

Article

Vehicle Registrations as a Leading Indicator of the Spanish Business Cycle: Machine Learning and Classical Forecasting Approaches

Rodrigo Gragera , Rufino Prieto and Carlos Poza * 

Department of Business Analytics, Universidad Francisco de Vitoria, 28223 Madrid, Spain; rgrageraflorez@gmail.com (R.G.); rufino.prieto@ufv.es (R.P.)

* Correspondence: carlos.poza@ufv.es

Abstract

Timely identification of business-cycle turning points is essential for policymakers, firms and investors, yet traditional macroeconomic indicators are often published with considerable delays. This study examines whether vehicle registrations constitute an effective Leading Economic Indicator (LEI) for the Spanish economy and proposes an integrated framework combining business-cycle analysis, monthly forecasting and daily nowcasting. The cyclical properties of registrations are analyzed using the Christiano–Fitzgerald band-pass filter, cross-correlation analysis and the Bry–Boschan dating algorithm. Monthly forecasting performance is evaluated by comparing TBATS and Prophet through rolling-origin cross-validation, while a novel machine learning-inspired distribution model (MD) is developed to transform monthly forecasts into daily estimates. The results show that passenger-car registrations anticipate industrial production by three to five months, with a ninety per cent bootstrap interval of one to six. TBATS consistently outperforms Prophet in one-step-ahead forecasting. An out-of-sample test confirms that lagged registrations reduce the forecast error of industrial production without recourse to two-sided filtering, although the reduction is modest in magnitude and is not found for quarterly GDP. The proposed MD-shallow model achieves the highest distributional accuracy, surpassing both Prophet and a persistence benchmark. The findings demonstrate that vehicle registrations are a valuable standalone Leading Economic Indicator and that integrating traditional economic indicators with advanced forecasting techniques improves real-time monitoring of automotive demand and business-cycle dynamics. The proposed framework provides a practical tool for policymakers, manufacturers, and researchers seeking earlier and more reliable assessments of economic activity.

Keywords: Leading Economic Indicators; forecasting; machine learning

JEL Classification: C22; E21; E32



Academic Editor: Periklis Gogas

Received: 30 June 2026

Revised: 30 July 2026

Accepted: 3 August 2026

Published: 9 August 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Accurately anticipating business-cycle turning points remains one of the main challenges in macroeconomics. Governments, central banks, firms and investors rely on timely information to support policy decisions, investment planning and risk management. However, traditional macroeconomic indicators, such as Gross Domestic Product (GDP), industrial production and employment, are published with considerable delays and are often

revised. Consequently, identifying variables capable of anticipating changes in economic activity has become a major research objective.

Leading Economic Indicators (LEIs) play a key role in reducing these informational delays by providing early signals of expansion and recessions. Among them, passenger-car registrations are particularly relevant because vehicle purchases are highly sensitive to consumer confidence, expected income and financing conditions. As a result, registrations frequently react before output and employment, making them a valuable indicator of future economic activity. Their inclusion in the OECD composite leading indicators further supports their relevance for business-cycle monitoring.

Recent advances in Big Data and machine learning have also transformed economic forecasting. Digital information from internet searches, social media and online news has considerably improved nowcasting and short-term prediction. Seminal contributions by Diebold (2000), Choi and Varian (2011), Scott and Varian (2012) and Varian (2014) demonstrated the value of large, high-frequency datasets for economic forecasting, while subsequent studies confirmed the predictive power of digital information for unemployment, labor market dynamics and recession forecasting (Antenucci et al., 2014; D'Amuri & Marcucci, 2017; Huang et al., 2018). More recently, Poza and Monge (2020) showed that combining digital information with advanced time-series techniques improves real-time monitoring of the Spanish economy. Nevertheless, relatively little attention has been devoted to analyzing passenger-car registrations as an independent leading indicator or to integrating their forecasting potential with modern machine learning techniques.

Against this background, this paper pursues three closely related objectives. First, it examines whether vehicle registrations contain robust leading information on the Spanish business cycle by analyzing their cyclical relationship with industrial production and GDP and evaluating their predictive content in a real-time forecasting setting. Second, it compares alternative forecasting approaches for monthly vehicle registrations, assessing whether a changepoint-based model (Prophet) outperforms a multiple-seasonality state-space model (TBATS) in a series characterized by pronounced structural breaks. Third, it develops and evaluates a novel machine learning-inspired distribution model for transforming monthly forecasts into daily nowcasts, testing whether greater model flexibility improves the accuracy of the intra-month allocation of registrations.

To achieve these objectives, the study adopts a multiscale empirical framework combining business-cycle analysis, out-of-sample forecasting and temporal disaggregation. The leading properties of vehicle registrations are assessed through cyclical analysis, turning-point detection and predictive validation, while monthly forecasting performance is evaluated using rolling-origin cross-validation. Finally, the proposed distribution model estimates the daily allocation of monthly registrations, allowing real-time monitoring from partial monthly observations.

The remainder of the paper is organized as follows. Section 2 reviews the literature on Leading Economic Indicators and recent forecasting techniques. Section 3 describes the data and methodology. Section 4 presents the empirical results, while Section 5 concludes and outlines the main implications and future research directions.

2. Literature Review

2.1. The Importance of LEIs to Anticipating Business Cycles

Monitoring business-cycle turning points has long been one of the central challenges in macroeconomics. Economic policy, business investment, and financial decision-making depend critically on the ability to identify changes in economic activity before they become visible in conventional macroeconomic statistics. However, most official indicators, including Gross Domestic Product (GDP), industrial production and employment, are

published with considerable delays and are frequently revised after their initial release. Consequently, economists have devoted substantial effort to identifying Leading Economic Indicators (LEIs) capable of anticipating expansions and recessions before they are reflected in official statistics.

LEIs are variables whose movements systematically precede fluctuations in aggregate economic activity, providing early information about future changes in the business cycle. Unlike coincident indicators, which merely describe the current state of the economy, or lagging indicators, which confirm events that have already occurred, leading indicators offer forward-looking information that reduces uncertainty and improves forecasting performance. As Baumohl (2009) explains, no single indicator perfectly predicts every recession or expansion, but carefully selected leading variables frequently provide valuable signals several months before turning points become evident in traditional macroeconomic series.

The theoretical rationale behind leading indicators is straightforward. Many economic decisions (including investment, durable goods purchases, hiring intentions and financial transactions) are taken before they materialize in aggregate output. Consequently, variables reflecting households' expectations, firms' investment decisions or credit conditions often anticipate subsequent changes in production and employment. This explains why institutions such as the OECD and The Conference Board have traditionally relied on composite leading indicators that combine several forward-looking variables to improve forecasting performance.

The growing availability of timely information has considerably transformed the construction of LEIs during the last two decades. Traditional composite indicators were mainly built using official macroeconomic variables released monthly or quarterly. Although these indicators remain useful, their relatively low frequency limits their ability to detect sudden changes in economic conditions. Diebold (2000) anticipated that the increasing availability of large datasets would fundamentally change macroeconomic forecasting, allowing economists to exploit much richer information sets than those employed by conventional econometric models. Similarly, Varian (2014) argues that Big Data has introduced entirely new possibilities for econometric analysis, particularly by incorporating digital traces generated continuously by millions of economic agents.

Among the earliest and most influential contributions to this literature, Choi and Varian (2011) demonstrated that Google search activity contains valuable real-time information capable of improving short-term predictions of several macroeconomic variables. Since internet searches reflect individuals' intentions before actual economic decisions are observed, Google Trends became one of the first large-scale digital datasets successfully incorporated into economic forecasting models. Their work established an entirely new research agenda linking online behavior with macroeconomic dynamics.

Subsequent studies have confirmed the predictive content of digital information across different economic variables. D'Amuri and Marcucci (2017) showed that Google searches significantly improve forecasts of U.S. unemployment, particularly during periods of rapid labor market deterioration. Likewise, Antenucci et al. (2014) demonstrated that social media activity contains useful information regarding labor market flows, allowing researchers to monitor employment dynamics almost in real time. These contributions illustrate how digital footprints frequently anticipate official statistics because they capture individuals' behavior before it becomes visible in conventional economic indicators.

Beyond internet searches and social media, researchers have increasingly explored additional sources of digital information. Huang et al. (2018) found that news sentiment extracted through textual analysis provides statistically significant early signals of recessions. Since news reflects both current events and expectations regarding future economic conditions, sentiment indicators often react more rapidly than traditional macroeconomic

variables. Similarly, Szármes (2014) argues that Big Data technologies enable the construction of entirely new business and economic indicators capable of complementing official statistics with high-frequency information. Collectively, these studies support the emergence of a new generation of digital LEIs that continuously monitor economic activity through online behavior, textual information and social interactions.

From a methodological perspective, these developments have been accompanied by significant advances in real-time forecasting. Scott and Varian (2012) proposed Bayesian variable selection techniques capable of identifying the most informative predictors from extremely large datasets, while Heinisch and Scheufele (2019) emphasized the importance of evaluating forecasting models using real-time information rather than revised historical data. Their results highlight that forecasting performance depends not only on model specification but also on the information actually available to forecasters when predictions are produced.

Although digital indicators have attracted considerable attention in recent years, traditional leading indicators continue to play a fundamental role in monitoring the business cycle. Among them, passenger-car registrations and vehicle sales constitute one of the most informative indicators of future economic activity. Purchasing a new vehicle represents one of the largest discretionary expenditures made by households and firms. Because these purchases can easily be postponed during periods of uncertainty or accelerated when economic expectations improve, vehicle registrations respond rapidly to changes in consumer confidence, disposable income, financing conditions and investment expectations. Consequently, they frequently anticipate fluctuations in aggregate demand before these become visible in GDP or industrial production.

The leading properties of vehicle registrations are also supported by their close relationship with several macroeconomic transmission mechanisms (Cantos-Sánchez et al., 2018). New vehicle purchases depend heavily on household expectations regarding future income, labor market stability, interest rates and access to consumer credit. Similarly, fleet purchases made by firms reflect expectations about future production, logistics demand and business expansion. As both households and firms adjust these decisions ahead of actual production changes, vehicle registrations naturally become forward-looking indicators of business-cycle dynamics.

