

December 2025



Working Paper

34.2025

A Real-Time Framework for Forecasting Metal Prices

Andrea Bastianin, Luca Rossini, Lorenzo Tonni

A Real-Time Framework for Forecasting Metal Prices

Andrea Bastianin (Department of Economics, Management, and Quantitative Methods, University of Milan and Fondazione Eni Enrico Mattei), **Luca Rossini** (Department of Economics, Management, and Quantitative Methods, University of Milan and Fondazione Eni Enrico Mattei), **Lorenzo Tonni** (Department of Economics, Management, and Quantitative Methods, University of Milan)

Summary

This paper develops a real-time forecasting framework for monthly real prices of four key industrial metals – aluminum, copper, nickel, and zinc – whose demand is rising due to their widespread use in manufacturing and low-carbon technologies. To replicate the information set available to forecasters in real time, we construct a new dataset combining daily financial variables with first-release macroeconomic indicators and use nowcasting techniques to address publication lags. Within this real-time environment, we evaluate the predictive accuracy of a broad set of univariate, multivariate, and factor-augmented models, comparing their performance with two industry benchmarks: survey expectations and futures-spot spread models. Results show that although short-run metal price movements remain difficult to predict, medium-term horizons display substantial forecastability. Indicators of manufacturing activity tied to primary metals—such as new orders and capacity utilization—significantly improve forecasting accuracy for aluminum and copper, with more moderate gains for zinc and limited improvements for nickel. Futures and survey forecasts generally underperform the real-time econometric models. These findings highlight the value of incorporating timely macroeconomic information into forecasting frameworks for industrial metal markets.

Keywords: First-Release Data; Energy Transition; Forecasting; Metals; Critical Raw materials

JEL Classification: C32; Q02; Q41; Q43; Q48

Corresponding Author:

Lorenzo Tonni
Department of Economics, Management, and Quantitative Methods
University of Milan
Via Conservatorio, 7, 20122, Milano
lorenzo.tonni@unimi.it

A Real-Time Framework for Forecasting Metal Prices

Andrea Bastianin^{a,b} Luca Rossini^{a,b} Lorenzo Tonni^{a,*}

December 17, 2025

Abstract. This paper develops a real-time forecasting framework for monthly real prices of four key industrial metals – aluminum, copper, nickel, and zinc – whose demand is rising due to their widespread use in manufacturing and low-carbon technologies. To replicate the information set available to forecasters in real time, we construct a new dataset combining daily financial variables with first-release macroeconomic indicators and use nowcasting techniques to address publication lags. Within this real-time environment, we evaluate the predictive accuracy of a broad set of univariate, multivariate, and factor-augmented models, comparing their performance with two industry benchmarks: survey expectations and futures-spot spread models. Results show that although short-run metal price movements remain difficult to predict, medium-term horizons display substantial forecastability. Indicators of manufacturing activity tied to primary metals—such as new orders and capacity utilization—significantly improve forecasting accuracy for aluminum and copper, with more moderate gains for zinc and limited improvements for nickel. Futures and survey forecasts generally underperform the real-time econometric models. These findings highlight the value of incorporating timely macroeconomic information into forecasting frameworks for industrial metal markets.

Key Words: First-Release Data; Energy Transition; Forecasting; Metals; Critical Raw materials.

JEL Codes: C32; Q02; Q41; Q43; Q48.

^(a) Department of Economics, Management, and Quantitative Methods, University of Milan, Milan, Italy.

^(b) Fondazione Eni Enrico Mattei, Milan, Italy.

^(*) *Corresponding author:* Lorenzo Tonni, Department of Economics, Management, and Quantitative Methods, University of Milan, Via Conservatorio, 7, 20122, Milan, IT. Email: lorenzo.tonni@unimi.it.

Acknowledgments: The authors gratefully acknowledge H. Bjornland, F. Ravazzolo, A. Viselli and the participants at 10th Annual Conference of the Society for Economic Measurement (SEM), held at the Athens University of Economics and Business, and the 2025 International Association for Applied Econometrics (IAAE) Conference, held in Turin for their useful feedback. The authors acknowledge financial support from the “Fund for Departments of Excellence” provided by the Ministero dell’Università e della Ricerca (MUR), established by the Stability Law (“Legge di Stabilità n. 232/2016, 2017”), within the projects of the Department of Economics, Management, and Quantitative Methods (University of Milan).


1 Introduction

Prices of several metals are under strong demand pressures due to their expanding use in electric mobility and renewable energy technologies (IEA, 2025a). As shown in Figure 1, global consumption of aluminum, copper, nickel, and zinc has doubled since 1995, while prices have remained highly volatile. This trend is expected to continue. Under the International Energy Agency’s (IEA) “Net Zero Emissions” Scenario, global aluminum production is projected to grow at a Compound Annual Growth Rate (CAGR) of 1.39% from 2022 to 2030.¹ Demand for other metals is even more dynamic in clean energy applications: zinc demand for clean technologies is projected to grow at 17.2% per year (2024–2030), copper at 17.4% (vs. 4.3% for total demand), and nickel at 22.5% (vs. 8.0% for total demand). These discrepancies between total and clean-energy demand highlight the growing influence of low-emission technologies on metal markets and the rising policy concerns around ensuring sustainable and resilient supply chains. They also underscore the urgent need for reliable short- and medium-term real-time metal price forecasts to support effective market monitoring and scenario analysis.

Timely and reliable data collection is therefore a prerequisite for both credible price forecasting and sound policy design in critical mineral markets. As emphasized by IEA (2025b), traceability systems and standardized data protocols are essential for integrating real-time information on production, trade flows, and sectoral demand, which are key inputs for anticipating market developments and identifying potential bottlenecks. Incomplete or outdated information undermines the accuracy of short- and medium-term price forecasts and limits the effectiveness of policy responses aimed at managing supply-chain risks, stabilizing markets, and guiding investment decisions (Wittenstein and Stuermer, 2023).

¹The compound annual growth rate is defined as $(V_t/V_0)^{1/t} - 1$, where V_0 and V_t denote the initial and terminal values, respectively, and t is the number of years. The IEA’s Net Zero Emissions by 2050 Scenario outlines a pathway for the global energy sector to achieve net-zero CO₂ emissions by mid-century, implying rapid clean energy deployment, no new fossil fuel development, and systemic changes across all sectors. See: <https://www.iea.org/reports/global-energy-and-climate-model/net-zero-emissions-by-2050-scenario-nze>. For aluminum: IEA (2023), *Global aluminum production in the Net Zero Scenario, 2010–2030*, IEA, Paris. <https://www.iea.org/data-and-statistics/charts/global-aluminium-production-in-the-net-zero-scenario-2010-2030-3>. For copper, nickel, and zinc: IEA, Critical Minerals Dataset, IEA, Paris, <https://www.iea.org/data-and-statistics/data-product/critical-minerals-dataset>. Aluminum projections refer to production; for copper, nickel, and zinc we use demand.

Figure 1: Monthly Price and yearly world consumption of copper, aluminum, nickel, and zinc from January 1995 to December 2024.



Notes: Prices are described in Section 2.1. World consumption data are sourced from the World Bureau of Metal Statistics.

This paper addresses this policy issue by developing a comprehensive database that brings together a range of predictors for forecasting the real price of key metals: aluminum, copper, nickel, and zinc. The database is designed to replicate the information environment faced by policymakers, investors, and practitioners when producing real-time forecasts at monthly sampling frequency. To this end, we rely on first-release values of macroeconomic indicators and employ nowcasting techniques to bridge data gaps due to publication lags. Within this framework, we evaluate the forecasting performance of several models relative to standard industry benchmarks, including survey-based forecasts and futures–spot spread models.

