DOCUMENTOS DE TRABAJO

Artificial Intelligence Models for Nowcasting Economic Activity

Jennifer Peña
Katherine Jara
Fernando Sierra

N° 1058 Diciembre 2025
BANCO CENTRAL DE CHILE

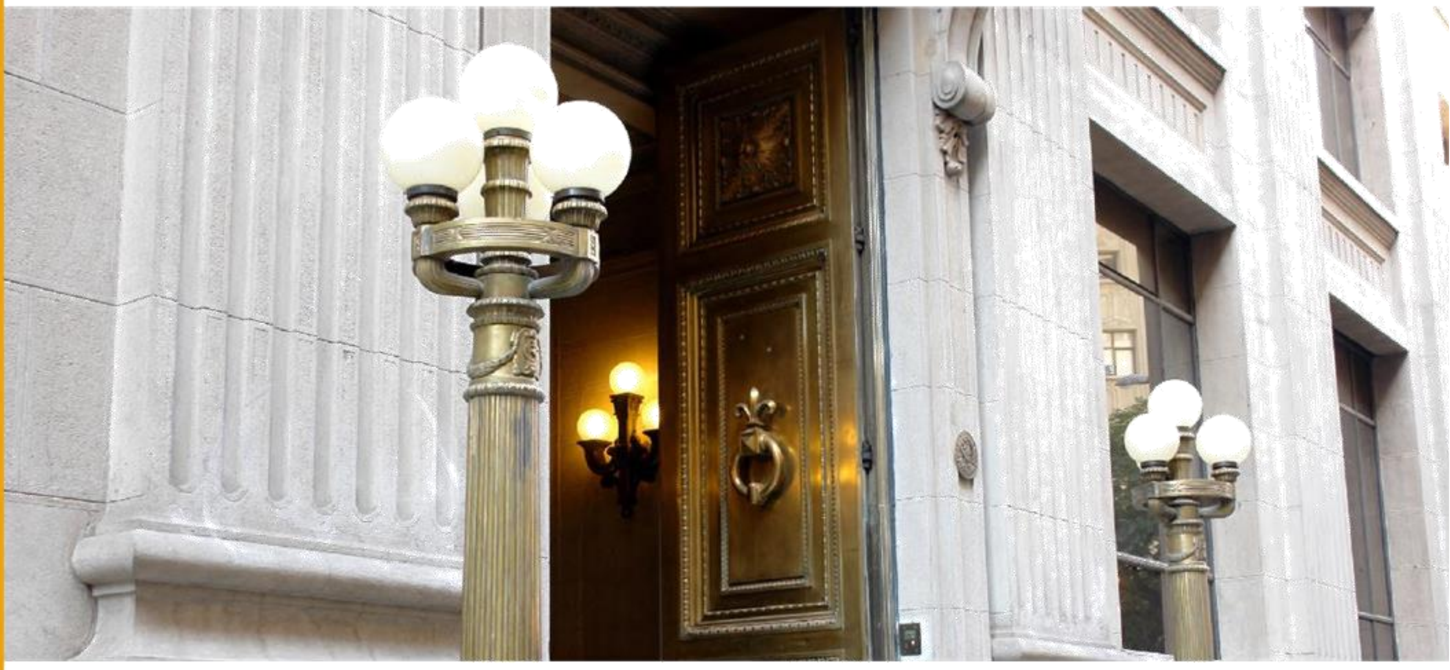

banco central Chile

Documentos de Trabajo del Banco Central de Chile
Working Papers of the Central Bank of Chile
Agustinas 1180, Santiago, Chile
Teléfono: (56-2) 3882475; Fax: (56-2) 38822311

Documento de Trabajo N° 1058 | **Working Paper N° 1058**

Artificial Intelligence Models for Nowcasting Economic Activity*

Jennifer Peña, Katherine Jara and Fernando Sierra†
Central Bank of Chile

Resumen
Este estudio analiza si las técnicas de inteligencia artificial—que abarcan modelos de aprendizaje automático y aprendizaje profundo—pueden mejorar la precisión de las estimaciones en tiempo real del índice mensual de actividad económica de Chile (IMACEC). El análisis se basa en un conjunto de datos amplio y diverso, que incluye tanto variables macroeconómicas tradicionales como datos de registros administrativos mensuales (provenientes de los registros tributarios electrónicos). Surgen tres hallazgos principales. Primero, los modelos no lineales—en particular XGBoost—presentan los menores errores cuadráticos medios, destacan también SVR y la regresión LASSO, por el lado de los lineales. Esto resalta el valor de los métodos no lineales y de los enfoques lineales regularizados al trabajar con datos heterogéneos. Segundo, las variables derivadas de los registros tributarios electrónicos—como el crédito comercial entre empresas y las ventas de las empresas sectoriales por región—se ubican de manera consistente entre los predictores más relevantes en todos los modelos. Tercero, los modelos con mejor desempeño—XGBoost, SVR y LASSO—logran errores más bajos que los referentes econométricos tradicionales, los cuales se basan únicamente en agregados macroeconómicos estándar y excluyen los datos no tradicionales. En general, los resultados muestran que los datos de registros administrativos oportunos, combinados con enfoques de inteligencia artificial, pueden mejorar significativamente el monitoreo económico y la toma de decisiones.

Abstract
This paper investigates whether artificial intelligence techniques—encompassing both machine learning and deep learning models—can enhance the accuracy of now-casts for Chile's monthly economic activity index (IMACEC). The analysis relies on a large and diverse real-time dataset that includes both traditional macroeco-nomic variables and high-frequency monthly administrative data (from electronic tax records). Three main findings emerge. First, nonlinear models—particularly XGBoost—achieve the lowest root mean squared errors, whereas linear regularized approaches such as SVR and LASSO also show competitive performance. This highlights the value of flexible nonlinear methods and regularized linear approaches when dealing with heterogeneous data. Second, features derived from electronic tax records—such as trade credit volumes and sectoral sales by region—consistently rank among the most important predictors across models. Third, the strongest-performing models—XGBoost, SVR, and LASSO—achieve lower errors than tra-ditional econometric benchmarks, which rely solely on standard macroeconomic aggregates and exclude non-traditional datasets. Overall, the findings show that timely administrative data, combined with AI approaches, can significantly improve economic surveillance and decision-making.

* We are very grateful for the valuable feedback received during the seminar at the Central Bank of Chile and the XX Interdisciplinary Social Sciences Conference in Spain, as well as for the useful comments provided by Markus Kirchner and Andrés Ceballos. The views expressed herein are those of the authors and do not necessarily reflect the views of the Central Bank of Chile. This study was developed within the scope of the research agenda conducted by the Central Bank of Chile (CBC) in economic and financial affairs of its competence. The CBC has access to anonymized information from various public and private entities, by virtue of collaboration agreements signed with these institutions. To secure the privacy of workers and firms, the CBC mandates that the development, extraction and publication of the results should not allow the identification, directly or indirectly, of natural or legal persons. Officials of the Central Bank of Chile processed the disaggregated data. All the analysis was implemented by the authors. The information contained in the databases of the Chilean IRS is of a tax nature originating in self-declarations of taxpayers presented to the Service; therefore, the veracity of the data is not the responsibility of the Service.
†E-mails: respectively, jpena@bcentral.cl, kjara@bcentral.cl, fsierra@bcentral.cl

1 Introduction

The ability to predict economic developments in real time under high uncertainty has taken unprecedented relevance with the COVID-19 pandemic and other recent shocks. Nowcasting procedures have been developed to deliver real time estimates of economic indicators using diverse data sources to address the challenge of providing timely economic insights. In the literature, nowcasting is defined as predicting the present or very near future/recent past, distinguishing it from traditional forecasting, which generally focuses on more distant future periods (Banbura et al., 2010).

Traditionally and following Breiman (2001) and more recently Jin-Kyu Jung (2018), nowcasting models have been oriented within the data modeling approach, which seeks to model the structural relationships between underlying economic factors and observed data to better understand the data-generating process. Meanwhile, the algorithmic modeling has recently gained relevance in economics, using algorithms to learn patterns directly from the data and optimize outcomes without explicit structural assumptions about the data-generating process; examples of this approach are artificial intelligence (AI) models.

Understanding the relationship between AI, machine learning (ML), and deep learning (DL) is essential. AI constitutes a broad field concerned with developing systems capable of performing tasks traditionally associated with human intelligence, such as reasoning, learning, and decision-making. Within this domain, ML represents a prominent subset focused on the construction of algorithms and models that allow computers to identify patterns from historical data and improve their performance on specific tasks autonomously without the need for explicit programming[1]. The primary objective of ML is to enable systems to generalize from past experiences to optimize future task performance (Hetland & Nelli, 2023). DL, in turn, is a specialized subfield of ML, distinguished by the use of multi-layered neural network architectures capable of extracting hierarchical representations from large and complex datasets. The relationship among these fields is thus inherently hierarchical: DL is a subset of ML, forming a central component within the broader framework of AI.

The literature on various ML techniques has grown in recent decades. Machine learning techniques have become popular among economists. Machine learning algorithms build statistical models to explain data and then optimize the fit based on past data (supervised learning). The algorithms continuously and autonomously update the parameters until a threshold of accuracy has been met. It is important for policymakers,

[1] Artificial intelligence encompasses not only machine learning but also other approaches such as expert systems and reinforcement learning frameworks.

particularly for central banks, to understand if these techniques can be leveraged for activity nowcasting. In the case of nowcasting, machine learning algorithms have the advantage that they can efficiently handle large datasets with lots of potential regressors. Thus, ML algorithms can solve the curse of dimensionality when there are many potential predictors and are a promising alternative to the usual time-series regression based methods used by central banks for nowcasting, so machine learning presents new opportunities for enhancing nowcasting accuracy.

In the literature, several studies have shown that machine learning techniques can improve GDP forecasting and nowcasting across diverse contexts. For example, Richardson et al. (2021) finds that boosted trees, support vector machines, and neural networks outperform autoregressive and dynamic factor models in New Zealand. Similar improvements are reported in Lebanon (Tiffin, 2016), Turkey (Bolhuis & Rayner, 2020), and the U.S. (Arro-Cannarsa & Scheufele, 2024); in each case, machine learning reduces forecast errors compared to traditional econometric approaches. Even with limited or less granular data, studies such as Jin-Kyu Jung (2018) show that methods like Elastic Net and recurrent neural networks provide superior accuracy. At the regional level, Barrios et al. (2021) demonstrate that machine learning can generate reliable quarterly GDP predictions in small, volatile economies such as Belize and El Salvador. Collectively, this evidence highlights the growing role of machine learning as a powerful alternative to conventional models in macroeconomic forecasting.

The contributions of this study are threefold. First, it advances the scarce Latin American literature on ML-based nowcasting by analyzing the Chilean case, where the Monthly Economic Activity Index (IMACEC)—the main short-term indicator of economic activity published by the Central Bank—is available only with a one-month lag [2]. Having a nowcast in advance is therefore valuable for keeping policymakers informed about the dynamics of economic activity. Second, the study extends the methodological scope by incorporating both machine learning and deep learning approaches—penalized regressions (LASSO, Ridge, Elastic Net), SVR, XGBoost, MLP, and 1D-CNN—thus providing a systematic comparison of their performance. Third, it integrates both traditional macroeconomic indicators and non-traditional data sources, such as firm-level sales, electronic payments, and climate variables (temperature and precipitation), which capture sectoral and geographic dynamics of the Chilean economy. By leveraging these diverse inputs, the study highlights how AI techniques can “let the data speak” and enhance the timeliness and accuracy of real-time monitoring.

Initially, we collected 187 variables as potential predictors from January 2018 to De-

[2]For further details, see IMACEC

cember 2024. These variables were transformed—when appropriate—into year-on-year growth rates, month-on-month variations, and seasonally adjusted series, resulting in a final set of approximately 400 transformed features. This extended feature set was used to evaluate the real-time performance of the proposed algorithms (see Appendix A for a detailed description). Importantly, to simulate a real-time nowcasting scenario, we ensured that only the information available up to the moment of each nowcast was used to construct and transform the predictors. The machine learning literature highlights the benefits of dimensionality reduction by selecting features that exhibit a positive and relatively strong correlation with the target variable, as this can enhance forecast accuracy (Kotsiantis, 2011). Therefore, a correlation analysis assessed the relationship between the IMACEC and the transformed features.

Finally, we compare the nowcasting performance of artificial intelligence models with that of traditional econometric approaches. Traditional models include SARIMA, BVAR, MIDAS, dynamic factor models, bridge models, among others, with their nowcasts obtained from Cobb & Peña (2020). To our knowledge, this is the first study to perform the nowcasting of Chilean economic activity using a comprehensive set of artificial intelligence models combined with non-traditional information sources. Furthermore, it is the first to assess the relative nowcasting performance of alternative machine learning methods using real-time data for Chile.

In summary, our results show that nonlinear machine learning methods—particularly XGBoost—consistently outperform other approaches in nowcasting Chile's economic activity. XGBoost achieves the lowest errors for the total IMACEC through gradient boosting and regularization. For the non-mining IMACEC, Support Vector Regression (SVR) with a linear kernel and LASSO also perform strongly, achieving errors comparable to XGBoost. Among linear models, SVR efficiently captures relationships in high-dimensional spaces, while LASSO stands out for reducing model complexity through variable selection, outperforming Ridge and Elastic Net. Deep learning models (MLP and 1D-CNN) show comparatively weaker performance, likely due to the short time series available. Monthly administrative records—such as trade credit and sectoral or regional sales from electronic tax data—consistently emerge as key predictors, underscoring the value of digital administrative data. Importantly, XGBoost, SVR, and LASSO outperform traditional econometric benchmarks, which rely exclusively on standard macroeconomic aggregates and exclude non-traditional datasets. These findings highlight the importance of flexible nonlinear methods, regularized linear techniques, and integrating monthly non-traditional datasets to enhance real-time economic monitoring and support timely policy decisions.

The remainder of the paper is structured as follows. Section 2 reviews the relevant literature, highlighting the role of machine learning and big data in nowcasting methodologies. Section 3 introduces the machine learning algorithms employed to predict the monthly growth rate of the IMACEC. Section 4 describes the data sources and provides the definitions of the variables used in the analysis. Section 5 presents the empirical findings, while Section 6 summarizes the main contributions and discusses policy implications.

2 Literature review

Following Breiman (2001), two cultures use statistical modeling to reach conclusions from data. One assumes that a given stochastic data model generates the data. The other uses algorithmic models and treats the data mechanism as unknown. The statistical community has been committed to the almost exclusive use of data models. Algorithmic modeling has developed rapidly in fields outside statistics (often called machine learning). It can be used both on large complex data sets and as a more accurate and informative alternative to data modeling on smaller data sets.

Recent studies have demonstrated the effectiveness of machine learning algorithms in nowcasting economic activity, particularly GDP growth. These models outperform traditional statistical approaches and can provide accurate real-time predictions, (Richardson et al., 2021). Machine learning nowcasting techniques have been successfully applied to diverse economies, including New Zealand, Belize, El Salvador, and India, showcasing their adaptability to different economic structures and volatility levels (Barrios et al., 2021, Ranjan & Ghosh, 2021). The models utilize various high-frequency indicators, including macroeconomic data, financial market information, and economic uncertainty indices, to generate timely forecasts. Machine learning algorithms have demonstrated the potential to improve official central bank forecasts (Richardson et al., 2021). These nowcasting tools are particularly valuable in bridging the information gap between the end of a quarter and the delayed publication of official macroeconomic indicators, enabling more agile policy decisions (Barrios et al., 2021).

2.1 International Literature

Machine learning methods have been implemented significantly in category prediction and have progressively begun to be used for nowcasting. In this context, Richardson et al. (2021) applies machine learning algorithms to approximately 600 national and international variables from 2009 to 2019 to predict New Zealand's GDP growth. The study proves that machine learning algorithms (driven trees, support vector machine, and neural networks) outperform a simple autoregressive model and DFM.