For this reason, vehicle registrations have been incorporated into several internationally recognized composite leading indicators, including those produced by the OECD. Their inclusion reflects decades of empirical evidence showing that durable goods demand typically leads aggregate economic activity by several months. Nevertheless, despite their recognized forecasting ability, surprisingly few studies have analyzed vehicle registrations as an independent leading indicator or explored their potential using modern forecasting techniques and high-frequency data.

Recent research has begun to combine traditional leading indicators with Big Data sources to improve forecasting accuracy. Poza and Monge (2020), for example, developed a real-time Leading Economic Indicator for the Spanish economy by integrating text mining techniques with fractional cointegration methods and continuous wavelet analysis. Their results demonstrate that combining digital information with advanced time-series methodologies substantially improves the monitoring of economic activity. This integrated approach suggests that traditional indicators and digital footprints should not be viewed as competing alternatives but rather as complementary sources of information capable of capturing different dimensions of the economic cycle.

Despite these important advances, two important research gaps remain. First, relatively limited attention has been devoted to exploiting passenger-car registrations as a standalone leading indicator of business-cycle turning points, despite their recognized role

within international composite indicators. Second, although Big Data has transformed macroeconomic nowcasting, relatively few studies have integrated traditional durable goods indicators with modern machine learning techniques to improve forecasting performance at different temporal frequencies. Addressing these gaps constitutes the main motivation of the present study, which examines both the leading properties of Spanish vehicle registrations and the forecasting potential of modern machine learning approaches for anticipating automotive demand. This focus naturally extends the recent literature on digital economic indicators while reaffirming the continuing relevance of carefully selected traditional LEIs for monitoring business-cycle dynamics.

2.2. *Advances in Forecasting Methodology to Monitor Business Cycle*

Vehicle registrations lead the business cycle, but exploiting that signal poses two distinct problems: modelling a series with several overlapping seasonalities and abrupt level breaks and moving from monthly to daily data by recovering how a within-month total otherwise known only at month-end is distributed across days. The literature has addressed the two separately; this paper joins them on the case of passenger-car registrations in Spain. The review surveys the cyclical behavior of demand, monthly forecasting models, multiscale analysis and temporal disaggregation, before identifying the gaps and the theoretical framework (Kim, 2023; Ledo-Lobato et al., 2024; Muth et al., 2025).

2.3. *Prior Work and Current State*

The automobile is a durable good whose purchase can be postponed; households defer the decision when income or confidence falls and bring it forward when they improve, so registrations move ahead of the economy and with greater cyclical amplitude. Dargay and Gately (1999) show that the relationship between income and the vehicle stock is non-linear: near saturation, each additional unit of income adds fewer new registrations. In an economy with an already-high stock, this calls for a model able to capture regime changes in the trend rather than a fixed income–demand relationship.

On which methods forecast best, the systematic review by Ranabhatt et al. (2025) of 63 studies in the sector under the PRISMA protocol concludes that classical time-series models remain competitive at short and medium horizons, and that external variables (GDP, fuel prices, consumer confidence) improve forecasts consistently. Cantos-Sánchez et al. (2018) further warn that scrappage schemes pulled purchases forward and distorted the Spanish series, which makes it necessary to isolate such policy effects and regulatory episodes, such as the 2018 WLTP standard.

Registrations combine the annual cycle, the end-of-month concentration and a working-day effect, together with abrupt level shifts; a seasonal ARIMA, which admits a single seasonality, does not suffice. De Livera et al. (2011) address this with TBATS, which stabilizes the variance through a Box–Cox transformation (Box & Cox, 1964), represents each seasonality with time-varying trigonometric terms and corrects the residual with an ARMA structure; it calibrates automatically, but is slow to react to breaks. Taylor and Letham (2018) target that weakness with Prophet, whose changepoint-based trend allows breaks to be specified explicitly, which is valuable a priori in a series shaped by 2008 and 2020. One might expect this capability to be decisive; the paper tests it by comparing the two models out of sample, through rolling-origin cross-validation at the one-step horizon, using MAE, RMSE, MAPE and the scale-free MASE (Hyndman & Athanasopoulos, 2021) and the Diebold–Mariano test (Diebold & Mariano, 1995). As documented in the results, the conclusion qualifies that intuition.

Analyzing the cycle at a single time scale wastes information: monthly data capture the business cycle and annual seasonality, while daily data capture the within-month distribution and the build-up towards month-end. Separating the two scales is supported in the decomposition literature: namely into STL (Cleveland et al., 1990), robust to outliers, and its multi-seasonal extension MSTL (Bandara et al., 2025). For seasonal adjustment, however, this paper adopts the statistical-office standard, X-13ARIMA-SEATS (Findley et al., 1998), applied identically to registrations and to the reference indicator, and isolates the cyclical signal with a band-pass filter.

The second problem is temporal disaggregation: moving from low to high frequency so that the disaggregated data sum exactly to the original total. The Spanish statistical office (INE) builds the quarterly national accounts this way, using Chow and Lin (1971) and Denton (1971). The former distributes the total by regression on a high-frequency indicator with AR(1) errors; Fernández (1981) and Litterman (1983) cover variants without a common trend. Without an auxiliary indicator, Boot et al. (1967) and Denton (1971) distribute the total by minimizing distortion of its shape; Denton's proportional variant is a demanding benchmark, since by construction it always respects the total. Dagum and Cholette (2006) systematize the tradition, including Denton–Cholette; Sax and Steiner (2013) implement it in the tempdisagg package, and Mosley et al. (2022) add a sparse variant. The closest Spanish antecedent, Cuevas et al. (2019, 2021), seasonally adjust the daily sales series of the tax agency, but work on the already-adjusted series rather than on the daily distribution of each month's total.

2.4. Gaps in the Literature and Motivation for the Study

Two under-explored lines follow from the above. The first is methodological: TBATS and Prophet have been compared in general settings, but seldom on a Spanish sectoral series hit by recent shocks, where the ability to detect breaks should be decisive; this paper provides that comparison, out of sample and at the one-step horizon, and its conclusion qualifies the initial intuition. The second, sharper, concerns disaggregation: the methods reviewed presuppose a stable linear relationship with an indicator and do not explicitly model the within-month distribution, nor do they reuse the same month's pattern from the previous year. The distribution model (MD) fills this gap. It inherits from the disaggregation tradition the constraint that the days sum to the total, but changes the object estimated: it learns the distribution over the simplex of working days rather than interpolating the level against an indicator. This combination has no direct precedent in the literature reviewed.

2.5. Theoretical Framework

The proposed framework builds on three conceptual foundations. The first is the family of models that represent a time series as a sum of unobserved components evolving over time, on which both TBATS and Prophet are built, as are the tools for extracting the cycle. Two notions of cycle are used: the classical cycle, the peaks and troughs in the level; and the growth cycle, the oscillation isolated by the Christiano–Fitzgerald band-pass filter (Christiano & Fitzgerald, 2003) in the 1.5-to-8-year band, with turning points dated by the Bry–Boschan algorithm (Bry & Boschan, 1971) in its Harding–Pagan version (Harding & Pagan, 2002). This signal separates the cyclical lead from high-frequency co-movement. The second foundation is the theory of temporal disaggregation, whose rule is that the estimated quantities sum to the total the MD inherits, even though it estimates the daily distribution rather than the level. The third is the scoring rule: because the MD predicts a distribution rather than a quantity, its error is measured by cross-entropy (Shannon, 1948), which penalizes errors more heavily on days of greater registration concentration and is the same function used to train models that predict probabilities (Goodfellow et al., 2016).

3. Data and Methodology

The paper addresses three questions, each stated as a falsifiable expectation and each answered in Section 4.

H1. *Vehicle registrations lead Spanish industrial activity. Operationally, the cross-correlation between the cyclical components peaks at a negative lag, and lagged registrations reduce the out-of-sample forecast error of industrial production relative to an autoregressive benchmark. Economic motivation is one of timeliness: the industrial production index is published with a lag of several weeks, whereas registrations are administratively available within days, so any predictive content they carry is available earlier than the official series it anticipates.*

H2. *Because the Spanish registration series contains pronounced structural breaks, notably the 2008 collapse, the scrappage schemes and the 2020 lockdown, a forecasting model with an explicit changepoint mechanism should outperform one without. Prophet is such a model and TBATS is not.*

H3. *Because the intra-month distribution of registrations is shaped by several interacting regularities, a model with greater functional flexibility should allocate the monthly total across days more accurately than a low-dimensional parametric one. The neural variant of the distribution model is such a model.*

3.1. Data

The empirical analysis draws on three blocks of data: the vehicle-registration series that is the object of study, the macroeconomic series that define the reference business cycle, and a set of independent sources used to validate both the registration data and the resulting leading-indicator properties. All series are public. The replication package listed in the data availability statement contains the code and the aggregated series required to reproduce the results.

3.1.1. Vehicle-Registration Data

The registration data come from two sources, chosen for the distinct demands of each block. The cyclical and validation analyses use the OECD/FRED passenger-car series (ESPSLRTCR03MLM, 1990M01–2023M12), which offers the longest homogeneous monthly history of new car registrations and permits international comparison; pre-1990 non-integer back-estimates are discarded, and the series was discontinued in 2023. The monthly-forecasting block instead uses the DGT vehicle register for 2021–2025, which is current where the OECD series ends and provides the microdata underlying the daily block. The two series differ in coverage as well as in source: the OECD series counts passenger cars only, whereas the DGT extract used here comprises registrations of all vehicle categories, since the DGT microdata carry no field distinguishing vehicle type. The DGT series therefore stands at roughly 1.68 times the level of car-only registrations, a ratio that is stable over the sample (coefficient of variation 5 percent). Because the forecasting comparison is scale-free, this level difference does not affect the ranking of models, but the DGT-based series is described throughout as vehicle registrations rather than as passenger cars. The intra-month block (Section 3.5) requires daily data; because the DGT removes its daily files after publication, only two consecutive months (February–March 2026) of DGT daily data are available. For statistically robust evaluation we therefore add an external proxy, a daily series of new-vehicle-financing contracts from an automotive-finance institution, 2010–2025, which, being formalized at registration, shares the calendar structure (day-of-week, holidays, month-end) the MD model targets; as it is confidential, it is used on normalized weights only.