Modelling and forecasting commodity prices has a long tradition in economics. Within the forecasting literature, a vast body of research has systematically examined crude oil prices (see, e.g., Alquist et al., 2013, for a survey) and, more broadly, energy prices (Ferrari et al., 2021), covering refined oil products (Bastianin et al., 2014; Baumeister et al., 2017), natural gas (Ruben et al., 2025; Baumeister et al., 2025), and electricity (Feroni et al., 2023; Gianfreda et al., 2020; Hauzenberger et al., 2025). In contrast, the literature on metal markets and critical raw materials remains notably fragmented – despite their growing strategic importance in the energy transition. An early survey on nonferrous metals is provided by Watkins and McAleer (2004).

Although this paper belongs to the broader literature on commodity price forecasting,

its contribution is particularly aligned with research that employs real-time data to forecast commodity prices and with studies modeling mineral prices and their role in the energy and digital transitions. Within the strand of research using real-time data for commodity price forecasting, crude oil continues to dominate the field (see Baumeister and Kilian, 2012, 2015; Ravazzolo and Rothman, 2013), with Baumeister et al. (2025) representing a notable exception for natural gas. In contrast, research on forecasting metal prices remains limited, and none of the existing studies, to our knowledge, relies on real-time data. Pierdzioch et al. (2013) analyze survey-based forecasts, while other works emphasize the role of macro-financial and structural drivers as predictors (Duarte et al., 2021; Gargano and Timmermann, 2014; Issler et al., 2014; Stuermer, 2017; West and Wong, 2014). Moreover, a growing number of recent studies have examined the role of metals in the ecological transition and the determinants of their prices using structural vector autoregressive (VAR) models (see, e.g., Bastianin et al., 2023; Baumeister et al., 2022; Boer et al., 2024; Considine et al., 2023).

Compared with the existing literature, this paper makes three main contributions. First, it develops a real-time forecasting framework for the monthly real prices of aluminum, copper, nickel, and zinc, supported by a newly constructed dataset that combines first-release macroeconomic indicators with nowcasting techniques. This approach reproduces the information actually available to forecasters at each point in time – a perspective absent from previous studies on metal price forecasting and crucial for policymakers and market participants operating under real-time uncertainty. Second, the paper provides a comprehensive evaluation of forecasting models, ranging from univariate benchmarks to multivariate specifications with macro-financial predictors, and compares their performance with industry standards such as survey expectations and futures–spot spread models. This comparison highlights the value of macroeconomic information relative to widely used market-based forecasts. Third, the analysis shows that the predictability of metal prices is highly horizon-dependent: short-term movements are largely erratic, while medium-term dynamics are systematically linked to manufacturing activity in sectors closely tied to primary metals. Taken together, these findings demonstrate that incorporating timely macroeconomic indicators substantially improves medium-term forecasts, while also clarifying the limits of predictability at shorter horizons. Our paper therefore extends real-time forecasting methods to an

important class of industrial commodities and provides new evidence on the macroeconomic forces shaping base metal prices.

The remainder of the paper is structured as follows. Section 2 describes the real-time database and compares nowcasting approaches. Section 3 presents the models and the forecasting experiment. Section 4 presents the main results, while Section 5 focuses on the time-varying performance of forecasts and assesses the benefits of forecast pooling. Section 6 concludes.

2 Real-Time Dataset Construction

This section describes the construction of the real-time dataset. We combine daily financial data with the first-release vintages of monthly macroeconomic predictors. These monthly series exhibit publication lags of varying length, so we employ nowcasting to make them usable in real-time as shown in Section 2.2.

2.1 Data

Daily data. We collect daily nominal spot prices for aluminum, copper, nickel, and zinc traded on the London Metal Exchange (LME) from the London Stock Exchange Group (LSEG). Since aluminum, copper, nickel, and zinc are storable commodities, we also source daily data on visible inventories from LSEG, representing the volume of stocks held in LME-approved warehouses.² Inventories proxy for market tightness and expected scarcity – the key state variable linking current market conditions to expectations about future supply and demand in the competitive storage model (Deaton and Laroque, 1992, 1996). Because they capture physical imbalances that drive short- and medium-term real price movements, inventories are routinely included in empirical models of commodity markets (Ruben et al., 2025; Alquist et al., 2020; Baumeister et al., 2025; Kilian and Murphy, 2014).

Both nominal price and inventory data are available in real-time and are not subject to revisions. In the forecasting exercise, daily observations are aggregated by taking calendar-

²Inventories are termed *visible* because they are publicly reported and correspond to the portion of global stocks available for delivery against LME contracts.

month averages.

Monthly data. Real-time data for selected monthly macroeconomic variables were retrieved from ALFRED (Archival Federal Reserve Economic Data), maintained by the Federal Reserve Bank of St. Louis. See Table 1 for the full list of variables. As shown in Table 1, we include five broad groups of predictors. Producer price indices (PPI) capture inflationary trends and cost pressures in metal-intensive industries: PPI series for primary and battery manufacturing reflect upstream price dynamics.

Since our focus is on forecasting real metal prices, nominal spot prices are deflated using the CPI vintage series of the corresponding period, with the two most recent missing observations nowcasted as described in the Section 2.2.

Indicators of industrial production, new orders, shipments, and manufacturing hours track cyclical demand for metals, while capacity utilization in the primary and semifinished metal sectors quantifies supply tightness. To account for structural shifts associated with the energy and digital transitions, we use shipments and producer prices in the battery-manufacturing sector as proxies for rising demand for copper, nickel, and aluminum in electrification and storage technologies.

Finally, real broad effective exchange-rate indices are included for Australia, Chile, China, Indonesia, Peru, the Philippines, and Russia – countries accounting for a large share of global production of aluminum, copper, nickel, and zinc. As shown by Chen et al. (2010), exchange rates of “commodity-currency” countries have robust predictive power for global commodity prices, since they anticipate future export earnings and terms-of-trade shifts.

With the exception of exchange rates, the remaining macroeconomic variables refer to the United States. This choice is motivated not only by the availability and ease of access to granular real-time data through ALFRED, but also by the prominent role of the U.S. economy in global mineral markets through its positions as a major consumer, refiner, and financial centre. Moreover, metal prices are quoted in U.S. dollars and respond strongly to U.S. macroeconomic and financial conditions.

Real-time data issues. As shown in Table 1, all macroeconomic series are released with publication lags of two up to three months. Consequently, the dataset exhibits a ragged-

Table 1: Real-time database structure: macroeconomic and financial variables.