As discussed in Soybilgen & Yazgan (2021), which focuses on economic forecasting and presented improved quarterly GDP growth prediction rates for the United States from the second quarter of 2000 to the fourth quarter of 2018. They predict economic outcomes with tree-based machine learning models, such as bagged decision trees, random forests, and stochastic gradient increase, a pressing problem we address in this study is the problem of roughness that occurs during the temporal alignment of observations from different sources, leading to missing observations or an incomplete set of consistent observations together. To address this problem, the authors used the so-called dynamic factor model to extract a small number of common (or latent) factors from 10 sets of financial and macroeconomic data, thereby markedly decreasing the dimensionality of the dataset.

It is also found in Kant et al. (2022) which discusses a selection of econometric and machine learning methods for real-time Dutch GDP forecasting from 1992 to 2018. The authors compare the performance of these methods using monthly data from 83 economic indicators. These variables include both measures of economic and financial activity. The methodology uses random forests, neural networks, and regularized regression models (LASSO and Ridge). Looking at the results shows that random forests produce the most accurate predictions, which leads to the discussion that high-level machine learning techniques are better for making short-term predictions, especially with large databases that have complex information dimensions.

In a related study, Yoon (2021) applies gradient boosting and random forest models to forecast Japan's real GDP growth over the period from 2001–2018. Using cross-validation for hyperparameter tuning and benchmarks from the IMF and the Bank of Japan, they find that both ML models consistently outperform institutional forecasts, with gradient boosting delivering the highest accuracy. This evidence highlights the potential of ensemble-based approaches to improve macroeconomic forecasting relative to traditional methods.

Continuing Arro-Cannarsa & Scheufele (2024) investigates the effectiveness of machine learning (ML) algorithms in immediate GDP prediction using a large mixed-frequency dataset of more than 1,100 indicators for Switzerland It compares traditional statistical models, such as principal component analysis (PCA) and forward subset selection (FSS), with machine learning techniques, such as LASSO, Ridge regression, elastic network, random forest, gradient increase, and support vector regression (SVR). The results show that ML models, especially crest regression, elastic network, and SVR with a linear kernel, outperform traditional benchmarks, with up to a 28% reduction in out-of-sample RMSE.

Finally, there is Németh & Hadházi (2024) which evaluates five artificial neural network (ANN) architectures that predict quarterly GDP growth in the United States. Federal Reserve Economic Data Monthly Database (FRED-MD) is used, comparing the performance of the Multilayer Perceptron (MLP), one-dimensional Convolutional Neural Network (1D CNN), Recurrent Neural Network (RNN), Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) models. The authors cover two assessment periods: one characterized by stable economic growth (2012:Q1–2019:Q4) and another that includes the recession caused by the COVID-19 pandemic (2012:Q1–2022:Q4). The variables used in the study include macroeconomic indicators, such as production, labor market, housing, consumption, prices, and interest rates, in their monthly frequency in contrast to the target variable, GDP growth in quarterly frequency. The results show that the 1D CNN model has a better performance in capturing patterns of the economic environment and highlights its accuracy compared to complex models such as LSTM and GRU.

2.2 Literature Focused on Latin America

In Latin America, few studies compare machine learning models' performance with traditional models in nowcasting monthly activity. However, recent studies have explored machine learning techniques for nowcasting economic activity and COVID-19 cases in Latin America. Barrios et al. (2021) present machine learning models fitted to predict quarterly GDP for Belize and El Salvador. Their results show that machine learning techniques can produce accurate quarterly GDP forecasts for two contrasting economies in an economic context marked by high volatility at both the national and international levels. These tools offer significant advantages, such as emphasizing out-of-sample performance, detecting nonlinearities in the data and, in conjunction with the above, modeling complex relationships among the variables of interest and predictors. They use around 120 input variables for Belize and Salvador.

Bolivar (2024) demonstrates the effectiveness of machine learning models in predicting quarterly GDP for Bolivia, using various indicators, including satellite imagery data (the study examines the relationship between the Global Index of Economic Activity and the 261 variables collected). These approaches help overcome delays in official economic data releases. Tenorio & Pérez (2024) evaluate the predictive accuracy of various machine learning models for nowcasting Perus's monthly GDP growth rate from 2008 to 2023. The main findings are: superior performance of machine learning models which reduce forecast errors by 20%-25%. Including unstructured data (e.g., sentiment variables derived from Google Trends) enhances predictive accuracy, especially during high-volatility periods like the COVID-19 pandemic and its subsequent economic recovery. The results emphasize the potential of machine learning methodologies and unstructured data to improve macroeconomics forecasting, suggesting their wider application in emerging economies like Peru.

We add to the growing literature by central banks and other authors that have applied machine learning for nowcasting economic activity. We aim to assess and validate this approach through a case study of Chile. Incorporating conventional and non-traditional data sources has become increasingly essential for comprehensively understanding a country's economic activity. In this paper, we integrate conventional inputs alongside non-traditional sources, broadening our analytical perspective. Among these non-traditional sources, we highlight electronic payment systems, which offer real-time insights into consumption patterns and electronic invoices from Chilean businesses. These invoices, mandatory by law, systematically record commercial transactions, enhancing transparency and facilitating tax oversight in Chile. Additionally, we include temperature and precipitation data from Chile's major agricultural and industrial regions, recognizing their significant influence on productivity, supply chain stability, and overall economic output. By leveraging these diverse sources, we aim to capture more granular and dynamic signals of economic performance, offering a robust framework for more accurate and timely analysis in an increasingly complex economic landscape.

3 Artificial intelligence models

To begin with, it is essential to define machine learning and understand its core principles. Machine learning constitutes a subset of artificial intelligence (AI) that deals with the development of algorithms and models that allow computers to learn from past data, improving their performance automatically, without being explicitly programmed to carry out certain tasks[3]. The main goal of machine learning is to develop systems that can learn from past experiences to improve performance on specific tasks in the future (Hetland & Nelli, 2023).

Machine learning can be divided into several categories, including supervised learning (where the model is trained on labeled data), unsupervised learning (where the model tries to identify patterns without labeled data), and reinforcement learning (where the model learns through feedback from an environment), (Hetland & Nelli, 2023).

[3]Artificial intelligence includes not only machine learning, but also other types of models, such as expert systems and reinforcement learning systems.

The machine learning process generally consists of the following steps: data collection, data preprocessing, model choice, model training, model evaluation, (Hetland & Nelli, 2023). In this section we provide an overview of the select models or algorithms used. To build an activity nowcast we use different supervised ML algorithms: based either on penalized-regression models (e.g., Ridge regression, LASSO regression, elastic net regression) or on tree-based methods (e.g., XGBoost), besides a recurrent neural network model (e.g. 1D CNNs).

3.1 Penalized Regressions Models

Regularization is a technique used to prevent overfitting in machine learning models by adding a penalty to the loss function. This penalty discourages the model from assigning too much importance to any feature, which can lead to overfitting. The goal of regularization is to find a balance between fitting the training data well and generalizing it well to new data. There are different types of regularization, including L1 and L2 regularization, and elastic net, which is a combination of both (these are the most common types). L1 regularization, also known as LASSO regularization, adds a penalty equivalent to the L1 norm (or the absolute value) of the coefficients or weights in the cost function. L2 regularization, on the other hand, adds a penalty equivalent to the L2 norm (or the square of the absolute value) of the coefficients or weights in the cost function, (Kavanagh, 2024).

3.1.1 Ridge Regression

The Ridge model is defined by adding a penalty bases on the sum of squares of the coefficients of the predictor variables. This penalty compels the coefficients to be very small, preventing them from taking extremely high values, thus reducing the influence of less relevant variables. To estimate the coefficients , the equation must be expressed as:

$$\min_{\beta} \left[\sum_{i=1}^{n} (y_i - \beta_0 - \sum_{j=1}^{p} x_{ij}\beta_j)^2 + \lambda \sum_{j=1}^{p} \beta_j^2 \right] \tag{1}$$

Where y_i is the observed value of the dependent variable for observation i, x_{ij} is the value of predictor variable j in observation i, β_j is the coefficient associated with predictor variable j, p is the number of predictor variables, and λ is the regularization hyperparameter that controls the magnitude of the penalty. The sum of the terms β_j^2 in the penalty prevents the coefficients from reaching large values, thereby contributing to stability and reducing the risk of overfitting.

3.1.2 LASSO Regression

The LASSO (Least Absolute Shrinkage and Selection Operator) model, introduced by Tibshirani (1996), employs a penalty based on the sum of the absolute values of the coefficients of the predictor variables. This penalty has the property of forcing some coefficients to exactly reach zero, resulting in the automatic selection of a subset of more relevant predictor variables and the elimination of less significant ones. The LASSO coefficients $\hat{\beta}^{Lasso}$ are estimated:

$$\min_{\beta}\left[\sum_{i=1}^{n}(y_i-\beta_0-\sum_{j=1}^{p}x_{ij}\beta_j)^2+\lambda\sum_{j=1}^{p}|\beta_j|\right] \tag{2}$$

The change lies in the hyperparameter λ which, by summing the absolute values of the coefficients $|\beta_j|$ in the penalty, leads to model selection and simplification by allowing some coefficients to be zero. This provides a more precise variable selection approach regarding the degree of importance of all variables.

3.1.3 Elastic Net Regression

The Elastic Net model appropriately combines the constraints of both the LASSO and Ridge models. In particular, Zou & Hastie (2005) mention that its advantage lies in correcting the model when the number of regressors exceeds the number of observations ($p > n$), which improves variable grouping. The penalty includes both the sum of the absolute values of the coefficients and the sum of the squares of the coefficients of the predictor variables. The equation for estimating the coefficients $\hat{\beta}^{Enet}$ is expressed as:

$$\min_{\beta}\left[\sum_{i=1}^{n}(y_i-\beta_0-\sum_{j=1}^{p}x_{ij}\beta_j)^2+\lambda\sum_{j=1}^{p}(\alpha|\beta_j|+(1-\alpha)\beta_j^2)\right] \tag{3}$$

Where λ is the global regularization hyperparameter and α is the hyperparameter that controls the mix between LASSO ($\alpha = 1$) and Ridge ($\alpha = 0$) penalties. The combination of both penalties in the Elastic Net model allows for a higher degree of flexibility in variable selection and coefficient alignment.

3.2 Support vector machines (SVM)

3.2.1 Support vector regression (SVR)

The SVR introduced by Vapnik (1997) is a supervised learning algorithm that is borne from the methodology of SVM for regression problems. Unlike traditional linear regression, the goal of SVR is not to minimize the error between predicted and actual values directly, but rather to find a function that fits the data within a specified tolerance range, using the ϵ parameter. This implies that SVR tends to be robust concerning outliers and data with high dimensionality. Next, the objective function of SVR is presented, which is to minimize the loss function, where it includes a regularization term to control the complexity of the model and avoid overfitting:

$$L = \frac{1}{2}||w||^2 + C \sum_{i=1}^{n} (\xi_i + \xi_i^*)$$

subject to the constraints:

$$y_i - (w \cdot x_i + b) \leq \epsilon + \xi_i, \quad (w \cdot x_i + b) - y_i \leq \epsilon + \xi_i^*$$

where:
- w is the weight vector defining the hyperplane.
- b is the bias term.
- ξ_i, ξ_i^* are slack variables to allow for margin violations.
- C is a regularization parameter that balances the trade-off between the model complexity and the tolerance for deviations.

SVR uses a kernel function $K(x_i, x_j)$ to transform data into a larger dimensional space, allowing linear or nonlinear relationships to be captured. Common kernel functions include the linear kernel, the polynomial kernel, and the radial base function (RBF) kernel. It should be noted that under the linear and polynomial kernel it allows us to calculate coefficients of importance of variables through weights, but on the other hand, RBF kernel is not directly interpretable given the transformation of space. The model predicts the target variable $\hat{y}_i$ using the following function:

$$\hat{y}_i = \sum_{j=1}^{n} \alpha_j K(x_i, x_j) + b$$

where α_j are the Lagrange multipliers optimized during training.

SVR can handle nonlinear relationships and overfitting, therefore it is a tool to consider in regression analyses, to a greater extent when presenting high-dimensional datasets together with complex patterns. As indicated SVR can model non-linear relationships but it depends on the kernel used, in this study the linear kernel was the most effective, making SVR behave as a linear model in most cases. The effectiveness of SVR is illustrated a two-dimensional example presented in Appendix D, where the model showed superior performance after variable selection with LASSO.

3.3 Decision Tree Models

3.3.1 eXtreme Gradient Boosting (XGBoost)

The XGBoost algorithm (Chen & Guestrin, 2016) is similar to Ridge regression and LASSO, which add a penalty term to the loss function to prevent overfitting. The XGBoost algorithm uses a similar approach, but with a different loss function and optimization algorithm. The XGBoost algorithm uses a gradient boosting approach to improve the performance of the model. The algorithm starts with a weak initial model and iteratively improves it by adding new trees that correct the errors of the previous model [4]. It utilizes the function to make trees by minimization of:

$$\hat{y}_i^{(t)} = \sum_{k=1}^{t} f_k(x_i) \tag{4}$$

$$L^t = \sum_{i=1}^{n} l(y_i, \hat{y}_i^{(t)}) + \sum_{k=1}^{t} \Omega(f_k) \tag{5}$$

$$\Omega(f_k) = \gamma T + \frac{1}{2}\lambda \sum_{j=1}^{T} w_j^2 \tag{6}$$

Where:

- Each function f_k is a regression tree that corrects the residuals of previous trees
- T: number of leaves in the tree
- w_j: score (weight) of leaf j
- γ: penalty for the number of leaves
- λ: L2 regularization on leaf weights

3.4 Neural Network

In recent years, the integration of deep learning methodologies has been employed in the analysis and forecasting of time series, effectively modeling intricate and non-linear patterns. This integration has introduced novel tools that expand upon the conventional statistical methods. To learn about the concepts of Neural Networks to understand their formulation and evolution, we start with the perceptron, which was the first neural network proposed by Rosenblatt in 1958, where it was used to solve binary classification

[4] There are several ensemble methods. They mainly differ in the way they process the data. Ensemble methods differ in how they combine decision trees. Bagging builds independent trees from data subsets (Random Forest being a special case), while boosting trains trees sequentially to correct previous errors. Both enhance predictive accuracy but through distinct strategies (see Ciaburro (2024))

problems based on linear combinations of input and activation functions (Rosenblatt, 1958). This model had a limitation with issues that are not linearly separable, such as an XOR gate[5]. After implementing multilayer Neural Networks (MLPs) that incorporated hidden layers and nonlinear activation functions, they developed the ability to solve the previous problem, using the error backpropagation algorithm discovered in Rumelhart et al. (1986), which allows minimizing the cost function through the chain rule. Despite solving more complex problems even including nonconvex functions, they were not efficient when applied to sequential data with temporal dependencies. Recurrent Neural Networks (RNNs) solve the previous problem by incorporating cyclic connections in their architecture, retaining information from previous states, which renders them adequate in the modeling of temporal sequences, but they have the disadvantage that they do not capture the long-term temporal dependence due to the fading of the gradient and the burst Bengio et al. (1994), which negatively impacts the generation capacity of series that have a temporary structure.

One way to face this problem is with the use of one-dimensional Convolutional Neural Networks (1D-CNN), created as an alternative to the analysis of time series in the short term. But it should be noted that they originally arose to detect computer vision, where they extract patterns from historical data arranged as one-dimensional vectors, through the use of convolutional filters, detecting local patterns and structures implicitly.

The use of this model is appropriate when it comes to capturing conjunctural information or detecting recent patterns that influence the prediction of activity. It should be noted that the 1D-CNNs offer a simple and less computationally expensive architecture but maintain the high learning capacity on temporal signals. Thus, it makes them efficient in modeling time series data.

3.4.1 MLP (Multilayer Perceptron) Regression Model

The Multilayer Perceptron (MLP) is an artificial neural network in its most basic form, commonly used in regression and classification tasks, introduced by Werbos (1989). When used for regression, the MLP model seeks to map an input $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ to a continuous output $\hat{y}$. The basic architecture of an MLP consists of multiple layers: an input layer, one or more hidden layers, and an output layer.