3.1.2. Reference-Cycle and Macro Data

We locate registrations against the cycle using two INE references. At monthly frequency, we use the industrial production index (IPI; total industry, unadjusted index, base 2021, 1975M01–2026M04), the standard monthly coincident indicator, whose frequency lets the lead be measured in months; we seasonally adjust it ourselves (Section 3.1.4) so that the registrations and reference receive identical treatment. At quarterly frequency, real GDP at market prices (seasonally and calendar-adjusted, base 2020, 1995Q1–2026Q1) is used to date the reference cycle and corroborate the IPI-based lead; INE builds these quarterly accounts by Chow and Lin (1971) and Denton (1971) temporal disaggregation, a benchmarking tradition to which our intra-month model (Section 3.5) is related but from which it departs in the object it estimates. As the authoritative external chronology, we adopt the recession datings of the Spanish Economic Association's Business Cycle Dating Committee, against which our own datings (Section 3.3) are compared.

3.1.3. External Validation Data

Two sources serve only to validate. ANFAC annual new-registration totals provide an administrative cross-check of the OECD/FRED series: aggregated to annual frequency, our monthly series reproduces them to within rounding (e.g., ≈ 1.61 M in 2007, ≈ 0.95 M in 2023) and records idiosyncratic events such as the April 2020 lockdown trough of 4163.

3.1.4. Data Construction and Treatment

All series enter in natural logs (proportional interpretation of components and growth). For the cyclical analysis, we align registrations with the IPI over their common 1990–2023 span, and with GDP over 1995–2023 (registrations summed to quarterly); the OECD series ends at 2023M12. Registrations and the IPI are seasonally adjusted identically with X-13ARIMA-SEATS, with additive outliers for the lockdown months (2020M03–M06) so the shock does not distort the seasonal factors; GDP enters in INE's already-adjusted form. For the daily block we construct, for each month, a day-of-week and working/non-working calendar incorporating Spain's national holidays, which the MD model uses to shape the within-month profile.

3.1.5. Descriptive Characterization

Registrations are markedly procyclical and far more volatile than output: over 1990–2023 the series peaks around 2007, roughly halves through the 2008–2013 contraction, recovers in the late 2010s, and collapses in 2020 (an April trough of 4163); its cyclical amplitude is several times that of the IPI or GDP, as expected of a big-ticket, postponable purchase that amplifies the cycle. It shows strong, stable within-year seasonality (Figure 1, raw vs. adjusted), which motivates the adjustment of Section 3.1.4. At daily frequency it concentrates sharply by day of week (working vs. non-working) and toward month-end, when dealers and fleets register before the monthly close the regularity the MD model exploits (Section 3.5). Two features need care: the COVID-19 lockdown months are extreme outliers (Section 3.1.4), and the 2018 spike reflects the WLTP regulatory transition (a pull-forward), not the cycle, so it is not read as a turning point.




Figure 1. Monthly passenger-car registrations, original and seasonally adjusted (X-13ARIMA-SEATS).

3.2. Conceptual Framework and Empirical Strategy

3.2.1. Operational Definitions

In the cyclical analysis (Section 3.3), registrations denotes newly registered passenger cars (the OECD series); in the monthly forecasting (Section 3.4) and the intra-month distribution (Section 3.5) it denotes newly registered vehicles of all categories (the DGT series), as set out in Section 3.1.1. We operationalize “the cycle” two ways: the classical cycle, the peaks and troughs in the level of the reference series; and the *deviation (growth) cycle*, the band-pass component of the log series isolated by the Christiano–Fitzgerald filter over the 18–96-month (1.5–8-year) band. A series is a *leading indicator* if its deviation cycle systematically precedes the reference’s; the *lead* is the displacement at which the cross-correlation between the two deviation cycles is maximized (positive when registrations move first). We distinguish this from *high-frequency co-movement*, the contemporaneous correlation of the two series’ prewhitened innovations: the lead asks whether the cycle precedes, co-movement only whether monthly shocks are synchronized, and the two need not coincide. Finally, the *intra-month distribution* is the vector of non-negative daily shares summing to one, the object the MD model forecasts (Section 3.5).

3.2.2. Rationale for the Multiscale Design

Registrations carry decision-relevant information at three time scales, and because each occupies a different frequency band, a single-frequency analysis cannot recover them jointly. At business-cycle frequencies (Section 3.3), registrations anticipate the macroeconomy, a signal invisible day to day, where calendar and seasonal variation dominates. At the monthly horizon (Section 3.4) the operational question is next month’s total. Within the month (Section 3.5), registrations follow a sharp working-day and end-of-month pattern, so the relevant object is how that total spreads across days, which lets the month-end figure be nowcast from partial-month data. The blocks are thus nested rather than independent: the cyclical analysis establishes registrations as a signal worth forecasting; the monthly model supplies the total; the intra-month model disaggregates and updates it as the month unfolds. Isolating each band filtering at cyclical frequencies and modelling the calendar at daily ones are what let each question be answered on its own terms.

3.3. Method 1: Cyclical Signal and Leading-Indicator Properties

The OECD composite leading indicator routinely includes national registration series among its components, but its headline lead of six to nine months refers to the composite rather than to registrations alone. Quantifying the lead of the registration component on its own, for Spain, is one of the contributions of this paper. We do not attempt a direct comparison against the composite, which would require the real-time vintages of its remaining components; that is noted as an extension. From the seasonally adjusted log series (Section 3.1.4) we extract each deviation cycle with the Christiano–Fitzgerald band-pass filter over the 18–96-month band (6–32 quarters for GDP), in its random-walk variant (appropriate to the $I(1)$ series), discarding the 36 months (8 quarters) at each end where the filter is least reliable.

We establish the lead two ways, kept distinct per Section 3.2.1. The cyclical lead is the cross-correlation between the registration and reference deviation cycles, computed without prewhitening, since the band-pass filter has already restricted attention to the cyclical frequencies and prewhitening would remove the very cyclical persistence whose alignment the exercise is designed to measure; the lead is the peak’s location. The quantity of interest is the location of the peak, which is the argmax of the cross-correlation over lags, and not the significance of the correlation at any individual lag. Because filtered components are strongly autocorrelated, the conventional band at plus or minus two over the square root of n does not apply; but even if it did, it would address the wrong hypothesis, since it tests whether a given correlation differs from zero rather than how precisely the peak is located. We therefore quantify the sampling variability of the argmax directly, through the specification grid and the moving-block bootstrap set out in Appendix A.3. Both resample the bivariate series jointly, so that the cross-series structure is preserved, and do so before filtering, so that the interval reflects the variability of the estimator rather than the dependence induced by the filter itself. This is reported against the IPI (monthly, fine resolution) and GDP (quarterly), on the pre-COVID-19 sample as baseline and the full sample for robustness (the lockdown being an extreme outlier).

Separately, high-frequency co-movement is measured on prewhitened growth innovations, isolating whether short-run shocks are synchronized. The inputs are the log differences of the seasonally adjusted series at their native frequency, monthly against the IPI and quarterly against GDP, with registrations summed over the three months. Following the Box–Jenkins procedure (Box & Jenkins, 1970), an autoregression is fitted to the registration series alone, with order selected by AIC over a maximum of twenty-four lags, and the same filter is then applied to both series before the cross-correlation is computed. The selected orders are eight for the monthly pairing over the full sample and two for the quarterly pairing and for the pre-COVID-19 monthly subsample; residual autocorrelation in the filtered registration series is not distinguishable from zero in any case (Ljung–Box test, twelve lags, p between 0.45 and 0.79; Ljung & Box, 1978). Unlike the cyclical lead, this correlation is computed on stationary innovations, so the conventional band does apply here.

We corroborate the lead by dating turning points with the Bry–Boschan algorithm in its quarterly (Harding–Pagan) form on the level of GDP for the classical cycle (checked against the Spanish Economic Association’s chronology) and on the deviation cycles for the growth cycle and measuring the average peak-to-peak and trough-to-trough displacement between the registration and reference chronologies (excluding 2020).

3.4. Method 2: Monthly Forecasting

The target is the monthly series of vehicle registrations. Two forecasting models are compared out of sample, with no prior preference: Prophet, a decomposable model

pairing a piecewise-linear trend with automatically detected changepoints, a multiplicative annual seasonal term (its amplitude scaling with the level of the series) and Spanish national-holiday effects, the weekly term disabled (the series being monthly); and TBATS, an exponential-smoothing state-space model that selects its Fourier seasonal harmonics automatically, stabilizes the variance with a Box–Cox transformation and carries ARMA errors. Both are evaluated on the DGT panel (2021–2025, which starts in 2021 to exclude the COVID-19 structural break) and, for statistical power, on the long OECD passenger-car series (1990–2023, Section 3.1.1), which spans several complete cycles; the protocol is rolling-origin cross-validation (Section 3.6). The more accurate model supplies the monthly total F that the MD model disaggregates (Section 3.5).