Group	ID	Description	Transformation	Missing Obs
Prices	CPI	Consumer Price Index for All Urban Consumers: All Items	$\Delta^2 \log(x_t)$	2
	PPI-M	Producer Price Index: Primary Metal Manufacturing	$\Delta \log(x_t)$	2
Ec. Act.	IP	US Industrial Production	$\Delta \log(x_t)$	2
	HEM	Hours of Employees, Manufacturing	$\log(x_t)$	2
	NO-M	New Orders: Primary Metals	$\log(x_t)$	2
	NO-A	New Orders: Aluminum and Nonferrous Metal Products	$\log(x_t)$	3
CU	CUP	Capacity Utilization: Primary & Semifinished Processing	$\log(x_t)$	2
	CUM	Capacity Utilization: Primary Metal	$\log(x_t)$	2
ET	PPI-B	Producer Price Index: Battery Manufacturing	$\Delta \log(x_t)$	2
	MVS	Manufacturers' Value of Shipments: Battery Manufacturing	$\log(x_t)$	3
Exchange Rates	AUS	Real Broad Effective Exchange Rate for Australia	$\Delta \log(x_t)$	2
	CHL	Real Broad Effective Exchange Rate for Chile	$\Delta \log(x_t)$	2
	CHN	Real Broad Effective Exchange Rate for China	$\Delta \log(x_t)$	2
	IDN	Real Broad Effective Exchange Rate for Indonesia	$\Delta \log(x_t)$	2
	PER	Real Broad Effective Exchange Rate for Peru	$\Delta \log(x_t)$	2
	PHL	Real Broad Effective Exchange Rate for Philippines	$\Delta \log(x_t)$	2
	RUS	Real Broad Effective Exchange Rate for Russia	$\Delta \log(x_t)$	2
Inventories	ALU-V	Aluminum Inventory Volume	$\Delta \log(x_t)$	–
	COP-V	Copper Inventory Volume	$\Delta \log(x_t)$	–
	NIC-V	Nickel Inventory Volume	$\Delta \log(x_t)$	–
	ZNC-V	Zinc Inventory Volume	$\Delta \log(x_t)$	–


Notes: In column 1, “EC. Act.” refers to Economic Activity; “CU” stands for Capacity Utilization, while “ET” refers to predictors related to the increasing use of metals in technologies involved in the energy and digital transitions.

edge structure (Wallis, 1986), as observations for different variables are available at different times. We address this issue by nowcasting the missing observations each time a real-time forecast is produced. Conducting a real-time forecasting exercise requires using data vintages exactly as they were available at the time each forecast was made.

As an illustration of the magnitude of data revisions, Figure 2 shows the history of the Manufacturers’ New Orders for Primary Metals series between February 2011 and January 2025. The initial release in 2011 reported a value of 20,800 million dollars (red dot). Over subsequent years, revisions alternated between downward and upward adjustments—rising to 20,400 million in 2013 and peaking at 21,800 million in 2020. In recent years, revisions have become smaller, stabilizing around 21,400 million dollars. This pattern highlights how successive data updates can alter both the level and trajectory of a series, reinforcing the importance of using first-release information in real-time forecasting.

To replicate the informational constraints faced by real-world forecasters, we collect the

Figure 2: Effects of revisions of Manufacturers’ New Orders for Primary Metals (black line) from February 2011 (red dot) to January 2025 (blue dot) in millions of dollars.



first-release value of each observation from ALFRED, constructing a dataset in which each vintage expands only through the addition of newly available observations (see Supplement for details on dataset construction). Table 1 also reports the transformations applied to each predictor as it enters forecasting models, as well as the number of missing observations requiring nowcasting.³

2.2 Nowcasting and backcasting

As shown in Table 1, most predictors are observed with a two-month lag relative to the forecast period, except for the New Orders for Aluminum products (NO-A) and Manufacturers’ Shipments (MVS) series, which exhibits three missing observations. To forecast the prices of aluminum, copper, nickel, and zinc, we first address missing values in the series by nowcasting and backcasting them using alternative time series models. To bridge these data gaps and replace missing observations with model-based estimates, we conduct a “horse race” among univariate time series models, estimated under both frequentist and Bayesian frameworks. This exercise is essential because, unlike for widely studied macroeconomic indicators such as CPI or industrial production, there is limited empirical evidence to guide the selection of

³Transformations follow, as closely as possible, those adopted by McCracken and Ng (2016), with deviations made when comparable series were unavailable. Additional information on data sources and series identifiers is provided in the Supplement.

appropriate nowcasting models for many of the commodity price series considered here.

Our benchmark model is the Random Walk with drift (RW-D), a common reference in forecasting due to its simplicity and robust empirical performance. As frequentist alternatives, we consider a simple autoregressive model with one lag (AR(1)) and an autoregressive model with a lag length selected each period by the Akaike Information Criterion (AR(AIC)), allowing for time-varying dynamics.

On the Bayesian side, we start with a Bayesian AR(1) with constant volatility (BAR(1)), and then extend it to account for stochastic volatility (BAR(1)-SV), following evidence that time-varying uncertainty improves macroeconomic forecasts (Carriero et al., 2019; Chan, 2020). Finally, since our sample covers the COVID-19 period, we also consider a Bayesian AR(1) model with stochastic volatility and outliers (BAR(1)-SVo), as proposed by Chan (2023) and Carriero et al. (2024), to better capture the presence of extreme shocks.

For each data vintage and nowcast horizon, we estimate the model parameters and project the process h -steps ahead, where h ranges from 1 to 3 depending on the publication lag of the series. The Bayesian and frequentist estimation follows a rolling window approach: as we move to the next data vintage, we shift the in-sample window forward by one period, incorporating a newly available first-release observation and repeating the process until the latest vintage, November 2023.

To assess the goodness of the forecasts and compare the model performance, we compute the Root Mean Square Forecasting Errors (RMSFEs) using forecast errors from all six models, measured in the original units of each variable. We report the RMSFEs for the baseline RW-D model, whereas for the other autoregressive models, we report the ratios between the computed metric of the current model over the computed metric of the baseline model. Then, entries of less than 1 indicate that the given current model yields nowcasts more accurate than those provided by the baseline random walk model. In line with Carriero et al. (2015) and Clark and McCracken (2015), rely the Diebold-Mariano test (Diebold and Mariano, 1995) to evaluate the forecast accuracy against the RWD model. The differences in accuracy that are statistically different from zero are denoted with one, two, or three asterisks, corresponding to significance level of 10%, 5%, and 1%, respectively.

Table 2 presents the RMSFE along with the ratios and the associated Diebold-Mariano

test for each variable and nowcasting horizon, with the lowest ratios highlighted in bold. All autoregressive models show similar nowcasting performances across variables and nowcast horizon. As the gains from more sophisticated models appear marginal and considering that a real-time forecaster cannot know in advance which models will perform best over the entire sample, we adopt the simple RW-D model for subsequent analysis. We notice that when the advantages of using sophisticated models is higher in terms of the ratio, the Diebold-Mariano test does not provide evidence of their significance with respect to the benchmark (e.g. PPI-M rows).

Table 2: Nowcasting of different predictors for each horizon and models.