The MLP model is mathematically defined as a function composed of linear transformations followed by nonlinear activation functions. The output of an MLP with a single hidden layer is:

$$\hat{y} = f(\mathbf{W}_2\sigma(\mathbf{W}_1\mathbf{x} + \mathbf{b}_1) + \mathbf{b}_2)$$

Where:

[5] An XOR (exclusive OR) gate is a digital logic gate that outputs true only when the inputs differ.

- $\mathbf{W}_1$ and $\mathbf{W}_2$ are arrays of weights.
- $\mathbf{b}_1$ and $\mathbf{b}_2$ are the bias vectors.
- $\sigma(\cdot)$ is an activation function, commonly the sigmoid function or ReLU.
- $\hat{y}$ is the predicted output.

It should be noted that for more complex models, with multiple hidden layers, the formulation expands to:

$$\hat{y} = f(\mathbf{W}_L\sigma(\mathbf{W}_{L-1}\ldots\sigma(\mathbf{W}_1\mathbf{x} + \mathbf{b}_1)\cdots + \mathbf{b}_{L-1}) + \mathbf{b}_L)$$

Where L is the number of hidden layers.

The optimization process in the training data of the MLP model is carried out by minimizing the loss function, in our case the mean square error (MSE) in for the regression, given by:

$$\mathcal{L}(\mathbf{W}_1, \mathbf{W}_2, \ldots, \mathbf{W}_L, \mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_L) = \frac{1}{N}\sum_{i=1}^{N}(\hat{y}_i - y_i)^2$$

Where N is the number of examples in the training set, $\hat{y}_i$ is the prediction for the i-th example, and y_i is the true output.

3.4.2 1D Convolutional Neural Network (1D CNNs)

The 1D CNNs, such as those introduced in LeCun et al. (1998), are used in one-dimensional signal processing tasks, including time series or sequences. In the context of regression, a 1D CNN performs an input sequence $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ to a continuous output $\hat{y}$. The model of a 1D CNN in regression applies convolutional filters to the input sequence, followed by a nonlinear activation function and a pooling layer. The output of a 1D CNN is calculated as follows:

$$\mathbf{h}_l = \sigma(\mathbf{W}_l * \mathbf{h}_{l-1} + \mathbf{b}_l)$$

Where:

- $*$ represents the convolution operation.
- $\mathbf{W}_l$ is the filter or weight of the l-th convolutional layer.
- $\mathbf{b}_l$ is the bias of the layer.
- σ is the nonlinear activation function, such as ReLU or tanh.
- $\mathbf{h}_l$ is the output of the l-th layer.

The output of the model, which is the prediction of the regression, is obtained as:

$$\hat{y} = \mathbf{W}_{out} \cdot \mathbf{h}_L + \mathbf{b}_{out}$$

Where L is the last layer of the network and $\mathbf{W}_{out}$ and $\mathbf{b}_{out}$ are the weights and bias of the output layer. The 1D CNN training process is performed by minimizing a loss function, as in the previous case using MSE. Where the loss function is defined as:

$$\mathcal{L}(\mathbf{W}_1, \mathbf{W}_2, \ldots, \mathbf{W}_L, \mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_L) = \frac{1}{N}\sum_{i=1}^{N}(\hat{y}_i - y_i)^2$$

Where N is the total observation points of the model, $\hat{y}_i$ is the prediction for the i-th example, and y_i is the true value.

4 Datasets

This section describes the characteristics of the data and the reasoning for arriving at the particular nowcasting framework.

4.1 Data

To perform the nowcast of the total and non-mining economic activity in Chile, we use monthly frequency predictor variables. The selected characteristics or predictors must capture the fundamental behaviors of the Chilean economy. This study distinguishes itself from previous research by incorporating traditional and non-traditional predictors. Among the traditional variables, we include indicators of key sectors such as mining, industry, and commerce, as well as variables from the fiscal, labor, monetary, financial, and international systems, in addition to consumption, expectations, and risks (both local and international), see **tables** A.1, A.2 and A.3. On the other hand, the non-traditional variables encompass information from invoices and electronic receipts that companies report to the Internal Revenue Service SII for commercial transactions. Another non-traditional source is the retail sales recorded on the main electronic payment platforms businesses use to process their transactions, which provides us with timely information on consumption patterns[6]. Finally, we consider climatic variables provided by the Meteorological Directorate of Chile (DMC), such as temperature and precipitation in the country. With all this, we obtain a comprehensive view of the dynamics of the economy, see **tables** A.4, A.5 for more details of variables.

[6]Central Bank of Chile's "Experimental Statistics" project. Available at: https://www.bcentral.cl/en/areas/estadisticas/estadisticas-experimentales

The dataset comprises 187 variables. Where applicable, we perform transformations such as seasonal adjustment and calculations of annual and monthly variations, resulting in a total of 400 variables for the period between January 2018 and December 2024.

Let us now examine the non-traditional variables in greater detail:

4.1.1 Electronic invoicing

In Chile, electronic invoices (facturas electrónicas in spanish) are digital tax documents that serve as legal proof of commercial business transactions. Mandated by the Chilean Internal Revenue Service (IRS) or Servicio de Impuestos Internos (SII) in Spanish, these invoices are generated, signed digitally, and transmitted electronically to the buyer and the tax authority. They are a key tool for fiscal transparency, enabling real-time monitoring of economic activities, reducing tax evasion, and simplifying tax compliance. Additionally, electronic invoices provide granular data on sales, purchases, and supply chain transactions, offering valuable insights for economic analysis and nowcasting/forecasting.

The sales data come from the electronic invoicing reported by firms to the IRS. Where the monthly panel dataset compiles information on Chilean firms from 2018 to July 2024, which allows us to compare the performance before, during and after the COVID-19 Pandemic.

Business trade credit also comes from electronic invoicing reported by firms. Incorporating inter-firm trade credit into a machine learning model for nowcasting economic activity is essential for capturing the financial dynamics of business operations and their impact on short-term economic fluctuations. Trade credit, often reflected in accounts receivable and payable, is a critical liquidity mechanism for firms, particularly during economic stress or uncertainty. In the Chilean context, where small and medium-sized enterprises (SMEs) play a significant economic role, trade credit can be an early indicator of financial health, supply chain stability, and business confidence. By integrating this variable into a machine learning framework, the model can leverage real-time transactional data to detect changes in credit flows, predict liquidity constraints, and anticipate potential disruptions in production and consumption cycles. This enhances the timeliness and accuracy of economic nowcasts, providing policymakers and analysts with a richer, data-driven perspective on the underlying trends driving short-term economic performance.

4.1.2 Electronic payment data

Transbank and Getnet provide the payment data that we use. For years, these companies have been the sole providers of transaction intermediation. In the last couple of years, others have appeared, but they are still the market's most representative companies. If a consumer pays by card in Chile, debit or credit, Transbank or Getnet will likely be managing the payment. The companies do not provide the raw data to the Bank for confidentiality reasons. The agreement with the provider is focused on acquiring data to evaluate macro issues, so it is consists only of daily sales disaggregated by the economic sector of the last link in the chain [7].

4.1.3 Climate variables

Incorporating climatic variables, such as monthly average temperature and precipitation, into a machine learning model for nowcasting economic activity in Chile is essential due to the country's strong reliance on climate-sensitive sectors like agriculture, mining, and energy production. Variations in temperature and rainfall directly influence crop yields, hydropower generation, and operational conditions in mining sites, creating significant short-term fluctuations in economic output. These variables can act as leading indicators, capturing real-time environmental shocks that traditional economic indicators may overlook or report with a delay. By integrating climate data, machine learning models gain access to non-traditional yet highly relevant predictors, enhancing their ability to detect subtle patterns, improve forecast accuracy, and provide more timely and reliable insights for policymakers and stakeholders navigating Chile's dynamic economic landscape.

4.1.4 Sectors

In the database, we include the International Standard Industrial Classification of all Economic Activities (ISIC) code for each firm. Through reference to the Economic Activity Code defined by EAC (CAE in Spanish) we obtain the economic sector to which each firm belongs. We exclude firms whose main sector of activity is related to mining as well as public administration, as the number of firms within each category does not comply with the requirement of statistical secrecy and the specific characteristics of these activities can bias the analysis.

[7] This information has been used in another work to study consumption: Cobb (2021)

4.1.5 Firm size

Chilean firms are classified by size bins according to their annual firm turnover in real terms, that is we convert the nominal amounts of Chilean Peso into Unidades de Fomento (UF). A firm is considered a micro firm if its turnover is below 2,400 UF, a small firm if turnover lies between 2,400 and 25,000 UF. They are categorized as medium sized firms when turnover lies between 25,000 and 100,000 UF. Finally they are considered large firms when turnover is above 100,000 UF. This classification is made according to what is currently established in the Statute of the Ministry of Economy of Chile (Law No. 20.416).

4.1.6 Frequency and variable transformations

The data frequency is monthly when constructing the model. Each variable was assessed regarding its predictive ability regarding monthly IMACEC growth. All monetary variables have been deflated and are expressed in constant terms. Sectoral deflators were used to deflate the variables. For example, the amounts in pesos for construction, commerce (wholesale, retail, and automotive), industry, and business services were deflated, occupying a specific indicator of each sector. When there is no deflator, the Unidades de Fomento (UF) is considered a deflator. The UF is a unit of account of the CPI index used in Chile. Where it makes sense, we also consider the seasonally adjusted series and the annual and monthly variations of the series as input to the models. Then, we transformed these variables into standardized ones to facilitate comparison and analysis. This standardization process allows us to maintain a common reference framework and ensure that different variables contribute equitably to the model.

Finally, including the transformations of the variables, we went from 187 to a total set of 400 predictors spanning from January 2018 to December 2024. However, literature on machine learning prediction suggests that feature reduction by filtering variables proving a positive and relatively strong correlation with the target variable enhance forecast performance, (see Kotsiantis, 2011, Bolivar, 2024). Therefore, a correlation analysis was conducted to examine the relationship between the IMACEC and the 400 variables collected (Annexes C.1 and C.2). The evaluation and selection of optimal predictors will be conducted independently for each machine learning algorithm employed. We will specify how we handle the data for the nowcasting update process in subsection 4.2 and how we test the model accuracy comparison in subsection 4.3. This approach will enable us to refine the process of choosing the most efficient prediction model, thereby achieving enhanced performance.

4.2 Recursive nowcasting

We use this data set to perform recursive nowcasts for Chilean IMACEC growth from 2024M1 to 2024M12, with the last nowcast made for 2024M12 using data up to November, 2024. For most methods, we standardize all indicators in each nowcasting round, ensuring each indicator has a mean of zero and unit variance in our estimation sample. The first nowcast is done for 2024M1, where we use data from 2018M1 to 2023M12 to train and test via cross-validation for optimal hyperparameters in our machine learning models, which we then use to nowcast out-of-sample 2024M1. In the next period, we use data from 2018M1 to 2024M1 to train and test for the optimal model, and subsequently, use it to nowcast 2024M2, and so on.

More precisely, the database used is divided into different blocks (folds), and a subset of these blocks is chosen as the testing group, while the remaining blocks are used for training the model. The portion of the data reserved for cross-validation consists of both training and testing. Unlike traditional cross-validation, where data is randomly divided, we use the Anchored Walk-Forward Validation method, a cross-validation technique for time series in which the training window progressively expands without removing old data, while the test window moves forward[8]. The advantage of this approach is that it considers the time series nature of the data and provides a more realistic estimate of the model's performance. This is crucial in order to avoid the issue of data leakage, where future information influences the model's training.

For example, in the first fold, training is performed with data up to month t, and testing is conducted with months t+1, t+2, t+3, and t+4. Then, the training and testing sets are expanded to repeat the process. With the model calibrated using the training blocks, predictions are obtained and compared with the testing data to calculate the prediction error. This procedure is repeated iteratively with all training and testing folds, keeping a record of the prediction errors. In this way, the model and hyperparameters that minimize the prediction error across all possible combinations are selected. This results in a prediction model that optimizes out-of-sample performance. This methodology is similar to that used by Barrios et al. (2021), where a separate validation set is not incorporated, as walk-forward validation acts as a validation mechanism by continuously evaluating the model on data not previously seen by the model. By using only training and testing in a temporal sequence, data integrity is maintained, and predictive performance is optimized

[8]Unlike Rolling Window validation, where older data is discarded, Anchored Walk-Forward Validation ensures that training always starts from a fixed point (anchored) and increases with each new iteration. This approach allows us to assess how the model improves as more historical data becomes available and is particularly useful when the relationships within the data are assumed to be relatively stable over time.

in a real-time updated data environment. Additionally, because data is often limited, splitting it into more sets would further reduce the amount of information available for training the model.

Figure 1 illustrates the recursive nowcasting process. In this validation scheme, we incrementally expand or slide the training window by one or more folds at a time:

1. Initial Split: we start by dividing the dataset into k folds. Each fold consists of a contiguous sequence of observations.

2. Training and Testing: In each iteration, one fold is designated as the test set, and the preceding folds are used for training. This simulates the model's ability to generalize to unseen future data.

3. Iterative Process: The process is repeated k times, each time moving the test set to the next contiguous sequence of observations. This ensures that each data point is used for both training and testing, preventing data leakage and providing a realistic evaluation of model performance.

Figure 1: Recursive nowcasting

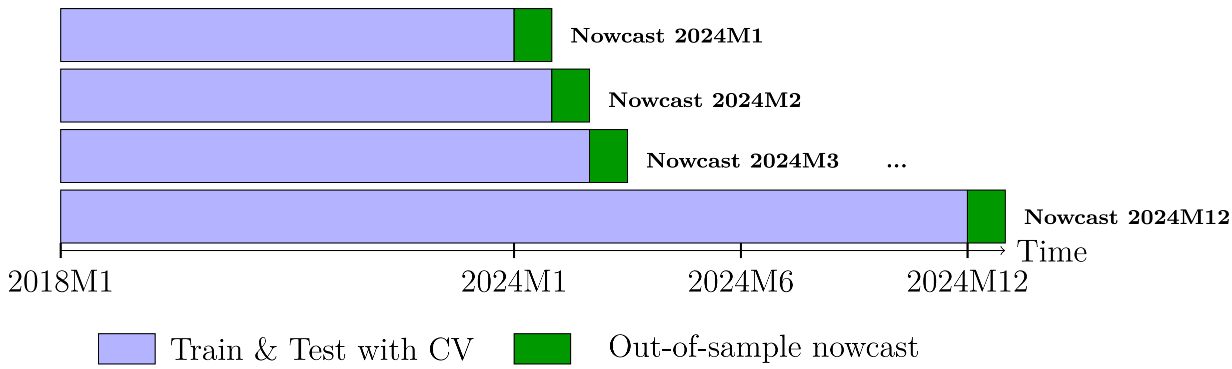

Notes: Graph illustrates recursive nowcasting. The green area called OOS stands for Out-Of-Sample nowcasting. The first OOS nowcast is made for 2024M1, whereas the last one is done for 2024M12. The purple area called "Train & Test with CV" is used to select the optimal tuning of the model used for OOS nowcasting. The model is chosen via walk-forward validation with an expanding window: it illustrates the expanding window variant of the walk-forward validation, which is also called anchored walk-forward validation. As you can see, we are incrementally increasing the size of the training and test set.
Sources: Researcher's own calculations.

In this paper, we use Anchored Walk-Forward Validation with 4-fold cross-validation. Additionally, in each iteration of the walk-forward process, we employ the Grid Search optimization technique to find the best hyperparameters for each model, ensuring optimal tuning without compromising the temporal integrity of the data. Table 1 provides an overview of the model-specific hyperparameters that are described in **section 3**. The table summarizes the hyperparameters tuning and their corresponding ranges from which we search for the optimal value via 4-fold walk-forward validation.