3.5. Method 3: Intra-Month Distribution/Nowcasting: The MD Model

Having shown that registrations lead the cycle, the value of that signal lies in its timeliness, which is the motivation for modelling the series daily rather than monthly. Once the within-month distribution is known, the month-end total can be estimated before the month closes as the observed cumulative count divided by the predicted share of the elapsed working days. This reads the leading signal weeks ahead of the official figure and updates it continuously rather than in monthly steps. The daily model is thus the operational counterpart of the leading-indicator result, and it also underpins a methodological correction, since the strong working-day and end-of-month effects can only be removed once the within-month distribution is modelled. What Section 4.3 validates is the accuracy of that distribution rather than a full nowcasting exercise, the two being linked; and because the DGT does not archive its daily files reliably, the validation uses a daily financing proxy (Section 3.1). Given the selected monthly forecast F for month i (TBATS, Section 4.2), the MD (Modelo de Distribución) model estimates a non-negative daily weight $\hat{\rho}_d$ for each working day, so that $\hat{y}_d = F \cdot \hat{\rho}_d$. It disaggregates the monthly total to days, but distinctively: where the regression-based methods of Section 3.1.2 (Chow–Lin, Denton) interpolate a level series against a related indicator, the MD learns the within-month distribution, itself a weight vector over the simplex of working days, calibrated by cross-entropy.

$$S_d = \left[(1 - w_i) q_d + w_i \rho_d^{y-1} \right] \{holiday_d\} \quad (1)$$

where $\hat{\rho}_d = \frac{S_d}{\sum_{d=1}^D S_d}$. If the day d is a holiday, then $\{holiday_d\} = 0$; otherwise, $\{holiday_d\} = 1$. A structural calendar component q_d and the historical weight ρ_d^{y-1} of the same day one year earlier are mixed by a per-month weight w_i . The holiday indicator, applied on the target year's calendar, zeroes holidays removing the prior year's holiday-inheritance bias. As $w_i \rightarrow 1$ the model reduces to the historical baseline; as $w_i \rightarrow 0$ it relies on the structural term.

The structural component is a logistic of calendar regressors, $q_d = \sigma(\alpha \beta_{dow}) = \frac{1}{1 + e^{-(\alpha \beta_{dow})}}$, with β_{dow} being the day-of-week log-effect (each weekday's share relative to $1/7$). The structural term uses the day-of-week effect; this is the specification we retain and denote MD-shallow. The general form below (Equation (2)) admits further calendar regressors. We explored an end-of-month position term and a first/second-fortnight indicator, neither of which improved out-of-sample accuracy, and MD-deep replaces the single logistic with a hidden-layer network.

The parameters (w_i and the α 's) are calibrated by minimizing the categorical cross-entropy between predicted and realized normalized weights, via L-BFGS-B subject to $w \in [0.2, 0.95]$, $\alpha \in [0.1, 20]$. Calibration is a sliding window: the reference month is the same calendar month of the previous year, which supplies both β_{dow} and $\rho_d^{(y-1)}$, and the target month is held out. Day-of-week intensities are computed on that reference month

as the share of units registered on each weekday over its total. The historical weights are carried across to the target month by working-day position: the n -th working day of the target month receives the weight of the n -th working day of the reference month, and when the target month contains more working days than the reference month the surplus days receive the mean weight of the reference month's working days. Non-working days receive a score of zero by construction, so the cross-entropy is evaluated over working days alone. In use, the predicted daily path is a real-time, calendar-correct benchmark for the in-progress month (Section 3.2.2).

In its general form the daily score combines N additive calendar regressors with p multiplicative campaign/holiday components, preserving the logistic structure and the day-level normalization:

$$s_d^{(y)} = \left[(1 - w_i) \sigma \left(\sum_{j=1}^N \alpha_j \beta_j(i, d, y) \right) + w_i \rho_d^{(y-1)} \right] \prod_{k=1}^p \gamma_k \quad (2)$$

where $\hat{\rho}_d = \frac{s_d^{(y)}}{\sum_{d=1}^D s_d^{(y)}}$. The structural coefficients α_j and the historical weight w_i are calibrated by minimizing the categorical cross-entropy with L-BFGS-B. The objective is differentiable in every parameter: because $\hat{\rho}_d$ is a normalization of $s_d^{(y)}$, the gradient simplifies to

$$\frac{\partial J}{\partial \theta} = - \sum_{d=1}^D \frac{\rho_d}{\hat{\rho}_d} \frac{\partial \hat{\rho}_d}{\partial \theta}, \quad \frac{\partial \hat{\rho}_d}{\partial \theta} = + \frac{1}{\sum s_d^{(y)}} \left(\frac{\partial S_d}{\partial \theta} - \hat{\rho}_d \sum \frac{\partial S_d}{\partial \theta} \right) \quad (3)$$

where the partial derivatives of S_d are defined as (4) and (5):

$$\frac{\partial S_d^{(y)}}{\partial \alpha_j} = (1 - w_i) \sigma \left(\sum_{j=1}^N \alpha_j \beta_j \right) \left(1 - \sigma \left(\sum_{j=1}^N \alpha_j \beta_j \right) \right) \beta_j \quad (4)$$

and

$$\frac{\partial S_d^{(y)}}{\partial w_i} = \rho_d^{(y-1)} - \sigma \left(\sum_{j=1}^N \alpha_j \beta_j \right) \quad (5)$$

On non-holiday days $\prod_{k=1}^p \gamma_k = 1$. Thus, the gradient reduces to the two partials above; these supply the analytic gradient that L-BFGS-B uses to minimize J .

In this generalized form the MD model is a parametric learner with a sigmoid activation, a simple-normalized cross-entropy objective and gradient-based calibration, structurally a single-layer neural network adapted to the intra-month allocation problem, with MD-deep adding one hidden layer. MD-deep replaces the single logistic of the structural component with a two-layer feedforward network. Its inputs are three calendar features per day: the day-of-week log-effect, the position within the month and a fortnight indicator. A hidden layer applies a logistic activation to H units, with $H \in \{2, 3, 4, 5\}$, and the output layer a further logistic to their weighted sum; scores are normalized over the simplex of working days as in MD-shallow, and w_i is estimated jointly with the network parameters. Calibration is by L-BFGS-B under box constraints of ± 10 on the network weights and $[0.2, 0.95]$ on w_i , with 1000 iterations and ten restarts per architecture, seeded at $42 + i$, retaining the restart with the lowest training cross-entropy. We report this in full because the negative result in Section 4.3 requires the failing model to have been estimated seriously.

3.6. Evaluation Protocol

All predictive claims are assessed out of sample, respecting chronological order. For the monthly forecasts (Section 3.4), TBATS and Prophet are compared by rolling-origin

cross-validation: the origin advances one month over an expanding window and each model forecasts one step ahead, so accuracy is assessed at the one-month horizon. The procedure runs on the DGT panel (36 origins) and, for power, on the long OECD series (384 origins, 383 for Prophet). Accuracy is summarized by MAE, RMSE, MAPE and the scale-free MASE, and equal predictive accuracy is tested with the Diebold–Mariano statistic under the Harvey–Leybourne–Newbold small-sample correction (Harvey et al., 1997), under both squared and absolute loss. The more accurate model is selected. The MD model (Section 3.5) is evaluated on the normalized daily weight probability distributions over the month’s working days, for which categorical cross-entropy is the primary metric, complemented by the Kullback–Leibler divergence (Kullback & Leibler, 1951) over the realized entropy ($KL = CE - H$). The powered validation uses the financing proxy: a sliding-window scheme runs across its 168 month–year pairs (2012–2025), each calibrated on prior data only and held out, supplying the statistical power the two-month DGT span lacks; equal accuracy is again tested by the Diebold–Mariano statistic (HLN-corrected, with a Newey–West long-run variance; Newey & West, 1987) and a sign test. MD-shallow is benchmarked against the persistence baseline ($w = 1$), Prophet’s implied daily distribution and the neural MD-deep. The two available months of DGT daily data serve as an operational illustration (Section 4.3), not a benchmark.

Section 4.4 adds a predictive test of the leading property itself. Lag orders are re-selected by BIC at each origin from information available at that origin alone, with a maximum of twelve own lags and six lags of registrations; the origin advances one observation at a time over an expanding window; and equal predictive accuracy is tested by the Clark–West statistic (Clark & West, 2007), which is the appropriate test for nested models. Newey–West long-run variance is used, with a bandwidth that accommodates the overlap induced by the year-on-year transformation. The Diebold–Mariano statistic is reported alongside but is conservative in this setting, since under the null the unrestricted model estimates parameters that are zero in population and its mean squared prediction error is inflated accordingly. Three evaluation windows are reported: the full sample, a window excluding any forecast whose origin-to-target interval intersects March 2020 to December 2021, and a window restricted to origins before 2020.

4. Empirical Results

4.1. Result 1: Cyclical Signal and Leading-Indicator Properties

Vehicle registrations lead the Spanish business cycle. The cross-correlation of the registration and IPI deviation cycles peaks at a negative lag within a coherent hump (Figure 2), placing registrations three months ahead of industrial production on the pre-COVID-19 sample (peak correlation 0.69, $n = 324$), with a near-flat plateau over three to four months. Adding 2020 shortens the estimated lead to two months and raises the peak correlation to 0.75, as the shared pandemic shock pulls the two series into closer contemporaneous alignment; the pre-COVID-19 figure is therefore taken as the baseline. Note that the filtered sample ends in December 2020 rather than December 2023, since the thirty-six-month edge trimming required by the band-pass filter discards the final three years, so 2020 is the only period distinguishing the two samples.




Figure 2. Cross-correlation of the cyclical components of vehicle registrations and industrial production, pre-COVID-19 sample (1993M01 to 2019M12, $n = 324$). The conventional significance band is not reported, since it does not apply to filtered components; the sampling variability of the peak location is quantified in Appendix A.3.

Because the lead is the argmax of a cross-correlation over lags, a discrete statistic with no analytical standard error, we quantify its sampling variability directly (Appendix A.3). Across a grid of filter bands, filter families and trimming lengths, the pre-COVID-19 lead ranges from four to six months and the full-sample lead from one to four, the difference between samples exceeding the dispersion within either. A moving-block bootstrap on the pre-COVID-19 sample, re-executing the full filtering and cross-correlation pipeline on each of two thousand replicates, returns a modal lead of four months with a ninety per cent interval covering lags one to six. The estimate is therefore three to five months depending on the treatment of the sample edge, imprecisely determined but robust in sign, and we report it as a range rather than as a point.

Turning-point dating corroborates this. On the growth cycle, the correspondence between registration and GDP turning points is positive in direction but cannot be resolved finely at quarterly frequency, where the minimum measurable displacement is one quarter. Of the twelve matched pairs, six are exactly contemporaneous, the median displacement is zero and the mean is 2.8 months, the latter driven by two episodes of fifteen and eighteen months. The eighteen-month value coincides exactly with the maximum displacement admitted by the matching rule and should be read as a bound rather than as an estimate. The individual displacements are reported in Appendix A.3. We therefore treat the turning-point evidence as corroborating the direction of the lead established against industrial production, not as an independent measurement of its length. Against quarterly GDP the cross-correlation peaks at one to two quarters (0.67), consistent in direction but coarser in resolution. At the 2008 classical cycle peak the lead reaches eighteen months. The classical datings agree with the chronology of the Spanish Economic Association's Dating Committee.