ID	Horizon	RW-D	AR(1)	AR(AIC)	BAR(1)	BAR(1) - SV	BAR(1) - SVo	ID	Horizon	RW-D	AR(1)	AR(AIC)	BAR(1)	BAR(1) - SV	BAR(1) - SVo
IP	1	2.50	1.17***	1.23***	1.14***	1.07***	1.06***	PPI-B	1	1.80	1.02***	0.98	1.02***	0.99**	1**
	2	4.18	1.05***	1.21***	1.04***	1.04***	1.05***		2	2.69	1.03***	0.94	1***	1**	1***
	3	-	-	-	-	-	-		3	-	-	-	-	-	-
NO-M	1	1126.48	0.92	0.92	0.92	0.92	0.92	CHL	1	1.98	1.04***	1.04***	1.04***	1.03***	1.03***
	2	1529.06	0.94	0.92	0.94	0.95	0.95		2	2.94	1**	1***	1**	0.99	0.99*
	3	-	-	-	-	-	-		3	-	-	-	-	-	-
NO-A	1	739.94	1.06***	1.06***	1.06***	1.07***	1.07***	AUS	1	1.49	1.02***	1.01***	1.02***	1.01***	1***
	2	1067.93	1.01***	1***	1.01***	1.03***	1.03***		2	2.15	1.04***	1.04***	1.03***	1***	1***
	3	1336.36	0.98**	0.97	0.98*	1***	1***		3	-	-	-	-	-	-
CPI	1	0.00	0.97	0.87	0.97	0.97	0.98	CHN	1	1.52	0.98	0.98	0.98	0.98	0.98
	2	0.01	1.07***	0.91	1.08***	1.11***	1.11***		2	2.41	1**	1***	0.99*	1***	1**
	3	-	-	-	-	-	-		3	-	-	-	-	-	-
HEM	1	0.32	1.01***	1.03***	1.02***	1***	1.01***	IDN	1	1.51	1.02***	0.96	1.03***	1.02***	1.02***
	2	0.44	0.97	0.97	1**	0.98*	0.98		2	2.31	1.01***	0.97	1.01***	1.01***	1.01***
	3	-	-	-	-	-	-		3	-	-	-	-	-	-
MVS	1	50.74	1.05***	1.06***	1.05***	1.04***	1.04***	PER	1	2.20	0.97*	0.97*	0.97*	0.97*	0.97*
	2	68.92	1.01***	1.02***	1.01***	1***	1**		2	3.31	0.99	0.99	0.98	0.98	0.98
	3	76.58	1.02***	1.06***	1.02***	1.01***	1***		3	-	-	-	-	-	-
CU-M	1	1.84	1.09***	1.06***	1.09***	1.09***	1.09***	PHL	1	1.36	1.01***	1**	1.01***	1.02***	1.01***
	2	3.17	1.01***	1***	1.01***	1.01***	1.01***		2	2.07	1.04***	1.04***	1.03***	1.04***	1.04***
	3	-	-	-	-	-	-		3	-	-	-	-	-	-
CU-P	1	1.17	1.04***	1.15***	1.04***	1.04***	1.04***	RUS	1	6.05	0.78	0.8	0.78	0.75	0.75
	2	1.84	1.02***	1.24***	1.02***	1***	1.01***		2	9.18	0.85	0.86	0.85	0.86	0.86
	3	-	-	-	-	-	-		3	-	-	-	-	-	-
PPI-M	1	6.61	0.76	0.76	0.76	0.75	0.75	-							
	2	11.89	0.8	0.8	0.79	0.78	0.78								
	3	-	-	-	-	-	-								

Notes: For the baseline (RW-D) model, we report the RMSFE, whereas for the other models, the ratio relative to the baseline model. A ratio below one indicates that the model outperforms the RW-D model. ***, **, and * indicate that ratios are significantly different from 1 at the 1%, 5%, and 10% based on the Diebold and Mariano (1995) test. The best-performing model at each horizon is highlighted in bold.

3 Forecasting Models and evaluations

3.1 Model Description

We consider both univariate and multivariate models for the real price of aluminum, copper, nickel and zinc. Moreover, we include various predictors as outlined in Section 2.1 and we consider forecast at different horizons. As a natural comparison for the univariate and multivariate models, we also rely on forecasts already available in the market, such as those

implied by future contracts (Alquist and Kilian, 2010) or produced by professional forecasters (Ang et al., 2007; Croushore, 2010).

3.1.1 Univariate models

As a baseline specification, we consider the RW-D model. In addition, we include two alternative univariate time series models: an AR(1) model and an AR model in which the number of lags is selected according to the AIC (AR(AIC)).

Moreover, to leverage information from the set of real-data predictors discussed in Section 2.1, we augment the AR model with exogenous variables. Thus, we consider Autoregressive Distributed Lags (ARDL(p,s)) models, where we include the p lags of the dependent variable and s lags of the predictor variables:

$$\hat{y}_{t+h} = \hat{c}_h + \sum_{i=1}^p \hat{a}_{i,h} y_{t-i+1} + \sum_{j=1}^s \hat{\Gamma}_{j,h} \mathbf{Z}_{t-j+1}, \quad (1)$$

where $\hat{y}_{t+h}$ is the real metal prices forecasted h steps ahead, $\hat{c}_h$ is the estimated intercept, $\hat{a}_{i,h}$ the estimated autoregressive coefficient for each lag $i = 1, \dots, p$. The term $\mathbf{Z}_{t-j+1}$ is a $K \times 1$ vector of exogenous predictors at lag $t - j + 1$ and $\hat{\Gamma}_j$ is the corresponding $1 \times K$ vector of coefficients for lag $j = 1, \dots, s$.

As for the AR model, the ARDL model is estimated either with 1 lags for both autoregressive and predictor variables or with the number of lags selected using the AIC. Regarding the choice of the K real-time predictors, we include a single predictor at time or we use all predictors within the same categories as outlined in Section 2.1.

Another widely used approach in the forecasting literature to exploit information from a dataset with a large cross-sectional dimension is to estimate factors that summarize the information contained in the real-time dataset. Following Stock and Watson (2002), we augment the AR model with one or two factors extracted from the set of real-time predictors and we define the Autoregressive Diffusion Index (ARDI) model.

The specification is analogous to Equation 1, except that the vector of exogenous variables $\mathbf{Z}_{t-j+1}$ is replaced by an $r \times 1$ vector of estimated factors $\hat{\mathbf{F}}_{t-j+1}$ at lag j and the corresponding coefficients are collected in a $1 \times r$ vector $\hat{\Gamma}_j$, for $j = 1, \dots, q$. These factors

are estimated with the PCA approach by Stock and Watson (2002), where r is the number of factors. The $r \times N$ matrix of factor loadings $\hat{\mathbf{\Lambda}}$ is estimated as the product between $\sqrt{N}$ and the r normalized eigenvectors corresponding to the r -largest eigenvalues of the sample covariance matrix of the standardized data $\tilde{\mathbf{x}}_t$.⁴ Then, the factors, $\hat{\mathbf{F}}_t$, are obtained by projecting the estimated loadings onto the data, i.e., $\hat{\mathbf{F}}_t = \hat{\mathbf{\Lambda}}\mathbf{x}_t$.

3.1.2 Multivariate models

As a further set of specification, we employ on a Vector Autoregressive (VAR) models of order p :

$$\hat{\mathbf{Y}}_{t+h} = \hat{\mathbf{C}}_h + \sum_{i=1}^p \hat{\mathbf{A}}_i \mathbf{Y}_{t-i+1}, \quad (2)$$

where $\hat{\mathbf{Y}}_{t+h}$ is the $n \times 1$ estimated vector of endogenous variables at horizon $t+h$. This vector includes the price and inventory volume for the mineral of interest, along with metals' new orders in the manufacturing sector (NO-M), for a total of $n = 3$ endogenous variables. The aim is to jointly model a measure of the metal price, a measure of the metal stock quantity, and a proxy for metal demand. $\hat{\mathbf{C}}_h$ is the $n \times 1$ estimated vector of constants, $\mathbf{Y}_{t-i+1}$ is the $n \times 1$ vector of endogenous variables at lag $t - i + 1$ and $\hat{\mathbf{A}}_i$ is the associated $n \times n$ matrix of autoregressive coefficients.

Additionally, we consider an alternative multivariate specification where we augment the VAR model in Eq. (2) with either one or two factors constructed as in the ARDI model, thereby estimating a Factor-Augmented VAR (FAVAR)⁵. As in the ARDI framework, this approach has the advantage of automatically selecting the most relevant information from a large set of predictors by weighting them differently in the factor estimation. Consequently, it removes the burden of manually identifying the most informative predictors from the list in Table 1. Furthermore, because the factors are re-estimated at each iteration, the associated factor loadings are allowed to evolve over time, thereby offering some flexibility to accommodate potential structural breaks in the sample.