Table 1: Hyperparameters Overview

Model	Hyperparameters	Range
LASSO	λ	0.0005-1
Ridge	λ	0.01-10
Elastic net	λ	0.01 - 10
	α	0.01-0.99
XGBOOST	learning_rate	0.01 - 0.10
	n_estimator	100 - 300
	max_depth	2 - 4
	subsample	0.5 - 1.00
	min_child_weight	4 - 12
	colsample_bytree	0.34 - 0.75
	reg_alpha	0.25 - 0.50
	reg_lambda	0.25 - 0.50
SVR	svr___c	0.1-100
	svr___epsilon	0.01-0.2
	svr___kernel	lineal-rbf
	learning_rate	0.001-0.01
MLP	mlp___hidden_layer_sizes	50-100-200
	mlp___activation	relu-tanh
	mlp___solver	adam-lbfgs-sgd
	mlp___aplha	10^{-5}-10^{3}
	mlp___learning_rate	constant-adaptive
1D CNN	seq_lenght	12-24
	filters	128-64
	kernel_size	6-12
	units	64-32

Note: table 1 reports the hyperparameters (column 2) of various models that are estimated with four cross-validation from ranges reported in column 3.

4.3 Nowcast evaluation methodology

We evaluate model performance through an out-of-sample nowcasting exercise. For each target month—from 2024M1 to 2024M12—we train the model using an expanding window starting from 2018M1 up to the month immediately before the nowcast target.

Within each training window, we select the optimal hyperparameters using time-series cross-validation (as discussed in the previous **section 4.2**). Once selected, the model is re-estimated on the full training sample with these hyperparameters, and a nowcast is generated for the next month. For example, to nowcast 2024M1, we perform cross-validation within 2018M1–2023M12 to tune the model, then retrain it on the full period and generate the forecast for 2024M1.

Prediction accuracy is evaluated by comparing the model's nowcast with the observed value using the Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and Mean Absolute Error (MAE).

$$MSE = \frac{\sum_{t=1}^{T}(y_t - \hat{y}_t)^2}{T} \tag{7}$$

$$RMSE = \sqrt{\frac{\sum_{t=1}^{T}(y_t - \hat{y}_t)^2}{T}} \tag{8}$$

$$MAE = \frac{1}{T}\sum_{t=1}^{T}|y_t - \hat{y}_t|$$

Where:

- y_t: Actual value at time t.
- $\hat{y}_t$: Predicted value by the model at time t.
- T: Total number of observations.

5 Results

In this section we present the out-of-sample performance of the models from section 3.

5.1 Fine-Tuning hyperparameters

We focus exclusively on results after hyperparameter tuning because this process significantly improved the performance of all algorithms. Hyperparameter tuning mitigated overfitting by optimizing the models' parameters and enhancing their generalization ability to unseen data. Presenting results without this tuning could be misleading, as it would not reflect the full potential of each model. For example, in the context of nowcasting economic activity, studies such as Richardson et al. (2021) for New Zealand GDP and Tiffin (2016) for Lebanon demonstrate that hyperparameter tuning can improve model performance by reducing forecast errors and enhancing predictive accuracy. These findings highlight the critical role of tuning in machine learning applications for regression-based nowcasting, ensuring models are better aligned with the data and capable of capturing underlying nonlinear patterns. Showcasing results after hyperparameter tuning offers a more accurate and reliable assessment of the models' comparative performance.Thus, showcasing results post-tuning provides a more accurate and reliable evaluation of the models' comparative performance.

In Table 2, we present the out-of-sample performance metrics for each model, covering the period from January to December 2024. The first column lists the algorithm or model used, while the next three columns display the MSE, RMSE, and MAE for the Total IMACEC. The last three columns report the same metrics for the Non-mining IMACEC. The MSE and RMSE are particularly useful when large errors should be penalized more heavily, whereas the MAE offers a more robust assessment in the presence of outliers.

It is important to highlight that these results are obtained using the revised series of both the Total IMACEC and the Non-mining IMACEC. IMACEC revisions are adjustments made to preliminary figures to reflect more accurate and complete information, consistent with the updates of quarterly and annual national accounts. Initially, IMACEC is published with preliminary data that may later change as more detailed and complete information becomes available. Revised series, including both original and seasonally adjusted data, are then published to ensure accuracy and reliability. Revisions are thus essential for maintaining the credibility of economic statistics, allowing policymakers and analysts to make informed decisions based on updated information. For completeness, in a subsequent section we also present model results using the non-revised series, which allows us to evaluate the robustness of our findings across different vintages of the same indicator. It is also worth noting that if a model appears in the table with a specification in parentheses (e.g., SVR (LASSO)), this indicates that the model was estimated with that specific type of regularization.

For the Total IMACEC, after fine-tuning hyperparameters, XGBoost consistently achieves the best predictive performance, delivering the lowest MSE (1.91), RMSE (1.38), and MAE (1.05) among all evaluated models. This highlights the benefits of sequential gradient-boosted decision trees in capturing complex nonlinear relationships and heterogeneous effects present in high-frequency macroeconomic data. LASSO and SVR (LASSO) also perform competitively, with SVR achieving slightly better results than LASSO in terms of MSE and RMSE. Ridge and Elastic Net regressions show substantially higher errors, while deep learning models — MLP and 1D CNN — rank in the middle, outperforming Ridge but not surpassing XGBoost or SVR. These findings suggest that ensemble-based and kernel-based approaches are particularly effective for short-term nowcasting of aggregate economic activity.

For the Non-mining IMACEC, the performance ranking changes. SVR (LASSO) emerges as the top performer, achieving the lowest MSE (0.97), RMSE (0.98), and MAE (0.88). LASSO follows closely, with strong predictive accuracy across all metrics. XGBoost also delivers robust results, matching LASSO in MSE (1.65) and RMSE (1.28) and achieving a competitive MAE (1.00). This shift in ranking suggests that kernel-based models may be more effective in this disaggregated setting, potentially due to fewer nonlinear interactions or greater stability in the non-mining series. Deep learning models again exhibit larger errors, particularly the 1D CNN, which performs worst in both IMACEC variants. Ridge regression remains the weakest performer among penalized linear methods.

Overall, the results consistently show that nonlinear models outperform linear penalized regressions in nowcasting Chilean economic activity, particularly when revisions are incorporated into the training process. XGBoost's iterative boosting, regularization, and hyperparameter optimization allow it to model subtle short-term dynamics while mitigating overfitting. Similarly, among linear approaches, SVR with a linear kernel effectively captures relationships in high-dimensional settings through regularization and margin-based optimization. While linear models such as LASSO, Ridge, and Elastic Net provide interpretability and efficient variable selection, they are inherently less capable of capturing the complex, nonlinear interactions present in heterogeneous economic data. Ensuring appropriate regularization, cross-validation, and feature scaling remains crucial across all model classes to maintain stability and predictive accuracy in real-time macroeconomic nowcasting. It is also worth noting that the relatively weaker performance of deep-learning models such as MLP and 1D CNN may be explained by their high data requirements. These architectures typically benefit from long time series and large samples that allow them to capture complex temporal dependencies.

Finally, the **figures 2 and 3** compare linear and nonlinear machine learning models used for nowcasting the annual growth rate of the Monthly Economic Activity Indicator in Chile. Panels (a) show the results of linear models such as LASSO, Ridge, Elastic Net and SVR for total activity and non-mining activity, respectively, contrasted with the observed values (black line). Panels (b) present nonlinear models applied to the same variables, including MLP and XGBoost. Overall, nonlinear models tend to capture volatility and extreme variations more accurately than linear models, which generally smooth fluctuations. This suggests incorporating nonlinear relationships improves predictive performance in a highly volatile economic environment.

Table 2: Fine-Tuned Out-of-Sample Performance Metrics of ML Models for IMACEC
Results with IMACEC Revisions, 2024m01–2024m12

Algorithm	**Total IMACEC**			**Non-mining IMACEC**		
	MSE	RMSE	MAE	MSE	RMSE	MAE
LASSO	2.16	1.47	1.19	**1.10**	**1.05**	**0.85**
Ridge	3.78	1.94	1.72	3.93	1.98	1.65
Elastic net	3.78	1.94	1.71	3.78	1.94	1.56
SVR (LASSO)	2.30	1.52	1.18	**0.97**	**0.98**	**0.88**
MLP (Total/LASSO)	3.97	1.99	1.74	**1.65**	**1.28**	**1.10**
1D CNN (LASSO)	4.51	2.12	1.48	5.43	2.33	2.07
XGBoost	**1.91**	**1.38**	**1.05**	**1.65**	**1.28**	**1.00**

Note: Period covered 2024m01–2024m12. Machine learning models were trained using 4-fold cross-validation and include IMACEC revisions. **SVR (LASSO)** uses Lasso regularization for both total and non-mining IMACEC. **MLP (Total/LASSO)** applies Lasso only to the non-mining IMACEC. **1D-CNN (LASSO)** applies Lasso to both variants. Among the nonlinear models, **XGBoost** shows the best overall performance, benefiting from sequential tree boosting. **SVR (LASSO)** follows closely, effectively capturing patterns depending on the kernel used, in this study the linear kernel was most effective in high-dimensional settings. Red boxes indicate models with all error metrics below 2.0. *Source:* Own elaboration.

5.2 ML Models vs. Traditional Nowcast Benchmarks

To complement the previous results obtained with revised IMACEC series, this section evaluates model performance using real-time vintages—the unrevised figures actually published each month. This contrast matters because revised data provide a cleaner view of underlying dynamics, whereas policy and market decisions must be taken with preliminary information. Comparing both perspectives helps assess the robustness and real-time usefulness of the models.

Table 3 reports out-of-sample MSE, RMSE, and MAE for Total and Non-mining IMACEC during 2024, computed on vintages (no subsequent revisions). The row labeled Traditional/Econometric Models corresponds to the benchmark nowcast, constructed as the RMSE-weighted composite of bridge models, SARIMA, BVAR, MIDAS, and dynamic factor models, following Cobb & Peña (2020). Rather than relying on a single specification, this weighted approach acknowledges that the relative performance of individual models can shift abruptly with changing economic conditions. Combining models therefore helps stabilize results and systematically reduces nowcast errors compared to any single econometric model, in line with the recommendations in the literature. Importantly, these benchmarks rely exclusively on standard macroeconomic aggregates (such as national accounts, sectoral indicators, prices, and financial variables) and do

Figure 2: YoY Growth of Total IMACEC: Observed vs Nowcast

(a) Linear models

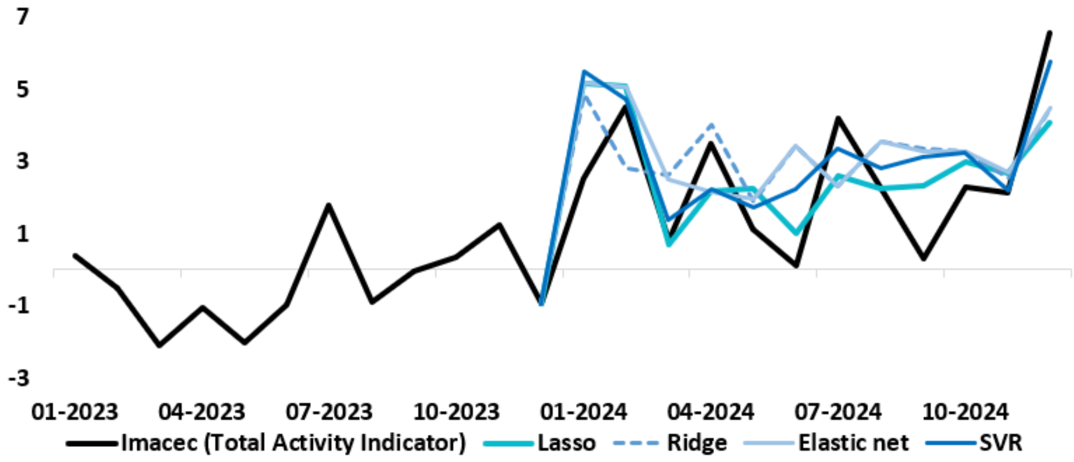

(b) Nonlinear models

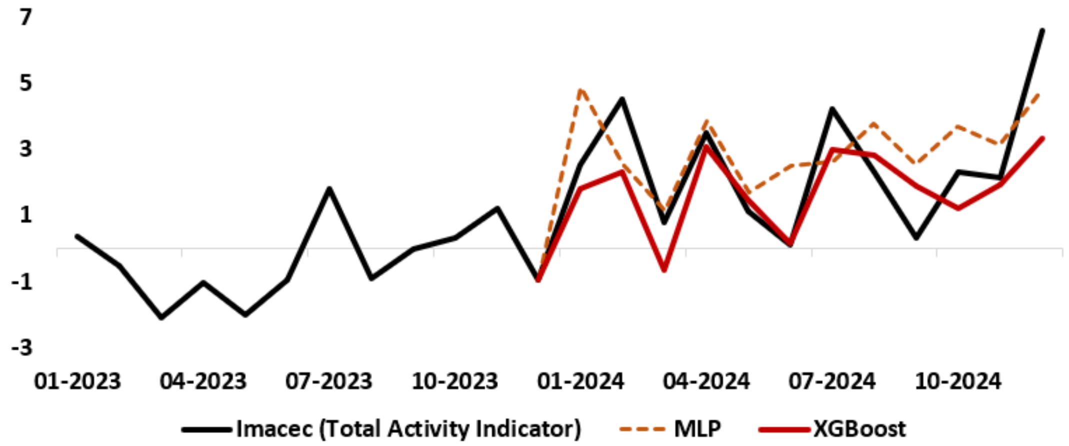

Notes: In panel (a), we show the observed annual growth rate of the total activity indicator (black line) and the nowcast of each linear model for each month: LASSO, Ridge, Elastic Net and SVR. In panel (b), we show the observed annual growth rate of the total activity indicator (black line) and the nowcast of each nonlinear model for each month: MLP and XGBoost.
Sources: Own calculations.

not incorporate experimental statistics or microdata (e.g., firm-level sales) discussed in Section 4, which makes their information set more restricted than in our AI specifications.

Overall, XGBoost achieves the best vintage-based performance for the Total IMACEC, posting the lowest MSE, RMSE, and MAE (MSE=1.98, RMSE=1.41; MAE=1.05). SVR (LASSO) and LASSO follow with competitive errors, while Ridge and Elastic Net are clearly weaker. Deep learning models (MLP) show intermediate or lower accuracy, consistent with their higher data requirements and the short time span available for monthly nowcasting. Relative to the RMSE-weighted traditional composite from Cobb & Peña (2020), XGBoost delivers lower errors, indicating that non-linear ML

Figure 3: YoY Growth of Non-mining IMACEC: Observed vs Nowcast

(a) Linear models

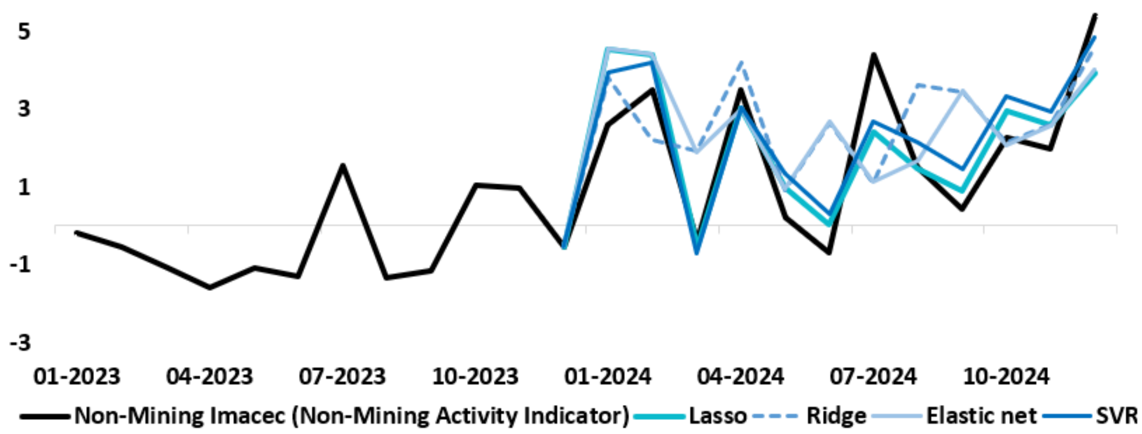

(b) Nonlinear models

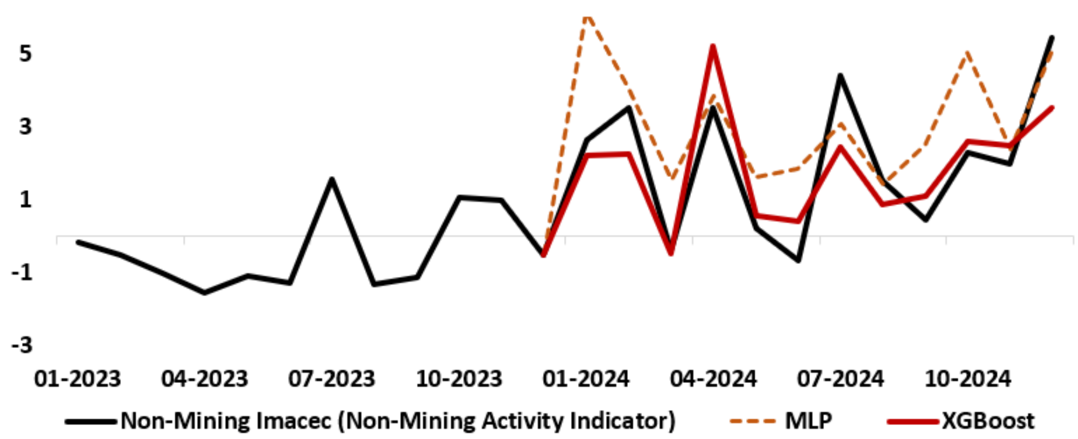

Notes: In panel (a), we show the observed annual growth rate of the non-mining activity indicator (black line) and the nowcast of each linear model for each month: LASSO, Ridge, Elastic Net and SVR. In panel (b), we show the observed annual growth rate of the non-mining activity indicator (black line) and the nowcast of each nonlinear model for each month: MLP and XGBoost.
Sources: Own calculations.

can better capture short-run nonlinearities and interactions in real-time conditions.

