Emotions. In: Machine Learning, Optimization, and Data Science (LOD 2021). Lecture Notes in Computer Science, vol. 13163, pp. 308–323. Springer (2022). https://doi.org/10.1007/978-3-030-95467-3_23

[3] https://github.com/Jocoplex6/IBEX-FinTime

5. Das, A., Kong, W., Sen, R., Zhou, Y.: A decoder-only foundation model for time-series forecasting (2024), https://arxiv.org/abs/2310.10688
6. Eurostat: Eurostat Database. https://ec.europa.eu/eurostat, European Commission. Accessed: May 2026
7. Federal Reserve Bank of St. Louis: Federal Reserve Economic Data (FRED). https://fred.stlouisfed.org/, Accessed: May 2026
8. Fischer, T., Krauss, C.: Deep learning with long short-term memory networks for financial market predictions. European Journal of Operational Research **270**(2), 654–669 (2018). https://doi.org/10.1016/j.ejor.2017.11.054
9. Godahewa, R., Bergmeir, C., Webb, G.I., Hyndman, R.J., Montero-Manso, P.: Monash Time Series Forecasting Archive. In: Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (NeurIPS 2021) (2021)
10. Grinsztajn, L., Flöge, K., Key, O., Birkel, F., Jund, P., Roof, B., et al.: TabPFN-3: Technical report (2026), https://arxiv.org/abs/2605.13986
11. Henrique, B.M., Sobreiro, V.A., Kimura, H.: Literature review: Machine learning techniques applied to financial market prediction. Expert Systems with Applications **124**, 226–251 (2019). https://doi.org/10.1016/j.eswa.2019.01.012
12. Investing.com: Financial Markets Data and News. https://www.investing.com/, Accessed: May 2026
13. Liu, C., Aksu, T., Liu, J., Liu, X., Yan, H., Pham, Q., Savarese, S., Sahoo, D., Xiong, C., Li, J.: Moirai 2.0: When less is more for time series forecasting (2026), https://arxiv.org/abs/2511.11698
14. Makridakis, S., Spiliotis, E., Assimakopoulos, V.: The M5 competition: Background, organization, and implementation. International Journal of Forecasting **38**(4), 1325–1336 (2022). https://doi.org/10.1016/j.ijforecast.2021.07.007
15. McCracken, M.W., Ng, S.: FRED-MD: A Monthly Database for Macroeconomic Research. Journal of Business & Economic Statistics **34**(4), 574–589 (2016). https://doi.org/10.1080/07350015.2015.1086655
16. Meseguer-Gómez, J.I., Reina-Jiménez, P., Jiménez-Navarro, M.J., Martínez-Ballesteros, M., Riquelme-Santos, J.C.: Exploring zero-shot learning for time series forecasting using large language models. In: Proceedings of Soft Computing Models in Industrial and Environmental Applications, SOCO 2025. Communications in Computer and Information Science. vol. 2806, pp. 217–226 (2026). https://doi.org/10.1007/978-3-032-19763-4_20
17. Moritz, S., Bartz-Beielstein, T.: imputeTS: Time series missing value imputation in R. The R Journal **9**(1), 207–218 (2017). https://doi.org/10.32614/RJ-2017-009
18. Müller, K., Xu, C., Lehbib, M., Chen, Z.: The Global Macro Database: A New International Macroeconomic Dataset (Version 2026-03). Working Paper 33714, National Bureau of Economic Research (April 2025). https://doi.org/10.3386/w33714, http://www.nber.org/papers/w33714
19. Nazareth, N., Ramana Reddy, Y.V.: Financial applications of machine learning: A literature review. Expert Systems with Applications **219**, 119640 (2023). https://doi.org/10.1016/j.eswa.2023.119640
20. Patelli, A., Zaccaria, A., Pietronero, L.: Universal Database for Economic Complexity. arXiv preprint arXiv:2110.00302 (2021)
21. Sezer, O.B., Gudelek, M.U., Ozbayoglu, A.M.: Financial time series forecasting with deep learning: A systematic literature review: 2005–2019. Applied Soft Computing **90**, 106181 (2020). https://doi.org/10.1016/j.asoc.2020.106181
22. Yahoo Finance: Historical Financial Market Data for IBEX-35 Companies. https://finance.yahoo.com/, Yahoo Inc. Accessed: May 2026