⁴The target metal price is excluded from $\tilde{\mathbf{x}}_t$, as it already enters the model via the dependent variable y .

⁵In this case, we exclude from $\tilde{x}_t$ the metal price, metal inventory volume, and metals' new orders in the manufacturing sector (NO-M), since these are already included in the vector of endogenous variables $\mathbf{Y}$.

3.1.3 Model-free forecasts

The model previously described are compared against a baseline model, but they could also be compared with the forecasts available in the market, such as those implied by the future contracts or produced by professional forecasters.

Forecasts based on futures prices. Since future contracts specify a price agreed upon in the present for the exchange of a commodity at a future date, their prices inherently reflect the market's expectations regarding future spot prices (Alquist and Kilian, 2010). Leveraging this characteristics, we employ LSEG Data & Analytics on future contracts with 3- and 15-month maturity⁶. The market's expectation for the metal spot real price at time $T + h$ is computed as

$$\hat{P}_{T+h|T} = \frac{F_T^{T+h}}{\hat{p}_{T+h|T}}, \quad (3)$$

where F_T^{T+h} denotes the futures price at time T for a transaction scheduled at time $T + h$, and $\hat{p}_{T+h|T}$ is the forecast at time T for the CPI index at time $T + h$ ⁷.

Professional forecasters. Another model-free forecast is provided by projections produced by professional forecasters. We compute the average forecast across individual forecasters published in the Consensus Economics survey. In detail, each monthly survey contains forecasts for fixed calendar dates – March, June, September, and December – regardless of the survey month. As a result, the forecast horizons are not constant but vary depending on the survey date. Since these forecasts follow a *fixed-event* structure, $y_{T+h|T}^{FE}$, we convert them into *fixed-horizon* forecasts, $y_{T+h|T}^{FH}$, using an interpolation method as in Dovern et al. (2012):

$$y_{T+h|T}^{FH} = \frac{d - |h - h_1|}{d} y_{T+h_1|T}^{FE} + \frac{d - |h - h_2|}{d} y_{T+h_2|T}^{FE}, \quad (4)$$

where d is the distance between two consecutive forecasted horizons in the survey and in our case is equal to 3 since forecasts are issued for a single month every quarter. The variable h denotes the fixed forecast horizon of interest, while h_1 and h_2 are the survey horizons immediately preceding and following h , respectively.

⁶Other maturities lower than 24 months were not available.

⁷We estimate $\hat{p}_{T+h|T}$ as $p_T(1 + \hat{g}_p^h)$, where $\hat{g}_p^h$ is the CPI sample growth rate over h periods.

3.2 Forecast evaluation methods

Based on the model described above, the datasets cover the period from January 2000 to November 2023, where the macroeconomic real-time observations start in April 2015⁸. To conduct the forecasting exercise, we set April 2015 as the last in-sample date for the first data vintage. Then, we forecast the real metal prices for different horizons, ranging from 1 to 24 months to better describe the short and medium term.

The metal prices are forecasted in growth rates and then we reconstruct the forecasts in the original measurement unit before computing the relative forecasting errors. As we move to subsequent data vintages, we shift the in-sample window one period forward and we repeat the procedure until November 2023 by adopting a rolling window approach. Since we are interested in forecasting up to 24 horizons, we use the direct forecasting method for the exogenous variables (see, e.g. Marcellino et al., 2006).

As stated in Section 2.2, we compute the Root Mean Square Forecasting Error (RMSFE) to evaluate the model performance. We report the RMSFEs for the baseline RW-D model, while for the other models we report the ratios between the computed metric of the current model over the one of the baseline model. Moreover, we apply the Diebold-Mariano t test (Diebold and Mariano, 1995), where we compare each model with respect to the benchmark model and we denote the difference in accuracy that are statistically different from zero with one, two, or three stars corresponding to 10%, 5%, and 1%, respectively. Finally, to jointly compare different model specifications, we employ the Model Confidence Set procedure proposed by Hansen et al. (2011) for a fixed horizon. The differences have been tested separately for each horizon and the MCS has been used to compare the predictive power without disentangling between univariate, multivariate or factor models. Results are presented in cells highlighted in light blue with confidence level α fixed at 0.25⁹.

⁸As shown in the Supplement, the start of real-time recording differs across series. To construct a balanced panel, we set April 2015 as the initial vintage, corresponding to the beginning of real-time availability for the Producer Price Index variables, which are the latest to begin real-time recording.

⁹Results for different values of α and for all models are available in the Supplement.

4 Forecasting the real prices of raw materials

Our forecasting evaluation is based on horizons up to 24 months ahead and compares a wide range of univariate and multivariate models, from simple benchmarks to factor-augmented specifications. The main results for copper and aluminum are summarised in Table 3 for selected horizons, while the corresponding full set of results is reported in the Supplement. For zinc and nickel, the key models are shown in Table 4. In all cases, we assess whether including real-time macroeconomic predictors delivers statistically significant gains in forecast accuracy relative to a random walk with drift (RW-D) baseline. The tables report RMSPEs (for RW-D) and RMSPE ratios (for all other models), alongside Diebold–Mariano tests and the Set of Superior Models (SSM) obtained from the MCS procedure.

Two general features emerge. First, different models outperform the baseline across metals and horizons, and no single specification dominates. This is consistent with the model selection uncertainty highlighted by the MCS results. Second, forecastability is horizon- and metal-specific: short-horizon gains are modest, while medium-term horizons (6–18 months ahead) are where most of the economically relevant improvements arise. Below we discuss each metal in turn and then compare these model-based forecasts to model-free benchmarks.

Copper. For copper, forecasting gains are particularly pronounced at medium-term horizons (6 to 18 months ahead), in line with the metal’s pivotal role for global industrial activity. As shown in the top panel of Table 3 and in the Supplement, models incorporating information from the U.S. manufacturing sector (NO-M) deliver the strongest improvements. In particular, the ARDL model using manufacturers’ new orders for primary metals as an exogenous regressor reduces the RMSPE by roughly 20% at the 6-month horizon and by about 30% at the 12-month horizon relative to RW-D, with both gains statistically significant and the model frequently included in the SSM.

Moving from univariate to multivariate specifications further improves performance, although the gains are somewhat smaller than for the best ARDL(NO-M) model. When new orders are included in a VAR alongside inventories and copper prices, medium-term horizons still exhibit RMSPE reductions of around 20%. This pattern is consistent with the idea that inventories transmit information on current and expected market tightness, but that new

orders remain the dominant driver of medium-term predictability.

Factor-augmented models (ARDI and FAVAR) confirm and refine this picture. Both one- and two-factor versions perform strongly and enter the SSM at several horizons. The first factor places high loadings on metals-related manufacturing activity, capacity utilisation and producer prices, capturing a broad demand-driven macroeconomic component. The second factor loads more heavily on inventories and other metals' prices, reflecting a market-specific component related to co-movements across metals and short-run supply adjustments. The forecasting gains of these factor models therefore arise from their ability to separate copper price movements into a macro-demand factor and a metals-specific factor.

Over short (1-month) and very long (24-month) horizons, the improvements relative to RW-D are more modest—on the order of 6–10% for the best models. By contrast, at medium-term horizons, the best-performing ARDL and factor-augmented models achieve RMSPE reductions in the range of 17–31% (from horizon 9 to 18), indicating that copper prices are substantially more predictable over these horizons when macroeconomic information is exploited.