For the Non-mining IMACEC, the ranking is similar. XGBoost again performs strongly (MSE=1.65; RMSE=1.28; MAE=1.00), closely matched by SVR (LASSO) and LASSO, which post very low MAE (down to 0.91–0.97). Deep learning models remain behind—particularly 1D-CNN—reinforcing the proposition that, with limited monthly history and noisy high-frequency signals, ensemble and kernel methods are more reliable than deep networks. Against the RMSE-weighted econometric composite [9], ML

[9]The projections correspond to the RMSE-weighted average of the models considered in Cobb & Peña (2020), since this approach assigns greater weight to those with lower forecast errors. This procedure

models—especially XGBoost and SVR (LASSO)—consistently outperform in real-time vintages.

Finally, note that ML “lets the data speak” without imposing rigid functional forms. Among ML methods, XGBoost stands out thanks to sequential gradient boosting with regularization and tuned hyperparameters, which captures complex nonlinearities and heterogeneous effects while controlling overfitting. This methodological flexibility, combined with broader information sets than those used in traditional econometrics, explains its robust real-time performance.

The results show that although machine learning models consistently outperform traditional econometric methods, their predictive power improves substantially when microdata and experimental statistics—such as electronic invoicing records—are incorporated. Without these inputs (Table E.7), prediction errors remain higher and the gains over traditional models are more modest. This contrast highlights the central role of high-frequency administrative data in unlocking the full potential of ML methods for macroeconomic nowcasting.

5.3 Statistical Robustness

The Diebold–Mariano test (Diebold & Mariano, 1995) is a statistical procedure designed to compare the predictive accuracy of two competing forecasting models. The null hypothesis states that both models have equal predictive performance. The test evaluates whether the differences in the models’ forecast errors are statistically significant, based on a chosen loss function—in this case, the squared error. In our robustness check, we use this test to formally determine whether the predictive accuracy of the machine learning models differs from that of the XGBoost benchmark, for both the Total and the Non-mining IMACEC.

XGBoost was selected as the benchmark model due to its outstanding performance in predicting both the Total IMACEC and the Non-mining IMACEC. In the evaluation of out-of-sample predictive errors (MSE, RMSE, and MAE), XGBoost consistently achieved the lowest error values among all tested algorithms. This superior performance can be attributed to its gradient boosting framework, which effectively captures complex, non-linear relationships and interactions among predictors, while controlling overfitting through regularization. Given these results, it is a natural choice to assess whether other machine learning models can statistically outperform this strong baseline.

establishes a benchmark within traditional approaches by emphasizing models with stronger historical performance. Moreover, this way of combining models is standard in the literature, as noted by the authors themselves.

Table 3: Fine-Tuned Out-of-Sample Performance Metrics of ML Models for IMACEC
Vintages only, without IMACEC Revisions, 2024m01–2024m12

Algorithm	**Total IMACEC**			**Non-mining IMACEC**		
	MSE	RMSE	MAE	MSE	RMSE	MAE
LASSO	2.35	1.53	1.26	**1.13**	**1.06**	**0.91**
Ridge	3.88	1.97	1.71	4.17	2.04	1.65
Elastic net	4.02	2.00	1.78	3.99	2.00	1.63
SVR (LASSO)	2.49	1.58	1.24	**1.12**	**1.06**	**0.97**
MLP (Total/LASSO)	4.10	2.02	1.73	**1.72**	**1.31**	**1.11**
1D CNN (LASSO)	4.66	2.16	1.54	6.34	2.52	2.24
XGBoost	**1.98**	**1.41**	**1.09**	**1.65**	**1.28**	**1.00**
Traditional/Econometrics Models	3.08	1.75	1.48	3.08	1.76	1.53

Note: Period covered: 2024m01–2024m12.Results from traditional models are constructed as an RMSE-weighted composite of bridge models, SARIMA, BVAR, MIDAS, and dynamic factor models, consistent with the approach proposed by Cobb and Peña (2020). Machine learning models were trained using 4-fold cross-validation and do not-include IMACEC revisions. **SVR (LASSO)** uses Lasso regularization for both total and non-mining IMACEC. **MLP (Total/LASSO)** applies Lasso only to the non-mining IMACEC. **1D-CNN (LASSO)** applies Lasso to both variants. Among the nonlinear models, **XGBoost** shows the best overall performance, benefiting from sequential tree boosting. **SVR (LASSO)** follows closely, although SVR can capture nonlinear patterns depending on the kernel used, in this study the linear kernel was most effective. Red boxes indicate models with all error metrics below 2.0. *Source:* Own elaboration.

Table **table 4** reports the Diebold–Mariano test results for both the total and non-mining IMACEC. For the Total IMACEC, the Diebold–Mariano test results indicate that Ridge (p = 0.0151), Elastic Net (p = 0.0247), MLP (p = 0.0287), and 1D CNN (p = 0.0338) present p-values below conventional significance thresholds (e.g., 5%). This means that, for these models, the difference in predictive accuracy relative to XGBoost is statistically significant. In all these cases, the sign of the DM statistic shows that XGBoost's forecast errors are smaller, confirming its superiority over these approaches in the Total IMACEC nowcasting task. For LASSO and SVR, the higher p-values suggest no statistically significant difference in performance.

In the case of the Non-mining IMACEC, the results are more mixed. Ridge (p = 0.0651) and Elastic Net (p = 0.0756) show p-values close to the 10% significance level, suggesting moderate evidence that XGBoost performs better. The 1D CNN, however, presents a much lower p-value (p = 0.0119), strongly supporting XGBoost's superiority for this sub-index. In contrast, LASSO, SVR, and MLP yield high p-values, indicating that their predictive performance is statistically similar to XGBoost for the Non-mining

IMACEC. These findings suggest that while XGBoost maintains a strong performance overall, its relative advantage varies depending on the sector-specific measure being forecasted.

Table 4: Comparing Predictive Accuracy of Two Nowcasts: The Diebold-Mariano Test

Algorithm	Total IMACEC	Non-mining IMACEC
LASSO	0.1777	0.7943
Ridge	0.0151	0.0651
Elastic net	0.0247	0.0756
SVR	0.2169	0.8081
MLP	0.0287	0.1552
1D CNN	0.0338	0.0119
Traditional/Econometric Models	0.0637	0.0800

Note: The Diebold–Mariano test (DM) (Diebold & Mariano, 1995) evaluates whether the predictive accuracy of two competing models is statistically different. The null hypothesis states that both models have equal predictive performance. A low p-value (e.g., < 0.05) indicates that the difference in forecast errors is statistically significant. In this table, a positive DM statistic indicates that the benchmark model (XGBoost) has lower forecast errors than the comparison model. *Source:* Own elaboration.

5.4 Feature importance

Despite having identified models with higher predictive capabilities for IMACEC growth, understanding the significance of individual predictors (features) is crucial. In Ridge, LASSO, and ElasticNet, coefficients in the feature importance constitute the influence magnitude and direction of features on the target variable. Positive coefficients signify a direct relationship, while negative coefficients indicate an inverse relationship. For other algorithms, the figure 5 reflects the relative importance of each feature, with individual weights summing up to 1. Since feature importance has different interpretations in each model, we only show the relative importance of each feature in each model to make it comparable. Figures 4 and 5 have the heatmaps with the importance of each characteristic in each algorithm to predict the total and non-mining IMACEC, respectively. Each feature has an individual weight ranging from 0 to 1. If the weight is 0, the characteristic has little influence on the model; however, if the weight is 1, it is a relevant variable in the algorithm.

In the case of the total IMACEC, the observations reflect the importance of business sales and trade credit between firms from the internal revenue service. Regarding sales, those in the commerce, and personal services sectors, as well as those in the central, northern, southern and southern macro-zones of the country stand out. The following variables are also relevant to the model for total activity: energy, truck sales, sales of non-durable goods, and consumer imports. For their part, within the employment variables, the dynamics of employment in the trade and construction sector stand out. On the other hand, the dynamics of mining production, manufacturing production, and copper prices are also considered relevant in most models. Another variable that also results from its potential to predict activity is the local risk indicator.

The INE prepares the Mining, Manufacturing Production, and Trade Activity indicators, which aim to measure the real evolution of those activities. Mining production indicates the physical production of a basket of products made by establishments whose main activity is mining. The trade index reflects the total sales at constant prices of companies whose main activity is trade. The manufacturing production indicator reflects physical production, manufacturing sales, and the number of person-hours employed in the production processes of the different establishments. The information is captured through surveys targeting at these establishments and is constructed on a fixed basis, using increasingly outdated weightings as the years pass.

In the case of non-mining activity, although the relevance of each variable depends on the model, the sales of companies reported by the Internal Revenue Service stand out once again. These include sales of personal services, the southern region, the central macrozone, and sales in commerce and industry. Also relevant are the index of non-durable goods sales, supermarket sales, truck sales, energy consumption, and the manufacturing production index.

From an economic perspective, these variables are crucial because they reflect the behavior of consumption and production in different sectors of the economy. For example, sales of non-durable goods and supermarket goods indicate household spending on everyday products, while truck sales and energy consumption are indicators of industrial and logistical activity.

On the foreign trade side, imports of consumer goods and intermediate goods, as well as consumer credit, stand out. Imports of consumer goods reflect domestic demand for foreign products. In contrast, imports of intermediate goods are essential for national production, as these goods are used as inputs to manufacture other products. Consumer credit, in turn, indicates households' confidence level and spending capacity.

Employment in the construction and commerce sectors is important in the labor market. Employment in construction is a key indicator of infrastructure and urban development investment. In contrast, employment in commerce reflects the dynamics of the service sector and its capacity to generate jobs.

Figure 4: Feature importance by algorithm for 45 variables- Total IMACEC

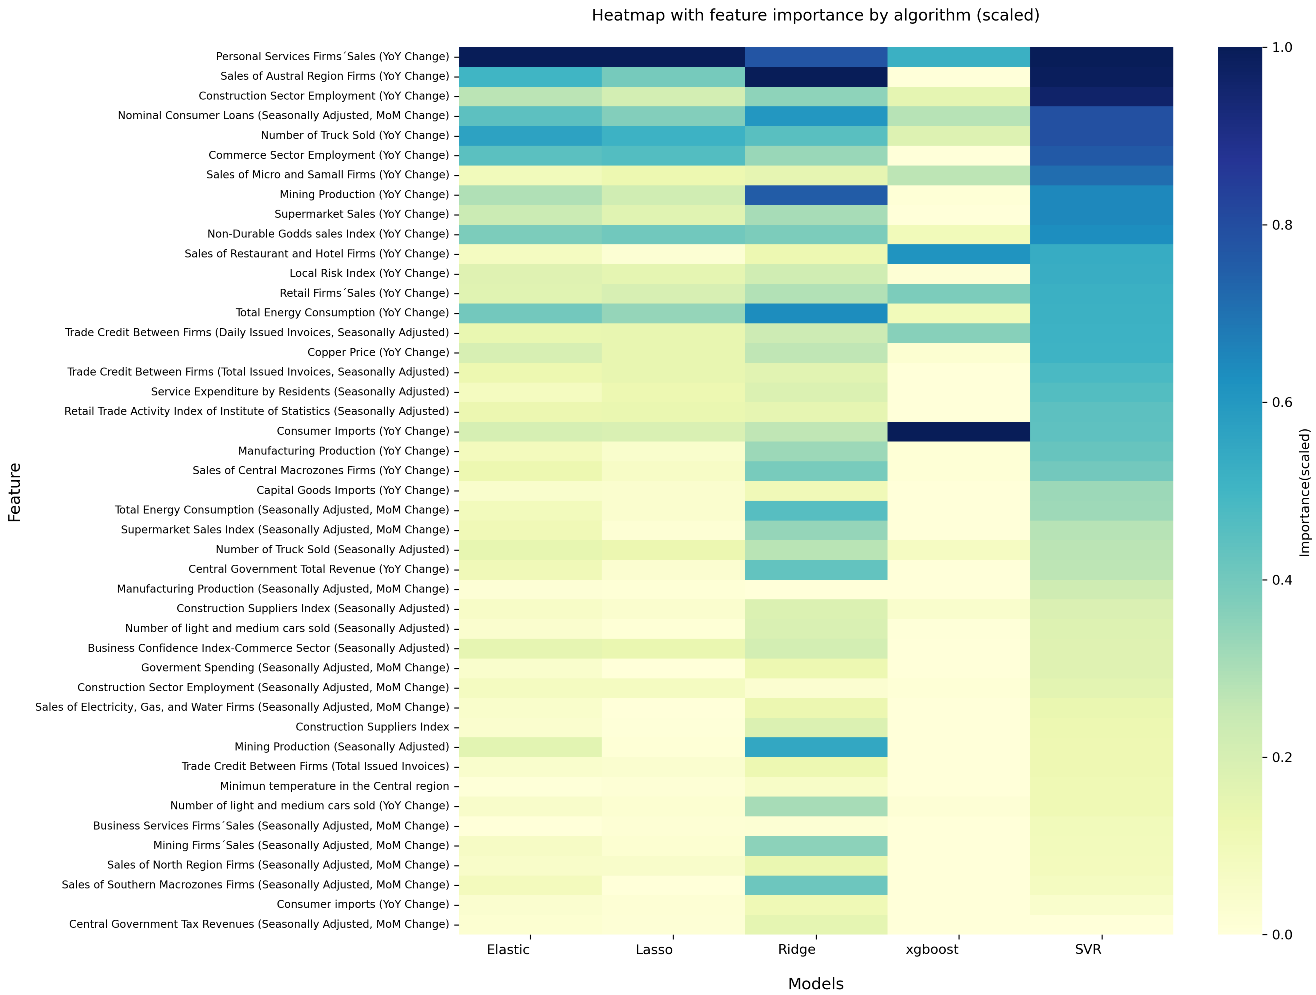

Notes: We considered the average score of each feature between January 2024 and August 2024, in each model. We plot the weights assigned to predictors in each algorithm, scaled between 0 and 1 for common variables. YoY (Year-over-Year) refers to the comparison with the same period in the previous year in order to identify long-term trends. MoM (Month-over-Month) refers to the comparison with the previous month, highlighting short-term changes or fluctuations.
Sources: Own calculations.