A prewhitened cross-correlation, which strips out the cyclical persistence where the lead resides, shows no anticipation. The peak lies at lag zero in all three pairings examined,

and the contemporaneous correlation before 2020 is weak, at 0.22 against industrial production, rising to 0.67 over the full sample and to 0.72 against quarterly GDP. The strength of the contemporaneous association therefore depends on the presence of large shared shocks such as the 2020 lockdown rather than on any systematic high-frequency synchronization. The lead is thus a property of the cycle, not of high-frequency shocks; the two must not be conflated.

The estimated lead of three to five months is specific to registrations and to Spain; the OECD's six-to-nine-month figure refers to its composite leading indicator, of which registrations are only one component.

Section 4.4 confirms that this leading property is not an artefact of the two-sided filter used to isolate it, although the resulting gain in forecast accuracy is modest.

4.2. Result 2: Monthly Forecasting

We compare TBATS and Prophet out of sample by rolling-origin cross-validation (Tashman, 2000): the origin advances one month at a time over an expanding window and forecasts one step ahead, so accuracy is measured at the one-month horizon relevant to the daily-disaggregation model (Section 4.3) and to near-term monitoring. The protocol is run on two series: the recent DGT panel (2021–2025; 36 one-step forecasts over 2023–2025) and, for statistical power, the long OECD series (1990–2023; 384 one-step forecasts over 1992–2023 after a two-year initialization, of which 383 are available for Prophet, whose estimation failed to converge at the January 1995 origin), which spans several complete cycles. Accuracy is summarized by MAE, RMSE, MAPE and the scale-free MASE, and equal predictive accuracy is tested with the Diebold–Mariano statistic under the Harvey–Leybourne–Newbold small-sample correction, under both squared and absolute loss.

TBATS is the more accurate model on both series (Table 1).

Table 1. Monthly series evaluation metrics.

Series	Origins	Model	MAE	RMSE	MAPE	MASE
DGT 2021–2025	36	TBATS	10,080	12,576	7.1%	0.67
DGT 2021–2025	36	Prophet	13,015	18,654	9.4%	0.87
OECD 1990–2023	384	TBATS	8452	11,999	14.1%	0.61
OECD 1990–2023	383	Prophet	14,184	20,103	24.5%	1.02

On the recent DGT panel it wins on every metric (MAPE 7.1% vs. 9.4%, MASE 0.67 vs. 0.87; one-step error path in Figure 3), though with only 36 one-step forecasts the Diebold–Mariano test does not reach significance ($p = 0.09$ – 0.18): the point estimates favor TBATS, but the sample is too short to certify the difference. The test is computed on the 383 origins for which both models produced a forecast.

The long OECD series settles it. TBATS beats the seasonal-naïve benchmark (MASE 0.61) where Prophet fails to (MASE 1.02), and the Diebold–Mariano test rejects equal accuracy overwhelmingly: $DM^* = 6.3$ ($p \approx 7 \times 10^{-10}$, squared loss) and $DM^* = 8.8$ ($p \approx 7 \times 10^{-17}$, absolute loss), both in TBATS's favor (per-year error paths in Figure 4). A one-step error of 7.1% on the recent series is small enough to be useful for near-term operational planning.




Figure 3. One-step rolling-origin absolute percentage error by month, DGT series, TBATS vs. Prophet (36 one-step forecasts, 2023–2025; the first 24 months initialize the expanding window).




Figure 4. Mean one-step rolling-origin absolute percentage error by year, OECD passenger-car series, TBATS vs. Prophet (384 one-step forecasts for TBATS and 383 for Prophet, 1992–2023; log scale—both models spike in 2020 under the COVID-19 shock).

The single 2025 hold-out and the rolling-origin design answer different questions. The hold-out trains once in 2021–2024 and forecasts the whole of 2025 at once in a one-to-twelve-month-ahead exercise, and on that year-ahead task Prophet is more accurate. Rolling-origin instead forecasts one month ahead at each origin; on that horizon TBATS is more accurate, and by a wide margin on the long series. TBATS's weak 2025 hold-out reflects a large multi-step miss, not poor short-horizon forecasting: its one-step errors are roughly half its year-ahead ones. The two findings do not conflict; they describe different horizons, and the near-term one is both the relevant and the only rigorously tested of the two.

The result is consistent with the models' structure. TBATS represents the annual cycle and its six-month harmonic with trigonometric terms under a Box–Cox stabilization, which suits a long history spanning several expansions and contractions; Prophet's piecewise-linear trend with changepoints and a single annual term is more reactive to recent breaks but generalizes less well over the full sample, where it does not improve on the seasonal-naïve. H2 is thus rejected: the changepoint mechanism does not pay for itself at the one-step horizon for which the comparison was designed, although Prophet remains competitive

in the year-ahead hold-out. TBATS is therefore selected as the monthly model, and its one-step forecast supplies the monthly total disaggregated by the MD model (Section 4.3).

4.3. Result 3: Intra-Month Distribution

We validate the MD model out of sample on 168 month–year pairs spanning 2012–2025, using a daily vehicle-financing series as a confidentiality-protected proxy for daily registrations. For each month the model is calibrated on prior data only and predicts that month’s daily shares, which are then rescaled to the month’s realized total; the evaluation therefore measures the quality of the daily distribution in isolation from the monthly forecast and is independent of which monthly model supplies that total. Accuracy is the categorical cross-entropy (CE) between predicted and realized working-day shares, together with its excess over the irreducible entropy of the realized distribution, the Kullback–Leibler divergence, $KL = CE - H$. The single-layer MD (MD-shallow) is compared against a deeper variant with one hidden layer (MD-deep), Prophet’s implied daily distribution, and a persistence baseline that reuses the prior-year same-month pattern (the limit $w = 1$).

MD-shallow is the most accurate model, with the lowest mean CE (2.98) and KL (0.10), ahead of Prophet (3.05/0.17), MD-deep (3.07/0.18) and the persistence baseline (3.29/0.41) (Table 2). Its advantage over each competitor is significant under both a sign test and the Diebold–Mariano HLN-corrected test, with a Newey–West long-run variance to absorb serial correlation in the loss differential. Two findings stand out. First, the hidden layer does not help: MD-deep is less accurate out of sample, winning only a third of the months, so its extra parameters do not generalize, consistent with overfitting and a pattern an in-sample fit cannot reveal. Second, MD-shallow improves on the persistence baseline by a wide average margin ($\Delta CE = 0.31$, DM $p < 10^{-5}$) while winning 62% of months: because the optimal mixing weight is high, the model leans heavily on the prior-year pattern, yet the small structural correction it adds is a genuine, significant gain over mere repetition. Against Prophet the margin in mean cross-entropy is small: Prophet alone accounts for roughly three quarters of the total improvement over the persistence baseline. What distinguishes MD-shallow is the consistency rather than the size of its advantage (Figure 5), which holds in 154 of the 168 out-of-sample months (92%). H3 is therefore rejected: the additional flexibility of the neural variant does not improve out-of-sample distributional accuracy, and in fact worsens it.

Table 2. Distributional accuracy over the financing-proxy test (working-day weights).

Model	Mean CE	Mean KL	DM* vs. MD-Shallow	P (DM)	MD-Shallow Wins	P (Sign)
MD-Shallow	2.98	0.10	—	—	—	—
Prophet	3.05	0.17	8.07	1.3×10^{-13}	154/168 (92%)	5.5×10^{-31}
MD-deep	3.07	0.18	5.48	1.5×10^{-7}	112/168 (67%)	1.9×10^{-5}
Baseline ($w = 1$), equivalent to proportional Denton	3.29	0.41	4.97	1.7×10^{-6}	104/168 (62%)	2.5×10^{-3}

Note: With a single monthly total as a constraint, the proportional Denton benchmark in first differences is algebraically identical to the persistence baseline. See Section 4.3.

The proportional Denton benchmark appears in Table 2 as the persistence baseline, with which it coincides exactly. This is an identity rather than a coincidence. Denton in first differences minimizes the sum of squared changes in the ratio of the target to the indicator subject to the target summing to the known total; that objective attains its lower bound of zero whenever the ratio is constant, and with a single constraint the constant is pinned down by the total alone, so the solution is the indicator profile rescaled to the total. When the indicator is the reference-month profile, which it must be for Denton to face the same information set as the other models, the result is precisely the persistence baseline. The two coincide to machine precision.





Figure 5. Difference in cross-entropy between MD-shallow and the persistence baseline across 168 month–year pairs. Negative values, in green, indicate that MD-shallow is more accurate. Gains concentrate in March, April and December; the March and April gains recur across the sample, whereas the December gains are concentrated in 2014, 2015, 2017 and 2018.

This degeneracy is itself informative. The canonical benchmark for temporal disaggregation collapses, on this problem, into the trivial model, which is why the intra-month allocation problem cannot be treated as one of bookkeeping disaggregation. The benchmark becomes non-degenerate only when several monthly totals are imposed jointly, so that the smoothness objective binds across month boundaries. We examined that variant over spans of two and twelve months; it performs worse (mean cross-entropy 3.35 and 3.37, respectively), and, since Denton imposes no non-negativity constraint, it returns negative daily values in three of the 168 months, which are not admissible proportions. We therefore retain the single-constraint form.

Finally, the model was applied to the two months of daily registration data released by the DGT, the operational target which it tracks closely (Figure 6). With only a single calibration month available, this is a deployment illustration rather than a further benchmark; the powered, out-of-sample evidence for the model is the proxy validation above.

As an illustration of the timeliness the daily model affords, the MD nowcasts the month-end total from partial-month data using its predicted shares. By mid-month it estimates the total to within $\approx 23\%$ on the proxy, against $\approx 36\%$ for naïve linear extrapolation; the MD's advantage is largest in the first half of the month, where the within-month pattern is most informative ($\approx 32\%$ vs. $\approx 48\%$ median error) (Figure 7), and both converge as the month completes. The absolute error reflects the noisiness of the daily proxy, but the consistent gain over linear extrapolation confirms that the modelled intra-month distribution carries real predictive content.