To illustrate the underlying mechanisms, Figure 3 plots the first two FAVAR factors for copper at the 12-month horizon.¹⁰ The left panel shows a broad macroeconomic factor largely driven by real economic activity, capacity utilisation and producer prices, while the right panel highlights a metals-market factor dominated by inventories and other metals' prices.

Aluminum. Turning to aluminum (bottom panel of Table 3), forecasting gains are again concentrated at medium-term horizons (roughly 6 to 15 months ahead). The largest improvements, of about 19–22%, are obtained from univariate ARDL models that incorporate manufacturing-sector indicators, especially capacity utilisation and new orders for primary metals. These results reflect the tight link between aluminum demand and industrial production, with capacity utilisation in metals-related manufacturing acting as a particularly informative proxy.

As with copper, factor-augmented models improve on the benchmark and enter the SSM


¹⁰The corresponding factors for the other commodities and for the ARDI models are reported in the Supplement.

Table 3: Copper and aluminum real prices point forecasts based on different models and horizons.

Metal	Model	Horizon 1	Horizon 3	Horizon 6	Horizon 9	Horizon 12	Horizon 15	Horizon 18	Horizon 21	Horizon 24
Copper	RW-D	292.26	610.34	918.10	1139.00	1363.14	1527.96	1628.31	1634.84	1625.21
	Futures	-	1.145***	-	-	-	1.004***	-	-	-
	Professional Forecasters	-	-	1.069***	1.000***	0.942	0.898	0.857**	-	-
	ARDL(1) - Ec. Activity	1.018***	1.029***	1.006***	0.941	0.860*	0.883	0.926	0.991*	1.058***
	ARDL(1) - Ex. Rates	0.975	0.980*	1.011***	1.045***	1.039***	1.052***	1.062***	1.069***	1.063***
	ARDL(1) - ETM Prices	0.984	1.011***	1.171***	1.197***	1.190***	1.141***	1.105***	1.077***	1.099***
	ARDL(1) - ET	0.973	0.926	0.931	0.933	0.929	0.950	1.000***	1.115***	1.199***
	ARDL(1) - CU	1.007***	1.068***	1.167***	1.199***	1.272***	1.280***	1.308***	1.445***	1.530***
	ARDL(1) - Inventories	0.976	1.007***	1.079***	1.106***	1.085***	1.149***	1.136***	1.167***	1.271***
	VAR(1)	0.985	0.893*	0.796**	0.772**	0.750***	0.797***	0.855***	0.939*	0.997*
	ARDI(1) - 1 Factors	0.942	0.856**	0.779**	0.717**	0.765**	0.832**	0.894**	1.004***	1.122***
	ARDI(1) - 2 Factors	0.943	0.861**	0.786**	0.691**	0.733**	0.815**	0.869**	1.002***	1.152***
	FAVAR(1) - 1 Factors	0.956	0.859**	0.784**	0.760**	0.743***	0.780***	0.841***	0.946*	1.035***
	FAVAR(1) - 2 Factors	0.957	0.868**	0.789**	0.734**	0.709***	0.766***	0.829***	0.965	1.097***
Aluminum	RW-D	89.91	192.38	276.64	350.61	424.51	483.25	525.46	529.15	523.40
	Futures	-	1.803***	-	-	-	1.093***	-	-	-
	Professional Forecasters	-	-	1.151***	1.054***	0.978	0.918	0.848**	-	-
	ARDL(1) - Ec. Activity	0.936	1.020***	0.989*	0.996**	0.970	0.970	0.998**	1.035***	1.090***
	ARDL(1) - Ex. Rates	0.961*	1.003***	1.016***	1.032***	1.052***	1.043***	1.036***	1.022***	1.013***
	ARDL(1) - ETM Prices	0.952	1.051***	1.143***	1.131***	1.130***	1.084***	1.053***	1.029***	1.040***
	ARDL(1) - ET	0.919	0.992	0.955	0.935	0.924	0.918	0.955	1.013***	1.042***
	ARDL(1) - CU	0.926	1.015***	1.046***	0.995**	0.954	0.938	0.955	1.012***	1.085***
	ARDL(1) - Inventories	0.951	1.030***	1.017***	1.024***	1.022***	1.010***	0.996*	1.001***	0.990
	VAR(1)	0.927	1.012***	0.939*	0.882**	0.863**	0.858***	0.884**	0.942	0.971
	ARDI(1) - 1 Factors	0.915*	0.948	0.906*	0.860*	0.843*	0.852	0.884	0.926	0.999***
	ARDI(1) - 2 Factors	0.982	0.993	0.942	0.871*	0.847*	0.830*	0.858	0.920	1.002***
	FAVAR(1) - 1 Factors	0.977	1.046***	0.997**	0.923	0.899*	0.885**	0.894*	0.939	0.974
	FAVAR(1) - 2 Factors	0.991	1.048***	1.001***	0.944	0.935*	0.885**	0.879	0.957	1.006***

Notes: For the RW-D model, we report the RMSPE, while for all other models, we present the RMSPE ratio relative to the RW-D model. RMSPE refers to the real price in USD per metric ton, with February 2015 CPI as base year. A ratio below one (in bold) indicates that the model outperforms the RW-D model. Asterisks denote statistical significance at the 1% (***), 5% (**) and 10% (*) based on the Diebold and Mariano (1995) test. Light blue cells indicate models included in the Set of Superior Models (SSM) at $\alpha = 0.25$ with the MCS procedure.

Figure 3: First (left) and second (right) estimated factors from the FAVAR model for real copper price



at several horizons. However, for aluminum the FAVAR specifications tend to underperform the best ARDI and simple VAR models. The VAR with inventories consistently beats RW-D and often ranks among the better performers, although it does not dominate across all horizons. The first factor in the ARDI and FAVAR models loads strongly on aggregate

manufacturing activity (industrial production, manufacturing hours and producer prices), summarising a demand-side component that aligns with the univariate results. The second factor assigns more weight to inventory volumes and co-movements with other base metals, capturing a supply-side or market-structure element.

Quantitatively, at short horizons (up to 3 months) and at the longest horizons (21–24 months), most models deliver only limited gains relative to RW-D, typically on the order of 8–10%. At medium-term horizons, however, aluminum price forecasts benefit more clearly from univariate models linked to capacity utilisation and from factor models that combine industrial demand and inventory information.

Table 4: Zinc and nickel real prices point forecasts based on different models and horizons.