5.5 Shapley Values

One way to see the important variables that favor prediction is the use of SHAP (SHapley Additive exPlanation) introduced for Lundberg & Lee (2017), the method is based on game theory and allows explaining predictions in complex models, especially when using Deep Learning models where the delivery of the contribution of the explanatory variables is not as intuitive as in classic time series models.

Figure 5: Feature importance by algorithm for 44 variables- non-mining IMACEC

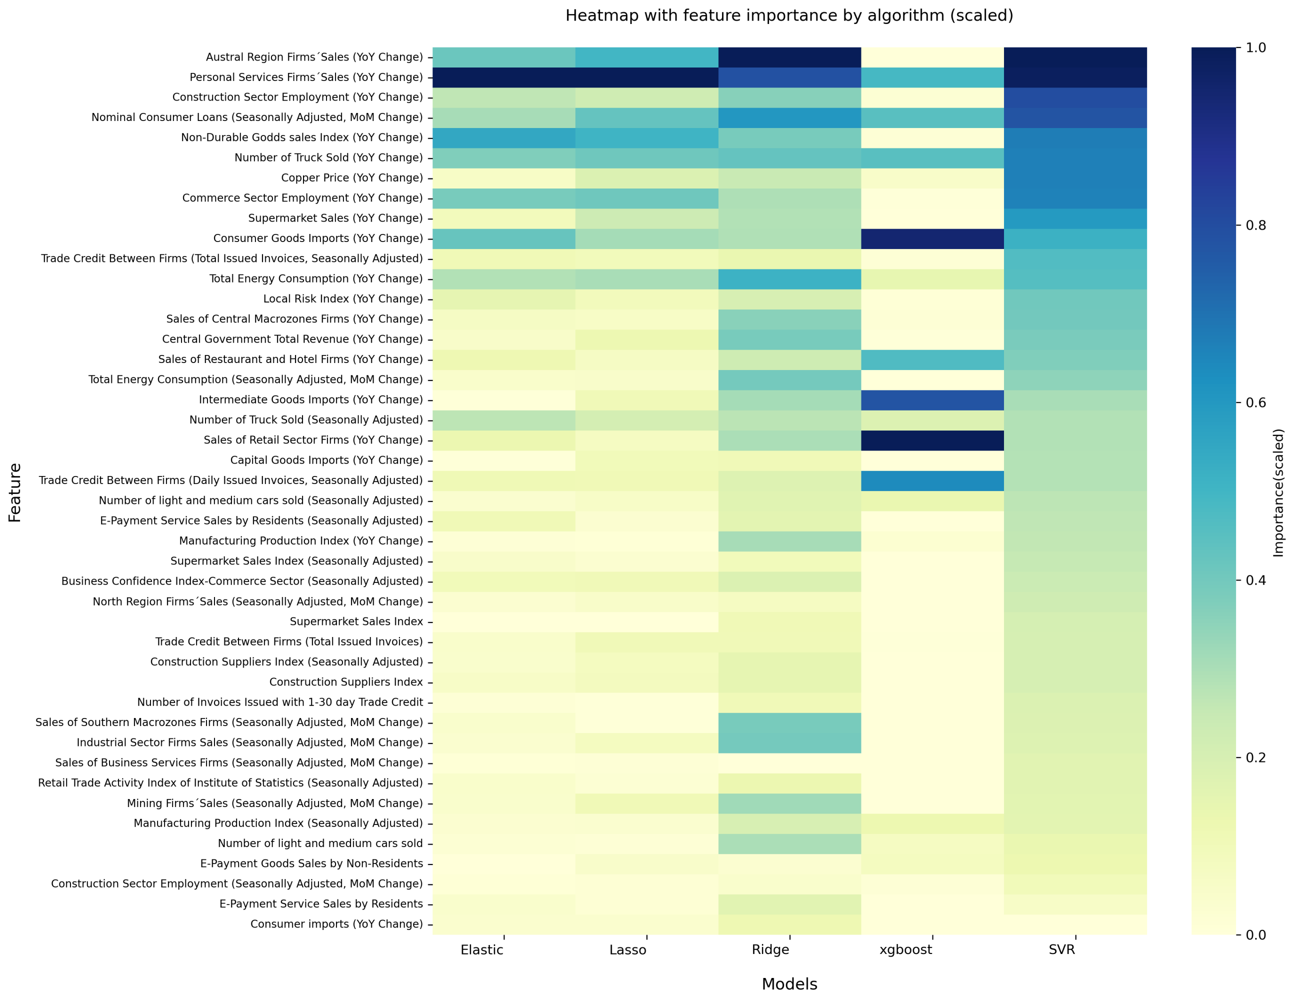

Notes: We considered the average score of each feature between January 2024 and August 2024, in each model. We plot the weights assigned to predictors in each algorithm, scaled between 0 and 1 for common variables.
Sources: Own calculations.

Shapley values are calculated as follows:

$$\phi_i = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!\,(|N|-|S|-1)!}{|N|!} \left[f(S \cup \{i\}) - f(S) \right]$$

Where:

- ϕ_i: SHAP value for the characteristic i.
- S: A subset of features without i.
- N: A complete set of features.
- $f(S)$: Prediction of the model using the features in S.

It can be noted that SHAP performs an evaluation of all combinations of characteristics, where it measures the change in prediction by including a new characteristic and then averages the weighted change. This exercise results in the average contribution of characteristics in the prediction of the model.

SHAP was also used in the one-dimensional convolutional neural network (1D CNN) for the same purpose of seeing how the features act at each instant of time. In this case an explainer optimized for neural networks based on TensorFlow is used, which acts on the trained convolutional model, and the training data is used as a reference to calculate the marginal contributions. SHAP approximates Shapley's values through network activations, capturing how each element of the time series influences the output, which allows us to analyze which moments in the entry window have greater relevance in the model's prediction, allowing the interpretation of temporal patterns. To simplify the analysis, a one-dimensional vector is calculated by first calculating a global importance and then obtaining the average importance of each characteristic without losing the relevant information.

Figure 6: Feature Importance by Algorithm - Total IMACEC

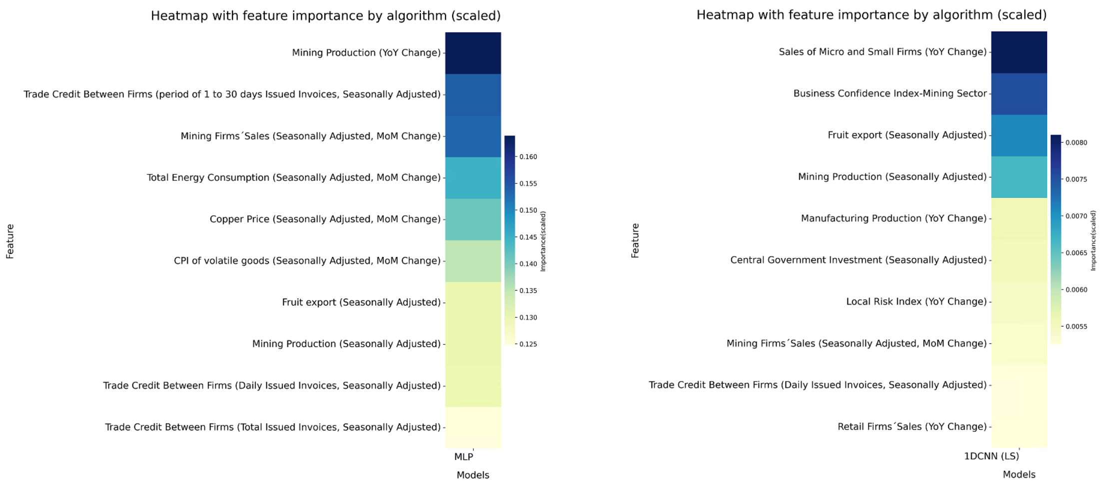

Notes: We calculate the importance of variables for the nonlinear MLP and 1D CNN models using SHAP. *Sources:* Own calculations.

It is necessary to differentiate between the important variables in obtaining nonlinear models, which vary according to the specific architecture of how the information is processed when obtaining the patterns. In the case of the MLP model, the important variables are obtained from the weights of the connections between the layers of the network. These values are independent, that is, they do not consider the sequential relationships, while the SHAP method, as pointed out, assigns the importance based on the impact of each variable in the prediction, considering the activation functions in each layer according to the architecture of the model. On the other hand, the 1D-CNN receives data sequentially, extracts the importance considering this context window and convolutional filters. Therefore, it is expected that CNN values will be lower, since CNN distributes the importance of the characteristics, reducing the variability in the SHAP values. Likewise, it should be considered that only the first two hundred sequences were used in the 1D-CNN network, in the same way, hierarchy of similar importance is observed. In terms of important variables, mining production, trade credit between firms and mining firms'sales contribute significantly to the prediction of the total IMACEC, as seen in figure 6.

On the other hand, for the prediction of the non-mining IMACEC, it can be seen in figure 7 that the variables that are common in the degree of importance are in a different order. Those that coincide to a greater degree are those variables of total energy consumption and manufacturing production, present in both models. The scale difference reflects the distinct architectures and processing approaches between the models. MLP processes all features globally and simultaneously, while 1D-CNN captures local temporal dependencies in reduced windows, revealing importance in these specific contexts. Therefore, in terms of variables that remain consistent across both models, these are systematically important.

Figure 7: Feature Importance by Algorithm - Non-mining IMACEC

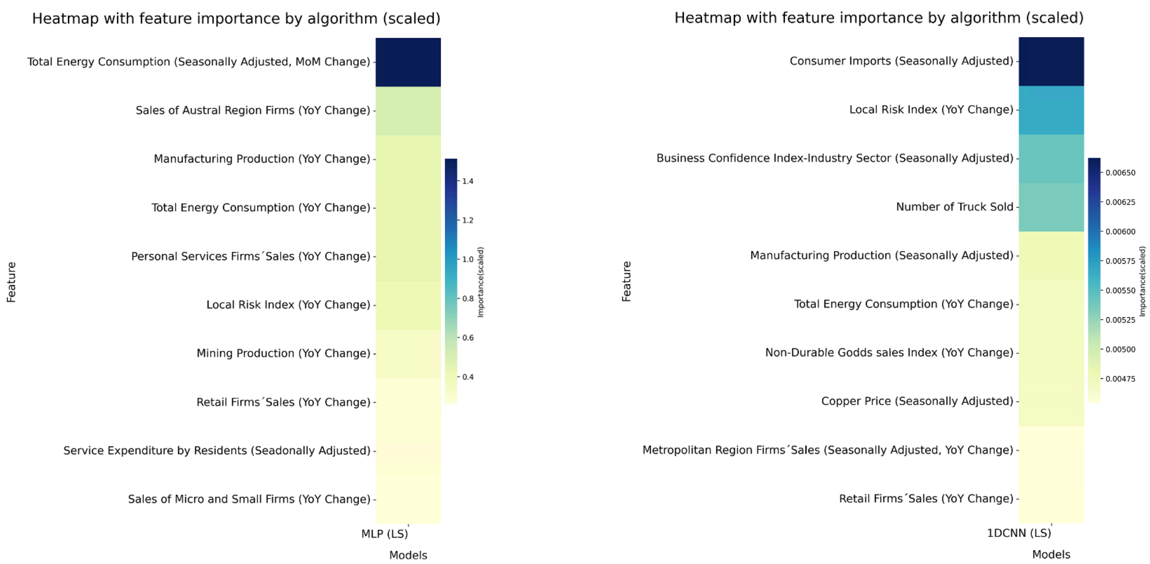

Notes: We calculate the importance of variables for the nonlinear MLP and 1D CNN models using SHAP. *Sources:* Own calculations.

5.6 Local Interpretable Model-agnostic Explanations (LIME)

In the case of the 1D-CNN, explainability methodology with LIME was used following Ribeiro et al. (2016), which is useful in interpreting predictions of complex models, as it is in this case. This technique gives local explanations around a specific instance using a linear model.

The process consists of generating perturbations given the instance X_0, where samples X_i are perturbed by introducing controlled variations in the attribute values of the initial instance, and the perturbations are created randomly in the input space of the feature set. Subsequently, each disturbed instance is introduced into the model, and the prediction is obtained. f corresponds to the prediction of the model:

$$\hat{y}_i = f(x_i)$$

It then defines the local weights and calculates the similarity between each perturbed instance and the original using the weighted Euclidean distance function. The proximity function $\pi_x(z)$ represents the similarity between the instance x and a perturbed sample z, where $D(x, z)$ is the distance between the instances, and σ controls the scale of the neighborhood:

$$\pi_x(z) = \exp\left(-\frac{D(x, z)^2}{\sigma^2}\right)$$

Then it uses the perturbed samples together with the predictions, and a weighted linear regression model is delivered to approximates the local behavior of the model. The interpretable model is observed in g in this case, the regression that minimizes the quadratic error.

$$\xi(x) = \arg\min_{g \in G} \sum_{i=1}^{N} \pi_x(x_i) \cdot (f(x_i) - g(x_i))^2$$

At this point, we can already interpret the coefficients of the fitted linear model, where each coefficient represents the local importance of the variable for the prediction in that passage of time.

The model is aligned with a locally adjusted linear framework, whereby the temporal sequence of characteristics is transformed into a flat vector. This transformation facilitates LIME's ability to execute controlled random perturbations during the original instance, thereby enabling the evaluation of the impact of each modification through the 1D-CNN model. The adjustment to a linear model allows to observe the relationship between the perturbations and the result of the prediction. The weight indicates how much the value affects the prediction, either positively or negatively, in addition to indicating the direction of the variation as it affects the prediction by considering those features. The following table 6 shows the results for the prediction of variables that are among the ten most important and which appear frequently throughout the iterations during the prediction period. In our case, the most frequent variables among the top contributors to the highest performance are found in January 2024 and October 2024 for the Total IMACEC and Non-mining, respectively.

The magnitude indicates that in the face of increases or decreases in the series the impact on the prediction can be positive or negative. In this study, the LIME model identifies the ten most salient variables influencing the monthly prediction. The hierarchy established by the model is based on those variables with the highest absolute weight in the local explanatory framework, with each being represented across a time window whose size is determined by the time sequence hyperparameter for each respective month. The prediction is formulated for the initial instance within the test set (index 0), where the LIME methodology generates numerous perturbed variations of this original sample. These perturbations maintain the temporal structure, preserving the dimensions [samples, time steps, features], and adhere to normalization. The values are perturbed within the normalized space (mean zero, variance one), thus conforming to the original distribution through the application of the original parameters.

The model employs a weighted linear regression to conform the perturbed samples, ascribing greater weight to perturbations that exhibit proximity to the original instance (utilizing similarity kernels) and weights that indicate how each variable, combined with a timestep, locally influences the prediction. Data normalization was conducted using a mean of zero and a variance of one; thus, the magnitude represents the standard deviations above or below the historical mean. In terms of weight, it signifies the marginal impact of specific conditions on the prediction. Furthermore, a tabular summary was produced since the variables recur in the estimates for each month. This recurrence indicates the relevance of these variables across multiple periods, albeit with varying importances according to the specific month of prediction. In essence, the relationship between a variable and the prediction is contingent upon the month in which the prediction occurs.

Table 5: Local Interpretable Model-agnostic Explanations (LIME) - Total IMACEC

Feature	**Average magnitude**	**Average weight**
Total Energy consumption (Seasonally Adjusted)	> 0.76	0.0409
Mining Firm's Sales (Seasonally Adjusted, MoM Change)	> 0.59	0.0391
Government Spending (SA, MoM Change)	> 0.61	0.0350

Note: Own elaboration. A 1D-CNN model was used, and LIME was applied to interpret the most frequent variables among the top ten contributors to the Total IMACEC in January 2024 (2024m01). This model yielded the best-performing prediction, achieving a Mean Absolute Error (MAE) of 0.09.