Source: own elaboration

Figure 6. MD daily prediction versus realized registrations, DGT data, March 2026.




Figure 7. Intra-month nowcasting error versus elapsed share of the month (median APE over 168 proxy months): MD-predicted shares versus linear extrapolation; shaded bands, interquartile range. Total refers to working-day registrations.

4.4. Result 4: Out-of-Sample Predictive Value

The cyclical analysis in Section 4.1 establishes that registrations lead the reference cycle, but they do so on a band-pass component extracted with a two-sided filter, which is not available at the end of the sample where a leading indicator would actually be used. This section therefore poses the operational question directly: does knowing past registrations reduce the error of a forecast of industrial production made in real time?

The target is the growth of the industrial production index over horizons of one, three, six and twelve months. A restricted benchmark regresses that growth on its own lags; the unrestricted model adds lags of registrations. Both series enter raw; registrations enter as the OECD passenger-car series of Section 3.1.1, whose 1990 start the exercise requires. The principal specification uses year-on-year log differences, so that no two-sided filter and

no seasonal adjustment procedure enters the exercise at any stage. The principal cell, the three-month horizon under the evaluation window excluding March 2020 to December 2021, was fixed before the results were examined; the full grid of forty-two cells is reported in the replication package.

In that cell, registrations reduce the mean squared prediction error by 1.9 per cent over 235 origins, and equal predictive accuracy is rejected (Clark–West 2.023, one-sided $p = 0.022$). The gain is not an artefact of an under-specified benchmark: with a maximum lag order of twelve, the selected order reaches that ceiling in none of the origins. Table A4 reports the four horizons under the three evaluation windows. The evidence is somewhat more consistent at six and twelve months ($p = 0.009$ and 0.013 excluding COVID-19, 0.010 and 0.014 before 2020) than at three, where the pre-2020 window falls just short of conventional significance ($p = 0.053$).

The magnitude of the gain is small. A reduction of one to two per cent in mean squared error is statistically detectable but of limited operational consequence, and we describe it as such. What the exercise establishes is that the leading property survives the removal of the two-sided filter, not that registrations transform the forecasting of industrial production.

Three findings run against the paper’s thesis and are reported here rather than relegated to a footnote. First, when the COVID-19 period is included in the evaluation window, registrations worsen the forecast, in the worst cells by more than half of the mean squared error. In April 2020 registrations fell by approximately ninety-five per cent against a fall of thirty-four per cent in industrial production; the approximate proportionality that holds over ordinary cyclical fluctuations does not survive an extreme non-cyclical shock. Registrations are informative about the timing of turning points, not about the magnitude of exceptional events. Second, no consistent predictive gain is found for quarterly GDP. Of eighteen cells, three reach the five per cent level, with no consistency across evaluation windows and one of them recording a mean squared error ratio above unity; the number of origins ranges from thirty-eight to sixty-three, so the exercise lacks power at quarterly frequency. We therefore make no predictive claim for GDP, and the displacement reported in Section 4.1 is a chronological result rather than a forecasting one. Third, an exploratory reverse test in which industrial production is used to predict registrations also rejects equal accuracy (one-sided $p = 0.031$), so predictability is bidirectional. What remains is an asymmetry of magnitude, the reduction in mean squared error being approximately five times larger in the proposed direction, 1.9 against 0.4 per cent. Mutual predictability is the expected outcome for two series sharing a common cycle, and unidirectional causality is not a property claimed for leading indicators generally, nor for the registration components of composite indicators.

Finally, this exercise does not establish that registrations add information beyond indicators already available to a forecaster. Testing that would require comparable real-time series for consumer confidence and related variables, which we identify as the natural extension of the result. Taken together with the cyclical evidence in Section 4.1, this supports H1, with the qualification that the estimated lead is imprecise and shortens when the pandemic period is included, and that the real-time gain is modest.

5. Conclusions and Discussion

This study addressed three closely related research questions concerning the role of vehicle registrations as a Leading Economic Indicator, the forecasting of monthly registrations and the real-time estimation of their intra-month evolution. Together, these objectives provide an integrated framework linking business-cycle analysis, monthly forecasting and daily nowcasting, thereby extending the role of administrative vehicle-registration data beyond conventional demand forecasting.

First, the empirical results support the hypothesis that vehicle registrations contain useful leading information about the Spanish business cycle. Across several complementary methodologies (including cyclical cross-correlation analysis, turning-point detection, bootstrap validation and out-of-sample predictive evaluation), registrations are found to anticipate industrial production by approximately three to five months. Nevertheless, the results also indicate that the estimated lead depends on the sample period and filtering procedure, while the predictive gains for quarterly GDP are considerably weaker. Rather than suggesting a universal leading relationship, these findings indicate that vehicle registrations constitute a valuable high-frequency indicator whose predictive content is strongest for industrial activity and whose magnitude should be interpreted in light of the underlying sample and methodology.

Second, the comparison between alternative forecasting approaches shows that forecasting performance depends on both the characteristics of the data and the evaluation framework. Although TBATS consistently provides more accurate one-step-ahead forecasts and statistically outperforms Prophet in the long OECD sample, Prophet remains competitive under the year-ahead hold-out evaluation. Contrary to our initial hypothesis, these findings indicate that explicit changepoint modelling does not necessarily lead to superior forecasting performance in the presence of structural breaks. Instead, they suggest that forecasting accuracy is determined not only by the ability to model structural change but also by the forecasting horizon and the statistical properties of the series.

Third, the proposed machine learning-inspired distribution model demonstrates that monthly forecasts can be transformed into reliable daily nowcasts through a data-driven temporal disaggregation procedure. However, contrary to the initial expectation, increasing model flexibility through a deeper neural architecture does not improve out-of-sample performance. The simpler MD-shallow specification consistently provides a better balance between predictive accuracy, robustness and computational efficiency. This result illustrates an important practical lesson: greater model complexity does not necessarily translate into better predictive performance when the underlying data-generating process is relatively structured and stable.

The present findings both confirm and complement previous research on automotive demand forecasting and business-cycle analysis. Consistent with [Kim \(2023\)](#) and [Muth et al. \(2025\)](#), vehicle registrations contain valuable predictive information, although our results show that forecasting performance depends critically on the forecasting horizon and the characteristics of the underlying series. In contrast to previous studies, the out-of-sample evidence indicates that the classical TBATS model outperforms Prophet for one-step-ahead forecasting of Spanish vehicle registrations. Whereas [Ledo-Lobato et al. \(2024\)](#) focused on vehicle registrations as a predictor of automotive demand, the present study examines their broader macroeconomic role by assessing their ability to anticipate business-cycle fluctuations. This positioning can be compared with established composite leading indicators. [The Conference Board \(2026\)](#) reports that its Leading Economic Index for Spain, which does not include vehicle registrations among its components, anticipates business-cycle turning points by approximately three months. Likewise, the [OECD \(2020\)](#) states that its composite leading indicator, which incorporates new passenger-car registrations as one of several component series, is designed to anticipate turning points by approximately six to nine months. Although these benchmarks refer to composite indicators rather than vehicle registrations in isolation, our estimated lead of three to five months lies between them, being close to the Conference Board's estimate and toward the lower bound of the OECD benchmark. Our contribution is therefore twofold. First, we isolate vehicle registrations from broader composite indicators and show that, on their own, they constitute a robust standalone Leading Economic Indicator for monitoring the Spanish

business cycle, rather than simply one component within a broader index or a predictor confined to automotive demand. Second, we advance the literature methodologically by integrating business-cycle analysis, monthly forecasting and daily nowcasting within a unified multiscale framework, demonstrating how traditional Leading Economic Indicators and modern forecasting techniques can be combined to provide operational real-time monitoring of economic activity.

Taken together, these findings have important methodological and practical implications. From a methodological perspective, they demonstrate that combining business-cycle analysis, modern forecasting techniques and temporal disaggregation substantially enhances the usefulness of traditional Leading Economic Indicators. From a practical perspective, the proposed framework allows policymakers, manufacturers, dealers and financial institutions to obtain earlier and more accurate information on automotive demand than would be possible using official monthly statistics alone. The possibility of transforming monthly forecasts into continuously updated daily estimates also opens new opportunities for real-time economic monitoring.

Despite these contributions, several limitations should be acknowledged. First, the empirical analysis focuses exclusively on Spain, and therefore the estimated leading relationships may not necessarily generalize to countries with different institutional settings, consumer behavior or automotive market structures. Second, although the monthly forecasting exercise benefits from long historical data, the validation of the daily distribution model relies on a proxy based on vehicle-financing operations because historical daily registration data are not systematically archived by the Spanish Directorate-General for Traffic (DGT). Third, the forecasting comparison is restricted to TBATS and Prophet. While these models represent two of the most suitable alternatives for highly seasonal time series, recent deep-learning architectures and transformer-based forecasting models were beyond the scope of the present study. Fourth, the analysis concentrates on registration data alone and does not explicitly incorporate complementary high-frequency information such as Google Trends, consumer sentiment indicators, mobility data or financial variables, which could further improve forecasting performance. Fifth, the predictive test establishes that registrations improve on an autoregressive benchmark, but not that they add information beyond indicators already available to a forecaster, such as consumer confidence. Predictability is also found to run in both directions, which is expected between two series sharing a common cycle but means that the evidence supports informativeness and timeliness rather than unidirectional causality.

These limitations suggest several promising directions for future research. Extending the proposed framework to other countries would allow researchers to assess the robustness and international applicability of the estimated leading relationships. Likewise, the availability of longer daily registration series would facilitate the validation and further refinement of the proposed distribution model under different market conditions. Future studies could also compare the proposed methodology with state-of-the-art deep-learning approaches, including recurrent neural networks, Temporal Fusion Transformers and foundation models specifically developed for time-series forecasting. Finally, integrating registration data with digital indicators derived from internet searches, social media, online news or other Big Data sources could lead to the construction of a new generation of hybrid Leading Economic Indicators capable of improving both macroeconomic nowcasting and business-cycle forecasting.