Metal	Model	Horizon 1	Horizon 3	Horizon 6	Horizon 9	Horizon 12	Horizon 15	Horizon 18	Horizon 21	Horizon 24
Zinc	RW-D	153.58	297.88	448.55	548.23	633.02	714.20	782.88	817.70	850.19
	Futures	-	0.973	-	-	-	0.863	-	-	-
	Professional Forecasters	-	-	1.017***	0.917	0.850	0.800*	0.769**	-	-
	ARDL(1) - Ec. Activity	1.037***	1.052***	1.049***	1.076***	1.110***	1.119***	1.128***	1.149***	1.104***
	ARDL(1) - Ex. Rates	0.993	1.003***	1.027***	1.045***	1.059***	1.046***	1.022***	0.977	0.939
	ARDL(1) - ETM Prices	1.007***	1.026***	1.134***	1.145***	1.212***	1.147***	1.054***	0.954	0.919
	ARDL(1) - ET	0.986	0.969	0.925*	0.897	0.883	0.889	0.909	0.940	0.950
	ARDL(1) - CU	1.014***	1.071***	1.190***	1.226***	1.219***	1.135***	1.079***	1.033***	1.011***
	ARDL(1) - Inventories	1.002***	1.038***	1.081***	1.045***	1.008***	0.983	0.958	1.046***	1.058***
	VAR(1)	1.005***	0.995	0.987	0.963	0.923**	0.895***	0.913**	1.012***	1.094***
	ARDI(1) - 1 Factors	0.973	0.955*	0.951	0.946*	0.925**	0.917**	0.935	0.959	0.955
	ARDI(1) - 2 Factors	0.955*	0.951*	0.966	0.950	0.913**	0.899**	0.910	0.925	0.937
	FAVAR(1) - 1 Factors	0.960	0.962	0.991	0.988	0.960	0.939	0.955	1.058***	1.148***
	FAVAR(1) - 2 Factors	0.957	0.965	1.000***	0.982	0.941	0.922	0.955	1.064***	1.183***
Nickel	RW-D	1557.23	2913.44	3526.04	3610.10	3936.27	4100.56	4675.75	4858.20	5283.20
	Futures	-	0.929	-	-	-	0.960	-	-	-
	Professional Forecasters	-	-	0.961	1.004***	0.955	0.968	0.927	-	-
	ARDL(1) - Ec. Activity	1.077***	1.047***	1.115***	1.195***	1.149***	1.156***	1.129***	1.141***	1.135***
	ARDL(1) - Ex. Rates	1.051***	1.019***	1.096***	1.067***	1.069***	1.091***	1.094***	1.066***	1.074***
	ARDL(1) - ETM Prices	1.054***	1.031***	1.187***	1.211***	1.281***	1.283***	1.226***	1.119***	1.129***
	ARDL(1) - ET	1.068***	1.012***	1.075***	1.010***	1.029***	1.045***	1.052***	1.092***	1.089***
	ARDL(1) - CU	1.065***	1.035***	1.127***	1.126***	1.118***	1.153***	0.972	0.845	1.053***
	ARDL(1) - Inventories	1.074***	1.066***	1.234***	1.249***	1.251***	1.254***	1.209***	1.145***	1.157***
	VAR(1)	1.061***	0.995	1.079***	1.067***	0.984	0.965	0.972	0.945	0.970
	ARDI(1) - 1 Factors	1.040***	0.978	1.052***	1.033***	0.940	0.926	0.889	0.861	1.007***
	ARDI(1) - 2 Factors	1.072***	1.032***	1.139***	1.076***	1.081***	1.105***	1.103***	1.075***	1.084***
	FAVAR(1) - 1 Factors	1.038***	0.987	1.081***	1.078***	0.998**	0.973	0.981	0.940	0.996*
	FAVAR(1) - 2 Factors	1.041***	0.990	1.093***	1.074***	0.987	0.947	0.959	0.903	1.001***

See notes in Table 3

Zinc. The picture is somewhat different for zinc, as shown in the top panel of Table 4. Here, model-free forecasts (futures and professional forecasts) perform relatively well at intermediate horizons and are selected into the SSM for horizons between 12 and 18 months. However, models that exploit macroeconomic information still yield meaningful gains, especially at medium and long horizons.

In particular, zinc forecasts benefit from models incorporating capacity utilisation, new

orders and inventory dynamics, with RMSPE reductions of roughly 12% at the 12-month horizon and up to about 21% at the 24-month horizon. By contrast, at short horizons (1–6 months ahead), the more complex specifications typically perform close to the RW-D benchmark. Capacity utilisation and new orders help anticipate zinc demand pressures and thus provide early signals about zinc consumption. Unlike for copper and aluminum, zinc inventories play an especially prominent role, reflecting their tight link to global galvanising cycles (construction, autos and infrastructure). Inventory behaviour therefore acts as an early indicator of impending supply tightness, with shorter and more frequent inventory cycles than for copper and aluminum.

Energy-transition (ET) indicators alone generate smaller gains than when they are incorporated into a broader factor structure. Including battery manufacturing producer prices and shipments is, however, useful to capture the growing importance of electrification and storage technologies. As with copper and aluminum, factor models (ARDI and FAVAR) deliver additional improvements, particularly in the ARDI specification with two factors. The first factor mirrors demand-driven components and loads on industrial production, new orders and producer price indices, while the second factor is dominated by inventories and co-movements with copper and aluminum. This confirms the value of combining standard macroeconomic indicators with metals-specific variables.

Nickel. Nickel behaves quite differently from the other three metals (bottom panel of Table 4). At short horizons, virtually no model systematically outperforms the RW-D benchmark, and model-free forecasts are often competitive or slightly better. Forecastability improves somewhat at medium horizons (12–18 months ahead), but only a limited set of models achieves statistically significant gains, and the magnitude of these gains remains modest compared with those observed for copper, aluminum and zinc.

The best-performing models for nickel are factor-augmented specifications, which deliver RMSPE reductions of roughly 12–15% at selected horizons. The first factor appears to capture aggregate macroeconomic activity but displays a relatively noisy pattern, making it difficult to associate it with a specific sector. The second factor relies heavily on inventory series and metals price co-movements. Even so, the overall evidence points to nickel being

intrinsically harder to forecast, likely due to its geographically concentrated and technologically fragmented production, together with heightened exposure to geopolitical shocks. These structural features are only imperfectly captured by the macroeconomic and market indicators in our dataset.

Cross-metal comparison and model-free benchmarks. Taken together, the results suggest that univariate models based on a single, well-chosen predictor often deliver the largest gains relative to RW-D, particularly for copper and aluminum. Factor models that summarise information from multiple macroeconomic series further improve forecast accuracy for some metals and horizons, especially when they combine demand-side indicators with inventories and other metals' prices. For zinc, inventories and co-movements with other base metals are crucial, while for nickel the scope for improving on the random walk is limited.

Comparing our model-based forecasts with model-free benchmarks highlights three additional points. First, futures prices are generally not reliable predictors of future spot prices, and their RMSPE is often higher than that of even simple macro-based models. This suggests that market participants do not fully incorporate available information – such as professional forecasts and macroeconomic signals – into futures prices. Second, across most horizons and metals, our models outperform both futures-based forecasts and those issued by professional forecasters, especially at medium-term horizons where macroeconomic predictors are most informative. Third, nickel remains particularly challenging: both our models and professional forecasters deliver relatively weak and often statistically insignificant gains over the RW-D benchmark, underscoring the intrinsic difficulty of forecasting nickel prices.

5 Model Comparison and Forecast Pooling

5.1 Model Confidence Set Information

Given the large number of competing models in our analysis, we use the Model Confidence Set (MCS) procedure of Hansen et al. (2011) to evaluate their relative performance. The MCS iteratively removes the worst-performing specifications until a set of models remains that cannot be statistically distinguished from the best one. The Supplement reports MCS

p -values across horizons, where one and two asterisks denote retention at the $\alpha = 0.10$ and $\alpha = 0.25$ levels, respectively.

A key result is that, for most metals and horizons, the SSM is large and typically contains many models. This indicates substantial model uncertainty: the real-time data do not provide sharp statistical discrimination among specifications. Only for nickel at the 24-month horizon does the SSM shrink meaningfully, reflecting clearer performance differences at very long horizons.

Although the MCS rarely identifies a small set of superior models, the associated p -values allow us to rank specifications. For copper, factor-augmented models (ARDI and FAVAR) consistently receive the highest rankings up to 18 months ahead, reinforcing the importance of broad macroeconomic information. At longer horizons, simpler models based on new orders or producer price indices often perform best. For aluminum, the top-ranked specifications are almost exclusively univariate ARDL models using capacity utilisation and metals-sector new orders, with producer price indices becoming more relevant at longer horizons.