Table 6: Local Interpretable Model-agnostic Explanations (LIME) - Non-mining IMACEC

Feature	Average magnitude	Average weight
Manufacturing Production (Seasonally Adjusted)	> 0.62	0.0366
Number of Truck Sold	≤ 0.51	-0.0438
Local Risk Index (YoY Change)	≤ -0.59	0.0365

Note: Own elaboration. In contrast, for the Non-Mining IMACEC, the prediction for October 2024 (2024m10) also performed well, with a MAE of 0.54.

6 Conclusions

This paper addresses the use of machine learning (ML) and deep learning (DL) algorithms to perform nowcasting of total and non-mining monthly economic activity in Chile, aiming to bridge the information gap between the end of each month and the publication of the Official Monthly Activity Indicator (IMACEC). Providing timely nowcasts is particularly relevant for policymakers, as it reduces the information lag and allows them to react more quickly and effectively to changes in economic conditions, thereby improving the design and implementation of macroeconomic policies. Additionally, it incorporates non-traditional data sources from digital tax documents, electronic payment methods (such as company sales), and climatic data (temperature and precipitation). In this way, it integrates information sources from the economic data ecosystem that is increasingly available promptly in line with the digital transformation process, challenging conventional paradigms centered on traditional macroeconomic data.

Our results show that nonlinear machine learning models—particularly XGBoost—achieve the best performance in nowcasting Chile's economic activity, outperforming linear alternatives and surpassing traditional econometric benchmarks. For the non-mining IMACEC, Support Vector Regression (SVR) with a linear kernel and LASSO regression also deliver competitive results, confirming the usefulness of regularized linear approaches in high-dimensional settings. Nonlinear models such as XGBoost prove especially effective at capturing complex interactions among heterogeneous predictors, while linear models, though more interpretable, display more limited capacity to approximate economic dynamics. Traditional econometric models remain valuable reference points, yet machine learning techniques offer greater flexibility and adaptability—attributes particularly advantageous in volatile or heterogeneous environments. Overall, these models should not be viewed as substitutes for classical econometric frameworks, but rather as complementary tools that enrich the evaluation and comparison of nowcasting methodologies based on artificial intelligence.

We find that trade credit variables—both in number and amount—along with sectoral and regional sales extracted from electronic tax records, consistently emerge as significant predictors across all models, with variations only in their relative importance. This highlights the potential of electronic activity records in improving nowcast accuracy. The integration of such high-frequency variables provides clear evidence of how digital transformation in measurement sources enhances both the timeliness and precision of economic monitoring. By systematically incorporating trade credit indicators and disaggregated sales data, economic nowcasting frameworks gain greater robustness and flexibility, underscoring the value of new data sources for real-time policy analysis.

Finally, while deep learning models like MLP and CNN did not perform as strongly—partly due to limited data length—the exercise highlights the trade-offs between model complexity and data availability. Overall, the findings suggest that combining machine learning with non-traditional data offers a powerful complement to traditional econometric approaches, narrowing the forecast gap and enhancing the capacity for timely policy responses.

References

Adam, K., Marcet, A. & Nicolini, J. P. (2016), 'Stock market volatility and learning', *The Journal of Finance* **71**(1), 33–82.

Arro-Cannarsa, M. & Scheufele, R. (2024), 'Nowcasting gdp: what are the gains from machine learning algorithms?', *SNB Working Papers* .

Banbura, M., Giannone, D. & Reichlin, L. (2010), 'Nowcasting', *Working Paper Series ECB* pp. 487–93.
URL: *http://dx.doi.org/10.2139/ssrn.1717887*

Barrios, J., Escobar, J., Leslie, J., Martin, L. & Peña, W. (2021), Nowcasting to predict economic activity in real time: The cases of belize and el salvador, Technical report, Interamerican Development Bank.
URL: *https://publications.iadb.org/es/nowcasting-para-predecir-actividad-economica-en-tiempo-real-los-casos-de-belice-y-el-salvador*

Bengio, Y., Simard, P. & Frasconi, P. (1994), 'Learning long-term dependencies with gradient descent is difficult', *IEEE transactions on neural networks* **5**(2), 157–166.

Bishop, C. M. (2006), *Pattern recognition and machine learning*, springer.

Bolhuis, M. & Rayner, B. (2020), Deus ex machina? a framework for macro forecasting with machine learning, Technical report, IMF Working Papers.
URL: *https://www.imf.org/en/Publications/WP/Issues/2020/02/28/Deus-ex-Machina-A-Framework-for-Macro-Forecasting-with-Machine-Learning-49094*

Bolivar, O. (2024), 'Gdp nowcasting: A machine learning and remote sensing data-based approach for bolivia', *Latin American Journal of Central Banking* **5**(2).

Breiman, L. (2001), 'Statistical Modeling: The Two Cultures', *Statistical Science* **16**(3), 199–231.
URL: *https://www.jstor.org/stable/2676681*

Chen, T. & Guestrin, C. (2016), Xgboost: A scalable tree boosting system, *in* 'Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)', ACM, pp. 785–794.

Ciaburro, G. (2024), *MATLAB for Machine Learning – Second Edition: Unlock the Power of Deep Learning for Swift and Enhanced Results*, Packt Publishing.

Cobb, M. (2021), 'Nowcasting chilean household consumption with electronic payment data', *Working Paper Central Bank of Chile* (931).

Cobb, M. & Peña, J. (2020), 'Proyecciones de corto plazo para el pib trimestral: Desempeño reciente de una serie de modelos estándar', *Working Paper Central Bank of Chile* (871).

Diebold, F. & Mariano, R. (1995), 'Comparing predictive accuracy', *Journal of Business Economic Statistics 13* **13**(3).
URL: *https://doi.org/10.2307/1392185*

Foundation, P. S. (2024), 'nowcast-lstm', `https://pypi.org/project/nowcast-lstm/`. Versión 1.0.0, Licencia MIT, aprobada por OSI.

Graves, A. (2013), 'Generating sequences with recurrent neural networks', *arXiv preprint arXiv:1308.0850* .

Hetland, M. L. & Nelli, F. (2023), *Beginning Python: From Novice to Professional*, O'Reilly Media.
URL: *https://www.oreilly.com/library/view/beginning-python-from/9798868801969/*

Hochreiter, S. (1997), 'Long short-term memory', *Neural Computation MIT-Press* .

Hopp, D. (2022), 'Benchmarking econometric and machine learning methodologiess in nowcasting', *UNCTAD Research Paper* (83).

Jin-Kyu Jung, Manasa Patnam, A. T.-M. (2018), An algorithmic crystal ball: Forecasts-based on machine learning, IMF Working Papers 269, IMF Working Papers.
URL: *https://www.imf.org/en/Publications/WP/Issues/2018/11/01/An-Algorithmic-Crystal-Ball-Forecasts-based-on-Machine-Learning-46288*

Kant, D., Pick, A. & de Winter, J. (2022), 'Nowcasting gdp using machine learning methods'. Working paper, November 3, 2022.

Kavanagh, K. (2024), *Google Machine Learning and Generative AI for Solutions Architects*, Packt Publishing.
URL: *https://books.google.cl/books/about/Google_Machine_Learning_and_Generative_A.html?id=*

Kotsiantis, S. (2011), 'Feature selection for machine learning classification problems: a recent overview', *Artificial intelligence review* **42**(1), 157–176.

LeCun, Y., Bottou, L., Bengio, Y. & Haffner, P. (1998), 'Gradient-based learning applied to document recognition', *Proceedings of the IEEE* **86**(11), 2278–2324.

Lundberg, S. M. & Lee, S.-I. (2017), 'A unified approach to interpreting model predictions', *Advances in Neural Information Processing Systems (NeurIPS)* **30**.
URL: *https://arxiv.org/abs/1705.07874*

Németh, K. & Hadházi, D. (2024), ‘Gdp nowcasting with artificial neural networks: How much does long-term memory matter?’, *arXiv e-prints* .
URL: *https://arxiv.org/abs/2304.05805*

Ranjan, A. & Ghosh, S. (2021), ‘A machine learning (ml) approach to gdp nowcasting: An emerging market experience’.

Ribeiro, M. T., Singh, S. & Guestrin, C. (2016), “‘why should i trust you?": Explaining the predictions of any classifier’, *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD ’16)* pp. 1135–1144.
URL: *https://dl.acm.org/doi/10.1145/2939672.2939778*

Richardson, A., van Florenstein Mulder, T. & Vehbi, T. (2021), ‘Nowcasting gdp using machine-learning algorithms: A real-time assessment’, *International Journal of Forecasting* **37**(2), 941–948.
URL: *ttps://www.sciencedirect.com/science/article/pii/S016920702030159X*

Rosenblatt, M. (1958), ‘The multidimensional prediction problem.’.

Rumelhart, D. E., Hinton, G. E. & Williams, R. J. (1986), ‘Learning representations by back-propagating errors’, *nature* **323**(6088), 533–536.

Soybilgen, B. & Yazgan, E. (2021), ‘Nowcasting us gdp using tree-based ensemble models and dynamic factors’, *Computational Economics* **57**(1), 387–417.
URL: *https://doi.org/10.1007/s10614-020-10083-5*

Tenorio, J. & Pérez, W. (2024), ‘Gdp nowcasting with machine learning and unstructured data’, *Working Paper. Banco Central de Reserva del Perú* .
URL: *https://www.bcrp.gob.pe/docs/Publicaciones/Documentos-de-Trabajo/2024/documento-de-trabajo-003-2024.pdf*

Tibshirani, R. (1996), ‘Regression shrinkage and selection via the lasso’, *Journal of the Royal Statistical Society: Series B (Methodological)* **58**(1), 267–288.

Tiffin, A. J. (2016), ‘Seeing in the dark: A machine-learning approach to nowcasting in lebanon’, *IMF Working Papers* (Working Paper No. 2016/056).
URL: *https://www.imf.org/en/Publications/WP/Issues/2016/12/31/Seeing-in-the-Dark-A-Machine-Learning-Approach-to-Nowcasting-in-Lebanon-43779*

Vapnik, V. (1997), ‘Support-vector networks’, *Machine Learning* **20**(3), 273–297.

Vapnik, V. (1998), *Statistical Learning Theory*, Wiley, New York.

Werbos, P. J. (1989), ‘Beyond the gibbs sampler: Neural networks, backpropagation, and gradient descent in the system modeling context’, *IEEE Transactions on Neural Networks* **1**(1), 1–13.

Werbos, P. J. (1990), ‘Backpropagation through time: what it does and how to do it’, *Proceedings of the IEEE* **78**(10), 1550–1560.

Yi, H. C., Choi, D. & Kim, Y. (2022), ‘Dynamic factor model and deep learning algorithm for gdp nowcasting’. Working paper, Bank of Korea.

Yoon, J. (2021), ‘Forecasting of real gdp growth using machine learning models: Gradient boosting and random forest approach’, *Computational Economics* **57**, 247–265.
URL: *https://doi.org/10.1007/s10614-020-10054-w*

Zou, H. & Hastie, T. (2005), ‘Regularization and variable selection via the elastic net’, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **67**(2), 301–320.
URL: *https://sites.stat.washington.edu/courses/stat527/s13/readings/zouhastie05.pdf*

Appendix

A Data

Table A.1: List of traditional variables included in the model (a)

No.	Variable	Units of Measure	Source
	Main indicators		
1	IMACEC	Index (base 2018=100)	BCCh
2	IMACEC Non mining	Index (base 2018=100)	BCCh
	Expectations		
3	IPEC: Total	Index	Adimark
4	IPEC: Current country situation	Index	Adimark
5	IPEC: Future country situation 1Y	Index	Adimark
6	IPEC: Future country situation 5Y	Index	Adimark
7	IPEC: Current personal situation	Index	Adimark
8	IPEC: Purchase household items	Index	Adimark
9	IPEC: Unemployment perception	Index	Adimark
10	IPEC: Inflation expectations 1Y	Index	Adimark
11	IMCE: Total	Index	UAI/ICARE
12	IMCE: Commerce	Index	UAI/ICARE
13	IMCE: Construction	Index	UAI/ICARE
14	IMCE: Manufacturing	Index	UAI/ICARE
15	IMCE: Mining	Index	UAI/ICARE
16	IMCE: Non Mining	Index	UAI/ICARE
	Financial Indicators		
17	IPSA	Index	BCS
18	IPSA real	Index	BCS
19	Consumer credits	billions of real pesos	BCCh
20	Mortage Loans	billions of real pesos	BCCh
21	Foreign trade credits	billions of real pesos	BCCh
22	Lending interest rate, 1 to 3 years	Percentage	BCCh
23	Lending interest rate, more than 3 years	Percentage	BCCh
	International Indicators		
24	Real exchange rate index	Index	BCCh
25	SOFR (ex Libor 6 months)	Index	Bloomberg
26	Copper price	Index	Bloomberg
27	VIX	Index	Bloomberg
28	Local risk index	Index	BCCh
29	External risk index	Index	BCCh

Table A.2: List of traditional variables included in the model (b)

No.	Variable	Units of Measure	Source
	Labor Market		
30	Real wage bill	Index	BCCh
31	ICL real	Index	INE
32	Unemployment rate	Percentage	INE
33	Wage earners	Thousands of people	INE
34	Employment in commerce	Thousands of people	INE
35	Construction employment	Thousands of people	INE
	Monetary Indicators		
36	CPI	Index (Base 2023 = 100)	INE
37	Core CPI	Index (Base 2023 = 100)	INE
38	Non-volatile goods	Index (Base 2023 = 100)	INE
39	Non-volatile goods excluding non-volatile foods	Index (Base 2023 = 100)	INE
40	Non-volatile foods	Index (Base 2023 = 100)	INE
41	Non-volatile services	Index (Base 2023 = 100)	INE
42	CPI Volatile	Index (Base 2023 = 100)	INE
43	Volatile foods	Index (Base 2023 = 100)	INE
44	Volatile Energy	Index (Base 2023 = 100)	INE
45	Remaining Volatile Items	Index (Base 2023 = 100)	INE
46	Volatile goods	Index (Base 2023 = 100)	INE
47	Volatile Services	Index (Base 2023 = 100)	INE
	Fiscal Indicators		
48	Total Central Government Revenues	CLP millions	Dipres
49	Central Government Tax Revenues	CLP millions	Dipres
50	Total Central Government Expenditures	CLP millions	Dipres
51	Total Central Government Investments	CLP millions	Dipres
	Activity Indicators		
52	Retail Trade Activity (IACM): Total	Index (Base 2018 = 100)	INE
53	IACM: Durable	Index (Base 2018 = 100)	INE
54	IACM: Non-Durable	Index (Base 2018 = 100)	INE
55	Supermarket sales	Index (Base 2018 = 100)	INE
56	Cars Sales	Units	ANAC
57	Truck sales	Units	ANAC
58	Mining output	Index (Base 2018=100)	INE
59	Concrete	Index	INE
60	Manufacturing output	Index	INE
61	Electric power generation	GWh	CNE
62	Supplier index	Index	CChC

Table A.3: List of traditional variables included in the model (c)

No.	Variable	Units of Measure	Source
	International Trade Indicators		
63	Total Imports (CIF)	USD millions	BCCh
64	Consumer Goods Imports	USD millions	BCCh
65	Durable Goods Imports	USD millions	BCCh
66	Semi-Durable Goods Imports	USD millions	BCCh
67	Other consumer goods Imports	USD millions	BCCh
68	Intermediates goods Imports	USD millions	BCCh
69	Energy products Imports	USD millions	BCCh
70	Other Intermediate Goods Imports	USD millions	BCCh
71	Capital goods Imports	USD millions	BCCh
72	Total Exports	USD millions	BCCh
73	Mining Exports	USD millions	BCCh
74	Agri./Livest., For. & Fish Exports	USD millions	BCCh
75	Manufacturing industry Exports	USD millions	BCCh
76	Total Exports excl. Mining	USD millions	BCCh
77	Copper Exports	USD millions	BCCh
78	Lithium carbonate Exports	USD millions	BCCh
79	Food products Exports	USD millions	BCCh
80	Fruits-farming sector Exports	USD millions	BCCh