Overall, this study demonstrates that passenger-car registrations remain a valuable and underexploited source of forward-looking information for monitoring economic activity. By integrating business-cycle analysis, monthly forecasting and daily nowcasting within a single multiscale framework, the proposed methodology transforms administrative

registration data into an operational high-frequency economic indicator. More broadly, the findings show that the value of traditional Leading Economic Indicators can be substantially enhanced when they are combined with modern forecasting techniques, providing both methodological advances and practical tools for real-time economic monitoring.

Author Contributions: R.G.: methodology, software, formal analysis, data curation, visualization, writing—original draft. R.P.: methodology (forecasting). C.P.: conceptualization, supervision, project administration, writing—review and editing. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The code and aggregated series required to reproduce the results are deposited at <https://doi.org/10.5281/zenodo.21607046>. Registrations for the cyclical analysis are the OECD/FRED series ESPSLRTCR03MLM; the industrial production index is INE Table 60272 (base 2021, original series) and quarterly GDP is INE Table 67824 (Quarterly National Accounts, chain-linked volume index, seasonally and calendar-adjusted); the 2021 to 2025 vehicle register extract is drawn from the DGT microdata portal. Access dates are recorded in the package README. The daily registration files published by the DGT are not archived at source and could not be deposited. The analysis in Section 4.3 uses a proprietary daily financing proxy which is not deposited: normalized daily shares are ratios of small integers from which monthly totals can be recovered by denominator search, so no transformation preserves both the analytical content of the series and the provider's confidentiality. The package contains instead the aggregate results on which Section 4.3 rests, together with the pre-computed outputs of every script that depends on non-redistributable data.

Conflicts of Interest: One of the authors was formerly an employee of the automotive-finance institution that provided the confidential daily vehicle-financing series used as a proxy in Section 4.3. The series was supplied for academic research purposes only; it is used exclusively in normalized form, and no raw data are shown, disclosed or deposited. The institution had no role in the design of the study; in the collection, analysis or interpretation of the data; in the writing of the manuscript; or in the decision to publish the results. The authors declare no other conflicts of interest.

Appendix A

This appendix collects supporting material referenced in the main text.

Appendix A.1. Forecast Horizon: Hold-Out Versus Rolling-Origin

The monthly comparison can be run either as a single hold-out training in 2021–2024 and forecasting all of 2025 at once or as rolling-origin cross-validation at the one-step-ahead horizon. The two answer different questions and rank the models differently. Under the hold-out, a one-to-twelve-month-ahead exercise, Prophet is the more accurate model (Table A1, Figure A1); under rolling-origin at the one-step horizon, TBATS is more accurate (Section 4.2). The reversal reflects the forecast horizon, not instability: TBATS's one-step errors are roughly half its year-ahead errors, so its weaker hold-out performance is multi-step degradation rather than poor short-horizon forecasting. The one-step horizon is the one relevant to the daily disaggregation and to near-term monitoring, and the only one tested with statistical power, which is why it is adopted in the main text.

Table A1. Single hold-out (train 2021–2024, forecast 2025): monthly forecast accuracy.

Model	MAE	RMSE	MAPE
TBATS	21,547.6	25,660.1	14.13%
Prophet	13,288.7	15,439.0	8.56%




Figure A1. Monthly registrations, actual versus TBATS and Prophet forecasts, single hold-out (train 2021–2024, test 2025). Prophet tracks the year-ahead path more closely; cf. the one-step results in Section 4.2.

Appendix A.2. MD Versus Prophet, Month by Month

The heatmap in the main text contrasts MD-shallow with the persistence baseline; the complementary contrast against Prophet is shown here. MD-shallow attains a lower cross-entropy in 154 of the 168 out-of-sample months (92%, Table A2; Figure A2).




Figure A2. Per-month difference, CE (MD-shallow)–CE (Prophet), across the 168 out-of-sample month–year pairs; green where MD-shallow is the better model.

Appendix A.3. Sensitivity and Uncertainty of the Estimated Lead

The lead reported in Section 4.1 is the argmax of a cross-correlation between filtered components. Because that statistic is discrete and non-linear, its sampling variability is assessed by a specification grid and by resampling rather than by an analytical band.

The grid varies the filter band, the filter family, the end-trimming and the sample period for a total of one hundred and twenty specifications, of which the full set is reported in the replication package. Specifications assuming stationary series are excluded from the summary below: with I(1) data the extracted component is dominated by the shared trend, yielding peak correlations above 0.97 that are invariant to the filter band. Table A2 reports the estimated lead at the baseline trimming of thirty-six months.

Table A2. Estimated lead in months by filter specification and sample. Trimming of thirty-six months at each end.

Filter	Band (Months)	Full Sample	Pre-2020
Christiano–Fitzgerald	12–96	1	4
Christiano–Fitzgerald	18–96	2	5
Christiano–Fitzgerald	24–96	4	6
Christiano–Fitzgerald	18–120	2	4
Baxter–King	12–96	1	4
Baxter–King	18–96	2	4
Baxter–King	24–96	4	5
Baxter–King	18–120	2	4
Hodrick–Prescott	n.a.	0	5

Note: The filters compared are Christiano–Fitzgerald (Christiano & Fitzgerald, 2003), Baxter–King (Baxter & King, 1999) and Hodrick–Prescott (Hodrick & Prescott, 1997). The Hodrick–Prescott cycle is a high-pass rather than a band-pass filter and retains high-frequency variation, so the band column does not apply. The full grid of one hundred and twenty specifications, including the stationary variants and the trimming lengths of twenty-four, twelve and zero months, is reported in the replication package. n.a. = not applicable.

Two regularities emerge. The lead is systematically longer on the pre-COVID-19 sample than on the full one, by three to four months, across every band and both band-pass families; and within each sample the dispersion across specifications is smaller than the difference between samples. The Hodrick–Prescott cycle, which is a high-pass rather than a band-pass filter and retains high-frequency variation, gives a lead of zero on the full sample and five before 2020.

Uncertainty is quantified by a moving-block bootstrap applied to the bivariate seasonally adjusted series in log differences, accumulated to levels, with the band-pass filter, cross-correlation and argmax re-executed on each replicate. Resampling is bivariate, so that the cross-series structure is preserved, and is performed before filtering, so that the interval reflects the variability of the estimator rather than the dependence induced by the filter. Seasonal factors are held fixed across replicates, an approximation adopted because re-estimating them two thousand times is not feasible. On the pre-COVID-19 sample with two thousand replicates the modal lead is four months at every block length examined, with a ninety per cent interval of one to six months at the longest.

The block length governs precision rather than location. A calibration exercise on a synthetic pair with a known lead of four months recovers a modal and median lead of exactly four at block lengths of thirty-six, sixty and eighty-four months, while the ninety per cent interval narrows from three-to-nine to three-to-five and the mass at the correct lag rises from fifty-six to seventy-nine per cent (Figure A3). The block bootstrap is therefore not biased downwards by the discontinuities between blocks, as might be feared at cycle frequencies, but longer blocks are more precise. This calibration uses a synthetic pair with a peak correlation near unity and does not establish the absence of bias under weak signal.

Table A3 reports the individual turning-point displacements underlying the discussion in Section 4.1.

Table A3. Individual turning-point displacements, deviation cycle. A positive value indicates that registrations turn before GDP.

Turning Point	GDP	Registrations	Displacement (Months)
Trough	1999 Q1	1997 Q3	+18
Peak	2001 Q2	2001 Q3	−3
Trough	2005 Q2	2005 Q3	−3
Peak	2008 Q1	2007 Q4	+3
Trough	2009 Q2	2009 Q1	+3
Peak	2011 Q3	2011 Q3	0
Trough	2014 Q1	2012 Q4	+15
Peak	2016 Q1	2016 Q1	0
Trough	2016 Q4	2016 Q4	0
Peak	2017 Q4	2017 Q4	0
Trough	2018 Q2	2018 Q2	0
Peak	2019 Q2	2019 Q2	0

Note: Turning points are dated by the Bry–Boschan–Harding–Pagan procedure and matched to the nearest counterpart within six quarters. The displacement of eighteen months coincides exactly with the maximum admitted by the matching rule and should be read as a bound rather than as an estimate. Six of the twelve pairs are exactly contemporaneous; the median displacement is zero and the mean is 2.8 months.




Figure A3. Bootstrap distribution of the estimated lead, pre-COVID-19 sample. Two thousand replicates per block length. The shaded band marks the ninety per cent percentile interval at the eighty-four-month block length. Lags above ten months account for less than one per cent of replicates and are omitted.

Appendix A.4. Out-of-Sample Predictive Test

This appendix reports the out-of-sample predictive exercise of Section 4.4 for industrial production under the principal (year-on-year) specification. The restricted benchmark is an autoregression in the year-on-year log growth of the IPI, with lag order re-selected by BIC at every origin from information available at that origin alone (maximum twelve lags); the unrestricted model adds lags of passenger-car registrations (maximum six). Forecasts are direct, the origin advances one observation at a time over an expanding window, and equal predictive accuracy is tested by the Clark–West statistic with a heteroskedasticity- and autocorrelation-consistent long-run variance whose bandwidth accommodates the eleven-month overlap induced by the year-on-year transformation. Three evaluation windows are

shown: the full sample; a window excluding any forecast whose origin-to-target interval intersects March 2020 to December 2021; and a window restricted to origins before 2020. A ratio below one favors registrations.

Table A4. Out-of-sample predictive test, industrial production, year-on-year specification. One-sided p -values. The three-month, excluding-COVID-19 cell (in bold) is the pre-specified principal cell.