Table A.4: List of non-traditional variables included in the model (d)

No.	Variable	Units of Measure	Source
Experimental Statistics			
81	Electronic Invoice (EI): Total	CLP	SII
82	EI: Agriculture and forestry	CLP	SII
83	EI: Mining	CLP	SII
84	EI: Wholesale and retail trade	CLP	SII
85	EI: Construction	CLP	SII
86	EI: Elec., Gas, Water & Waste Mgmt.	CLP	SII
87	EI: Restaurants and hotels	CLP	SII
88	EI: Manufacturing industry	CLP	SII
89	EI: Business services	CLP	SII
90	EI: Personal services	CLP	SII
91	EI: Transp. & Comm.	CLP	SII
92	EI: Northern Macrozone	CLP	SII
93	EI: Metropolitan Macrozone	CLP	SII
94	EI: Central Macrozone	CLP	SII
95	EI: South-Central Macrozone	CLP	SII
96	EI: Southern Macrozone	CLP	SII
97	EI: Austral Macrozone	CLP	SII
98	EI: Micro and Small Enterprises	CLP	SII
99	EI: Medium Enterprises	CLP	SII
100	EI: Large Enterprises	CLP	SII
101	Boleta electrónica (BE): Total	CLP	SII
102	BE: Stores	CLP	SII
103	BE: Supermarkets	CLP	SII
104	BE: Great stores	CLP	SII
105	BE: Clthg., Ftwr. & HH Eq.	CLP	SII

Table A.5: List of non-traditional variables included in the model (e)

No.	Variable	Units of Measure	Source
	Experimental Statistics		
106	Business-to-business trade credit: Total	Amount in billions of pesos	BCCh
107	B2B trade credit: Up to date	Amount in billions of pesos	BCCh
108	B2B trade credit: 1 to 30 days	Amount in billions of pesos	BCCh
109	B2B trade credit: 31 to 60 days	Amount in billions of pesos	BCCh
110	B2B trade credit: 61 to 90 days	Amount in billions of pesos	BCCh
111	B2B trade credit: Up from 91 days	Amount in billions of pesos	BCCh
112	Business-to-business trade credit: Total	No. of Inv. Issued (k units)	BCCh
113	B2B trade credit: Up to date	No. of Inv. Issued (k units)	BCCh
114	B2B trade credit: 1 to 30 days	No. of Inv. Issued (k units)	BCCh
115	B2B trade credit: 31 to 60 days	No. of Inv. Issued (k units)	BCCh
116	B2B trade credit: 61 to 90 days	No. of Inv. Issued (k units)	BCCh
117	B2B trade credit: Up from 91 days	No. of Inv. Issued (k units)	BCCh
	Climate Indicators (*)		
118	Regional accumulated precipitation	Millimeter	DMC
119	Regional average temperature	Celsius (°C)	DMC
120	Regional maximum temperature	Celsius (°C)	DMC
121	Regional minimum temperature	Celsius (°C)	DMC
	Electronic Payment Data		
122	Residents' total sales	CLP	TBK+GetNet
123	Residents' sales of goods	CLP	TBK+GetNet
124	Residents' sales of services	CLP	TBK+GetNet
125	Non-residents' total sales	CLP	TBK+GetNet
126	Non-residents' sales of goods	CLP	TBK+GetNet
127	Non-residents' sales of services	CLP	TBK+GetNet

(*) Precipitation, minimum, maximum and average temperature data correspond to regions 1 to 16.

B Glosary

Concept	**Definition**
ANAC	National Automotive Association of Chile.
BCCh	Central Bank of Chile.
BCS	Santiago Stock Exchange.
CChC	Chilean Chamber of Construction.
CNE	National Energy Commission.
CPI	Consumer Price Index.
DIPRES	Budget Directorate of Chile.
DMC	Meteorological Directorate of Chile.
GWh	Gigawatt-hour.
IACM	Retail Trade Activity Index.
ICL	Labor Cost Index.
IMACEC	Monthly Economic Activity Indicator.
IMCE	Monthly Business Confidence Indicator.
INE	National Statistics Institute of Chile.
IPEC	Economic Perception Index.
IPSA	Selective Stock Price Index.
UAI/ICARE	Adolfo Ibáñez University / ICARE.
SII	Internal Revenue Service of Chile.
SOFR	Secured Overnight Financing Rate.
TBK+GetNet	Transbank and GetNet.
VIX	Volatility Index.

C Correlations

Figure C.1: Correlation with the target variable: Total IMACEC

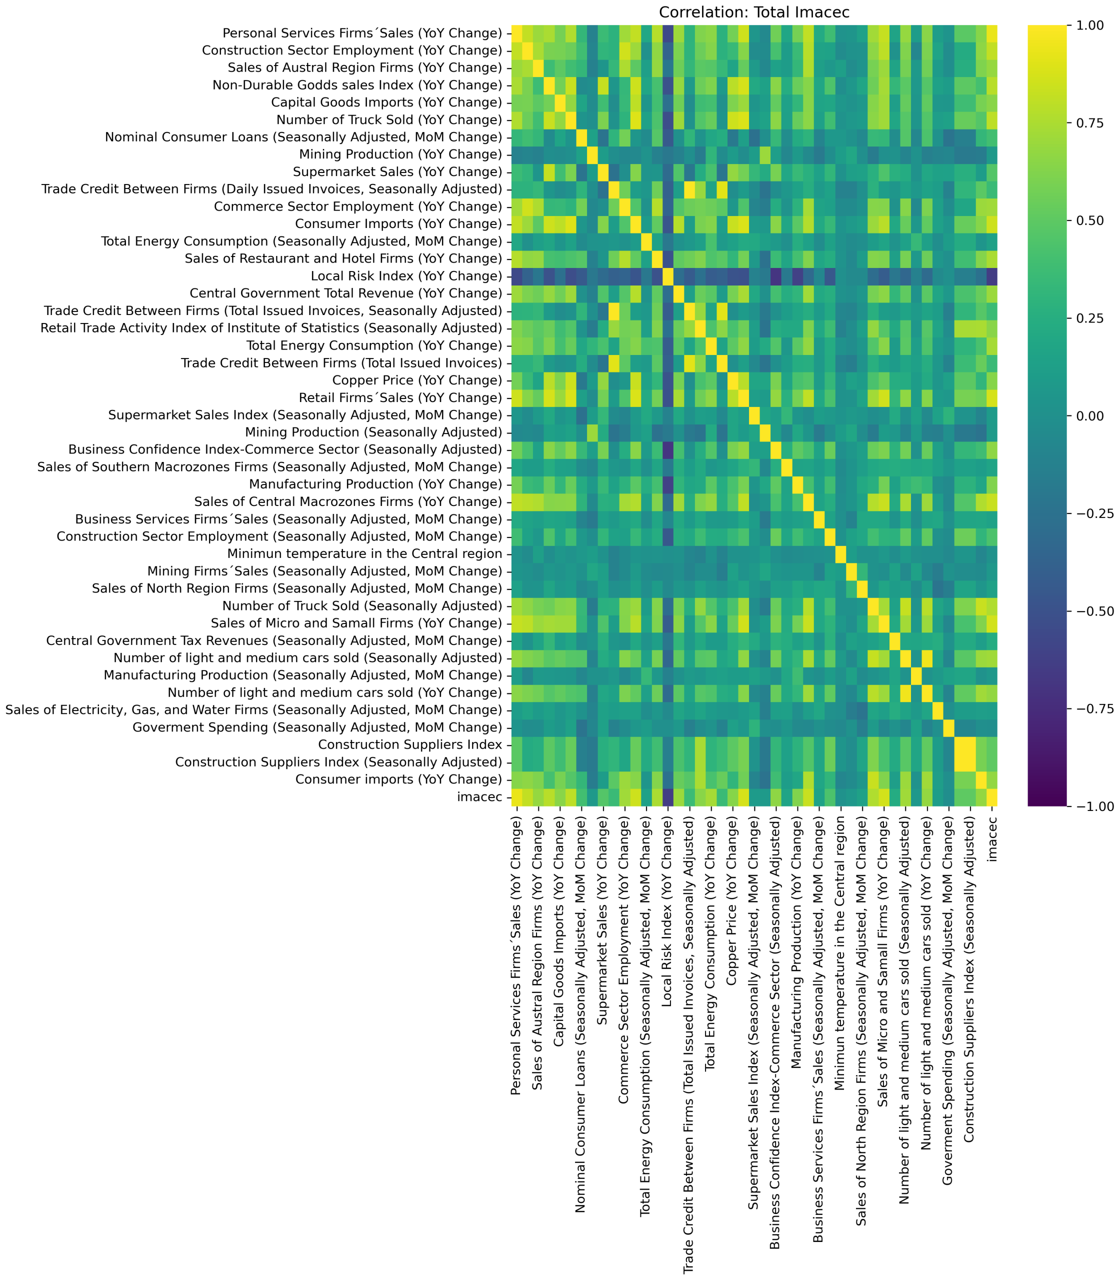

Notes: Yellow color positive correlation and dark negative correlation. *Sources:* Own calculations.

Figure C.2: Correlation with the target variable: Non-mining IMACEC

Notes: Yellow color positive correlation and dark negative correlation. *Sources:* Own calculations.

D SVR Example

Figure D.3: Support Vector Regressor Example

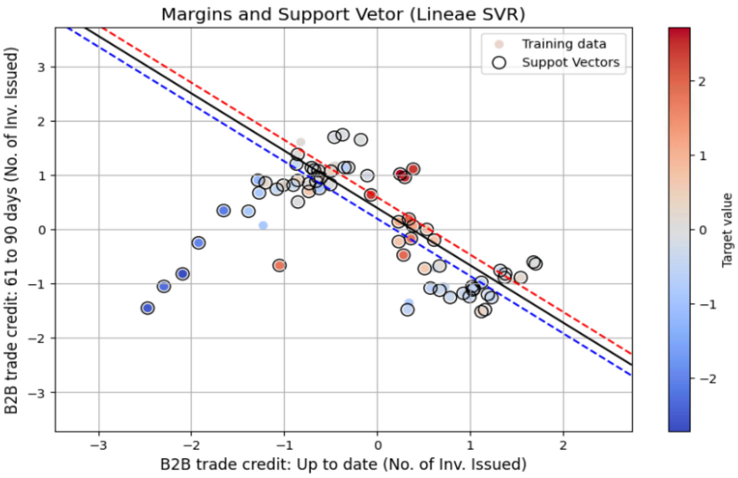

Notes: The graph explains how support vector machines for regression are visualized in a two-dimensional format. To explain the operation, two variables are considered, where the X-axis shows inter-company credits Up to date and the Y-axis shows credits from 61 to 90 days. The data is already normalized and used to predict the IMACEC.

The SVR is predicting based on two inter-company credit variables: those paid on time and those with much longer terms. The black line corresponds to the decision function. The negative slope shows the economic relationship, where more invoices paid on time indicate fewer invoices over 60 days overdue and a higher IMACEC, indicating a healthy economy with greater purchasing power.

The blue points indicate a low target variable and negative values, where periods of economic contraction show less general activity, also indicating an increase in invoices with longer payment terms. The right side shows red and orange points with more invoices paid on time and longer terms during periods of economic expansion with activity that can be interpreted as better cash flow and payment security. The central zone shows intermediate behavior.

The black circles correspond to support vectors, identifying inflection points. The dotted lines are the margins where points within them represent the normal relationship between credit and IMACEC. As presented, commercial credit between companies behaves as a leading indicator where early signals of activity can be observed.

E Models without electronic tax records

Table E.7: Fine-Tuned Out-of-Sample Performance Metrics by Algorithm for ML Models
Results without IMACEC Revisions or Experimental Statistics (excluding electronic tax records), 2024m01–2024m12

Algorithm	**Total IMACEC**			**Non-mining IMACEC**		
	MSE	RMSE	MAE	MSE	RMSE	MAE
LASSO	1.74	1.32	1.09	2.80	1.67	1.39
Ridge	5.76	2.40	1.93	6.99	2.64	2.01
Elastic net	1.62	1.27	1.07	4.72	2.17	1.61
SVR (LASSO)	1.46	1.21	1.09	3.23	1.80	1.48
MLP (EN)	3.14	1.77	1.43	3.21	1.79	1.63
1D CNN (EN/LASSO)	4.75	2.18	1.84	5.26	2.29	2.01
XGBoost	2.39	1.55	1.22	3.23	1.80	1.47
Traditional/Econometrics Models	3.08	1.75	1.48	3.08	1.76	1.53

Note: The machine learning models without electronic tax records. Results are maintained but with lower performance. Between 2024m01-2024m12. Source: Own elaboration.

Documentos de Trabajo **Banco Central de Chile**	**Working Papers** **Central Bank of Chile**
NÚMEROS ANTERIORES	PAST ISSUES
La serie de Documentos de Trabajo en versión PDF puede obtenerse gratis en la dirección electrónica: www.bcentral.cl/esp/estpub/estudios/dtbc. Existe la posibilidad de solicitar una copia impresa con un costo de Ch$500 si es dentro de Chile y US$12 si es fuera de Chile. Las solicitudes se pueden hacer por fax: +56 2 26702231 o a través del correo electrónico: bcch@bcentral.cl.	Working Papers in PDF format can be downloaded free of charge from: www.bcentral.cl/eng/stdpub/studies/workingpaper. Printed versions can be ordered individually for US$12 per copy (for order inside Chile the charge is Ch$500.) Orders can be placed by fax: +56 2 26702231 or by email: bcch@bcentral.cl.

DTBC – 1058
Artificial Intelligence Models for Nowcasting Economic Activity
Jennifer Peña, Katherine Jara, Fernando Sierra

DTBC – 1057
Clasificación de Riesgo de Crédito Bancario, Ventas y Estados Financieros en Base a Información Tributaria de Firmas en Chile
Ivette Fernández D., Jorge Fernández B., Francisco Vásquez L.

DTBC – 1056
Exogenous Influences on Long-term Inflation Expectation Deviations: Evidence from Chile Carlos A. Medel

DTBC – 907*
Earnings Inequality in Production Networks
Federico Huneeus, Kory Kroft, Kevin Lim

DTBC – 1055
Markup Distribution and Aggregate Dynamics Mario Giarda, Will Jianyu Lu, Antonio Martner

DTBC – 1054
Decoding Central Banks' Practices: A Closer Look at Inflation Expectations Surveys
Valentina Cortés Ayala, Karlla Muñoz Cáceres, Daniel Pérez Klein

DTBC – 1053

An Assessment of the Effects of Monetary Policy Communication in Chile

Mario González-Frugone, Ignacio Rojas

DTBC – 1052

This Time is Global: Synchronisation in Economic Policy Uncertainty Indices

Carlos Medel

DTBC – 1051

Beyond Costs: The Dominant Role of Strategic Complementarities in Pricing

Elias Albagli, Francesco Grigoli, Emiliano Luttini, Dagoberto Quevedo, Marco Rojas

DTBC – 1050

The Anatomy of Aggregate Productivity

Federico Huneeus, Yasutaka Koike-Mori, Antonio Martner

DTBC – 1049

The Aggregate Welfare Effects of Nonlinear Prices in Supply Chains

Luca Lorenzini, Antonio Martner

DTBC – 1026*

The Incidence of Distortions

David Atkin, Baptiste Bernadac, Dave Donaldson, Tishara Garg, Federico Huneeus

DTBC – 1006*

Input Price Dispersion Across Buyers: New Evidence and Implications for Aggregate Productivity

Ariel Burstein, Javier Cravino, Marco Rojas

DTBC – 1048

The effect of automation on the labor market: An approach using firm-level microdata

Camilo Levenier

DTBC – 1047

Firm-level CO2 Emissions and Production Networks: Evidence from Administrative Data in Chile

Pablo Acevedo, Elías Albagli, Gonzalo García-Trujillo, María Antonia Yung

DOCUMENTOS DE TRABAJO Diciembre 2025