Horizon	Window	Origins	MSPE Ratio	Clark–West	p
1	Excluding COVID-19	241	0.977	1.90	0.029
1	Pre-2020	216	0.988	1.64	0.051
1	Full	264	0.909	1.93	0.027
3	Excluding COVID-19	235	0.981	2.02	0.022
3	Pre-2020	212	0.988	1.62	0.053
3	Full	260	1.036	0.08	0.467
6	Excluding COVID-19	226	0.992	2.37	0.009
6	Pre-2020	206	0.993	2.33	0.010
6	Full	254	1.166	0.03	0.487
12	Excluding COVID-19	208	0.977	2.22	0.013
12	Pre-2020	194	0.977	2.19	0.014
12	Full	242	1.211	−0.83	0.797

Note: The full grid of forty-two cells, including quarterly GDP, the monthly specification with calendar dummies, and the exploratory reverse test, is reported in the replication package together with the forecast paths of every cell.

References

- Antenucci, D., Cafarella, M., Levenstein, M., Ré, C., & Shapiro, D. (2014, March). *Using social media to measure labor market flows* (Working Paper No. 20010). National Bureau of Economic Research. Available online: <http://www.nber.org/papers/w20010.pdf> (accessed on 2 August 2026).
- Bandara, K., Hyndman, R. J., & Bergmeir, C. (2025). MSTL: A seasonal-trend decomposition algorithm for time series with multiple seasonal patterns. *International Journal of Operational Research*, 52(1), 79–98. [CrossRef]
- Baumohl, B. (2009). *The secrets of economic indicators. Hidden clues to future economic trends and investment opportunities*. Wharton School Publishing.
- Baxter, M., & King, R. G. (1999). Measuring business cycles: Approximate band-pass filters for economic time series. *The Review of Economics and Statistics*, 81(4), 575–593. [CrossRef]
- Boot, J. C. G., Feibes, W., & Lisman, J. H. C. (1967). Further methods of derivation of quarterly figures from annual data. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 16(1), 65–75. [CrossRef]
- Box, G. E. P., & Cox, D. R. (1964). An analysis of transformations. *Journal of the Royal Statistical Society: Series B (Methodological)*, 26(2), 211–243. [CrossRef]
- Box, G. E. P., & Jenkins, G. M. (1970). *Time series analysis: Forecasting and control*. HoldenDay.
- Bry, G., & Boschan, C. (1971). *Cyclical analysis of time series: Selected procedures and computer programs*. National Bureau of Economic Research.
- Cantos-Sánchez, P., Gutiérrez-i-Puigarnau, E., & Mulalic, I. (2018). The impact of scrappage programmes on the demand for new vehicles: Evidence from Spain. *Research in Transportation Economics*, 70, 83–96. [CrossRef]
- Choi, H., & Varian, H. R. (2011, December 18). *Predicting the present with Google trends* [Technical Report]. Google. Available online: <http://people.ischool.berkeley.edu/~hal/Papers/2011/ptp.pdf> (accessed on 2 August 2026).
- Chow, G. C., & Lin, A. (1971). Best linear unbiased interpolation, distribution, and extrapolation of time series by related series. *The Review of Economics and Statistics*, 53(4), 372–375. [CrossRef]
- Christiano, L. J., & Fitzgerald, T. J. (2003). The band pass filter. *International Economic Review*, 44(2), 435–465. [CrossRef]
- Clark, T. E., & West, K. D. (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics*, 138(1), 291–311. [CrossRef]
- Cleveland, R. B., Cleveland, W. S., McRae, J. E., & Terpenning, I. (1990). STL: A seasonal-trend decomposition procedure based on Loess. *Journal of Official Statistics*, 6(1), 3–73.
- Cuevas, Á., Ledo, R., & Quilis, E. M. (2019). *Incorporando información fiscal de frecuencia diaria en la previsión macroeconómica a corto plazo* (Papeles de Trabajo del IEF, n° 11/2019). Instituto de Estudios Fiscales. Available online: https://www.ief.es/docs/destacados/publicaciones/papeles_trabajo/2019_11.pdf (accessed on 2 August 2026).

- Cuevas, Á., Ledo, R., & Quilis, E. M. (2021). Seasonal adjustment of the Spanish sales daily data. *Journal of the Spanish Economic Association*, 12, 687–708. [CrossRef] [PubMed]
- Dagum, E. B., & Cholette, P. A. (2006). Benchmarking, temporal distribution, and reconciliation methods for time series. In *Lecture notes in statistics* (Vol. 186). Springer.
- D’Amuri, F., & Marcucci, J. (2017). The predictive power of google searches in forecasting US unemployment. *Journal of Forecasting*, 33, 801–816. [CrossRef]
- Dargay, J., & Gately, D. (1999). Income’s effect on car and vehicle ownership, worldwide: 1960–2015. *Transportation Research Part A: Policy and Practice*, 33(2), 101–138. [CrossRef]
- De Livera, A. M., Hyndman, R. J., & Snyder, R. D. (2011). Forecasting time series with complex seasonal patterns using exponential smoothing. *Journal of the American Statistical Association*, 106(496), 1513–1527. [CrossRef]
- Denton, F. T. (1971). Adjustment of monthly or quarterly series to annual totals: An approach based on quadratic minimization. *Journal of the American Statistical Association*, 66(333), 99–102. Available online: <https://www.jstor.org/stable/2284856> (accessed on 2 August 2026). [CrossRef]
- Diebold, F. X. (2000). *Big data dynamic factor models for macroeconomic measurement and forecasting*. University of Pennsylvania and NBER.
- Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263. [CrossRef]
- Fernández, R. B. (1981). A methodological note on the estimation of time series. *The Review of Economics and Statistics*, 63(3), 471–478. [CrossRef]
- Findley, D. F., Monsell, B. C., Bell, W. R., Otto, M. C., & Chen, B.-C. (1998). New capabilities and methods of the X-12-ARIMA seasonal-adjustment program. *Journal of Business & Economic Statistics*, 16(2), 127–152. [CrossRef]
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press. Available online: https://www.deeplearningbook.org/front_matter.pdf (accessed on 2 August 2026).
- Harding, D., & Pagan, A. (2002). Dissecting the cycle: A methodological investigation. *Journal of Monetary Economics*, 49(2), 365–381. [CrossRef]
- Harvey, D., Leybourne, S., & Newbold, P. (1997). Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, 13(2), 281–291. [CrossRef]
- Heinisch, K., & Scheufele, R. (2019). Should forecasters use real-time data to evaluate leading indicator models for GDP prediction? German evidence. *German Economic Review*, 20, e170–e200. [CrossRef]
- Hodrick, R. J., & Prescott, E. C. (1997). Postwar U.S. business cycles: An empirical investigation. *Journal of Money, Credit and Banking*, 29(1), 1–16. [CrossRef]
- Huang, M. Y., Rojas, R. R., & Convery, P. D. (2018). News sentiment as leading indicators for recessions. *arXiv*, arXiv:1805.04160.
- Hyndman, R. J., & Athanasopoulos, G. (2021). *Forecasting: Principles and practice* (3rd ed.). OTexts. Available online: <https://otexts.com/fpp3/> (accessed on 2 August 2026).
- Kim, S. (2023). Innovating knowledge and information for a firm-level automobile demand forecast system: A Machine Learning perspective. *Journal of Innovation & Knowledge*, 8(2), 100355. [CrossRef]
- Kullback, S., & Leibler, R. A. (1951). On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1), 79–86. Available online: <https://www.jstor.org/stable/2236703> (accessed on 2 August 2026). [CrossRef]
- Ledo-Lobato, J., Navarro-Torres, A. M., & Mena-Nieto, Á. I. (2024). Forecasting the evolution of the vehicle fleet in Spain. *DYNA*, 99(5), 455–460. [CrossRef]
- Litterman, R. B. (1983). A random walk, Markov model for the distribution of time series. *Journal of Business & Economic Statistics*, 1(2), 169–173. [CrossRef]
- Ljung, G. M., & Box, G. E. P. (1978). On a measure of lack of fit in time series models. *Biometrika*, 65(2), 297–303. [CrossRef]
- Mosley, L., Eckley, I. A., & Gibberd, A. (2022). Sparse temporal disaggregation. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 185(4), 2203–2233. [CrossRef]
- Muth, M., Parsegyan, A. T., Litzinger, J., & Lingenfelder, M. (2025). Marketing predictions under macroeconomic volatility: Empirical evidence for automotive SMEs from a Machine Learning perspective. *Journal of Marketing Analytics*, 1–23. [CrossRef]
- Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708. [CrossRef]
- OECD. (2020). *Interpreting OECD composite leading indicators (CLIs)*. Available online: https://www.oecd.org/content/dam/oecd/en/data/methods/Interpreting_OECD_Composite_Leading_Indicators.pdf (accessed on 2 August 2026).
- Poza, C., & Monge, M. (2020). A real time leading economic indicator based on text mining for the Spanish economy. Fractional cointegration VAR and continuous wavelet transform Analysis. *International Economics*, 163, 163–175. [CrossRef]
- Ranabhatt, N., Barreto, S., Pimpão, M., & Prates, P. (2025). Demand forecasting in the automotive industry: A systematic literature review. *Forecasting*, 7(4), 73. [CrossRef]
- Sax, C., & Steiner, P. (2013). Temporal disaggregation of time series. *The R Journal*, 5(2), 80–87. [CrossRef]

- Scott, S. L., & Varian, H. R. (2012). *Bayesian variable selection for nowcasting economic time series* [Technical Report]. Google. Available online: <http://people.ischool.berkeley.edu/~hal/Papers/2012/fat.pdf> (accessed on 2 August 2026).
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379–423. [CrossRef]
- Szármes, P. (2014). Creating New types of business and economic indicators using big data technologies. *Science Journal of Business and Management*, 3(1-1), 18–24. [CrossRef]
- Tashman, L. J. (2000). Out-of-sample tests of forecasting accuracy: An analysis and review. *International Journal of Forecasting*, 16(4), 437–450. [CrossRef]
- Taylor, S. J., & Letham, B. (2018). Forecasting at scale. *The American Statistician*, 72(1), 37–45. [CrossRef]
- The Conference Board. (2026). *LEI for Spain*. Available online: <https://www.conference-board.org/topics/business-cycle-indicators/spain/> (accessed on 2 August 2026).
- Varian, H. R. (2014). Big data: New tricks for econometrics. *Journal of Economic Perspectives*, 28(2), 3–27. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.