Zinc displays greater heterogeneity: factor models rank highly at very short horizons, ARDI models with energy-transition variables perform well up to one year, and model-free forecasts (futures and surveys) enter the SSM for several medium-term horizons. At longer horizons, producer price indices and metals-sector indicators reappear among the best-performing models. Nickel remains the most difficult metal to forecast: factor models, capacity-utilisation specifications, and model-free forecasts achieve similar MCS rankings at short and medium horizons, while models including metals-sector new orders are favored at very long horizons.

Overall, the MCS results confirm the main finding of Section 4: no single model dominates across metals or horizons, but certain predictors—especially manufacturing new orders and capacity utilisation—repeatedly appear in the best-performing specifications.

5.2 Forecast pooling

It is well known in the forecasting literature that pooling forecasts from multiple models can improve predictive accuracy compared to relying on a single model (Bates and Granger, 1969). Forecast combination leverages the strengths of different models, mitigating the risk of selecting a single, potentially misspecified, model. By averaging or weighting forecasts based on their past performance, combination methods can enhance robustness and reduce forecast uncertainty. This is particularly relevant and useful in our analysis, as the MCS procedure in the previous section revealed that it is difficult to identify a set of models that are clearly superior – particularly for nickel and zinc. As a result, combining forecasts from multiple models may represent an effective strategy.

In this section, we perform model selection in real time using the MCS procedure and then combine the forecasts of the corresponding SSM at a confidence level equal to $\alpha = 0.25$. Specifically, during the first year of the evaluation period, we average the forecasts of all models considered in Section 4. Starting from the thirteenth month, we apply the MCS procedure to the past 12 observations and then construct forecasts using two different approaches: (i) averaging the forecasts of all selected models and (ii) averaging the forecasts of the two best-performing models, as determined by their MCS p-values.

As shown in Table 5, pooling forecasts using both methods consistently yields greater forecast accuracy compared to the RW-D model across all horizons and metals, except for nickel, where the random-walk outperform the pooling forecast up to 10% in the first 6 horizons. As highlighted by the gray cells, we notice that the forecasting pooling provides more accurate predictions than any individual model over the horizon ranging from 9 to 21 horizons depending from the metal.

For copper, however, individual models based on Manufacturers' New Orders for Primary Metals (NO-M) – such as the corresponding ARDL, VAR, and FAVAR specifications – continue to outperform forecast combination strategies in the short and medium term. This result is likely due to the exceptionally strong predictive power of new orders for copper prices relative to other metals, combined with the weak forecasting performance of other macroeconomic predictors. Indeed, the gap in forecasting accuracy between models based

on new orders and those based on other predictors is particularly pronounced for copper compared to the other three metals. As a result, pooling forecasts in the case of copper may mitigate the signal provided by a key indicator of planned production by incorporating information from variables with limited predictive content, thus introducing additional noise. Indeed, it is important in the long-term (from 21 horizons), when doing forecasting combination outperform the best individual models. To wrap up, we notice that in the long-term adding a combination tool leads to better results with respect to single models thus highlighting the importance of using different variables but also different models in the forecasting exercise when less information are available.

Table 5: Point forecasting measures by combining the forecasting based on MCS procedures.

Model	Horizon 1	Horizon 3	Horizon 6	Horizon 9	Horizon 12	Horizon 15	Horizon 18	Horizon 21	Horizon 24
Copper									
RW-D	292	610	918	1139	1363	1528	1628	1635	1625
SSM25	0.972	0.923**	0.845**	0.753**	0.744***	0.813***	0.823***	0.863***	0.882**
Top 2	0.974	0.897**	0.816**	0.742**	0.746***	0.824***	0.833***	0.831***	0.852**
Aluminum									
RW-D	90	192	277	351	425	483	525	529	523
SSM25	0.939	1.008***	0.922*	0.829***	0.819**	0.762**	0.772**	0.845***	0.906**
Top 2	0.938	1.015***	0.906*	0.805**	0.769**	0.761**	0.762**	0.800**	0.914**
Nickel									
RW-D	1557	2913	3526	3610	3936	4101	4676	4858	5283
SSM25	1.077***	1.039***	1.090***	1.035***	0.922*	0.958	0.814	0.759	0.926
Top 2	1.075***	1.038***	1.005***	0.976	0.889*	0.928	0.798*	0.685	0.961
Zinc									
RW-D	154	298	449	548	633	714	783	818	850
SSM25	0.989	0.988	0.981	0.909**	0.816***	0.787***	0.758***	0.779**	0.803*
Top 2	0.993	0.976	0.936*	0.864***	0.836**	0.811***	0.775***	0.747***	0.761*

Note: For the RW-D model, we report the RMSPE, while for all other models, we present the RMSPE ratio relative to the RW-D model. RMSPE refers to the real price in USD per metric ton, with February 2015 CPI as base year. A ratio below one indicates that the model outperforms the RW-D model. Asterisks denote statistical significance at the 1% (***), 5% (**) and 10% (*) based on the Diebold and Mariano (1995) test. Forecast combinations that outperform the RW-D model are highlighted in bold. Grey cells indicate forecast combinations that outperform the best individual model at the corresponding horizon.

5.3 Performance over time

All forecasting statistics analyzed so far evaluate model performance by averaging squared predictive errors over the entire out-of-sample period. However, the presence of structural breaks can significantly impact model performance over time, leading different models to

perform better in different subperiods. To account them, Figures 4, 5, and 6¹¹ illustrate the time evolution of the RMSPE ratios of various models relative to the RW-D benchmark at a 6-, 12- and 18-month horizon, respectively¹². A downward (upward) sloping line indicates that a given model is performing better (worse) than the benchmark at that point in time, while a flat line suggests comparable performance between the model and the RW-D. The value of the ratio at any given time indicates the average forecasting performance of the model up to that point.

In the short term (6 horizons), we notice that forecasting combination does not perform the other models along time for all the metals considered. For copper, models that include new orders as a single variable or as a factor outperform the random walk and improve after the COVID-19 outbreak. The same happens for aluminum, where the model with capacity utilization outperforms the other models and is in line with the forecasting combination methods after the COVID-19 outbreak. For nickel and zinc (bottom panels), the situation is chaotic with models closed to the benchmark model across time and no big differences after 2020–2021.


Moving to the medium horizon (see Figure 5), we notice strong co-movements after 2019 between Factor augmented, VAR, and ARDI with new orders models. Similarly, the two forecasting combination models appear to react similarly to shocks. For aluminum, the forecasting combination shows important results along time appearing the best model with similarity with ARDI models with capacity utilization. These findings highlight the importance of including different models for those metals since they capture business cycle momentum and expected future industrial demand. In the bottom panel, nickel remains unpredictable and it seems that using a simple model, such as the RW-D, is providing good results. However, zinc seems to improve moderately when forecasting combination is considered, although professional forecasts seems to outperform all models.

In the long horizon (see Figure 6), we notice that forecasting combinations increasingly outperform individual models for all the metals. Indeed, for copper (top left panel), forecast pooling strategies exhibit markedly improved performance, with cumulative ratios lower

¹¹We provide evidence for the other horizons where model-free forecasts are available in the Supplement.

¹²To allow the RMSPE ratio to stabilize, we skip the first 12 forecast periods.


Figure 4: Cumulative RMSPE ratios for different models (colored lines) and metals (panel) at horizon 6.



than those of any individual model for most of the sample, and in particular, they capture strongly the post-COVID price surge. Similar results are available for aluminum (top right), where forecast pooling – particularly when based on the two best-performing MCS models – outperforms all other forecasting strategies.


For nickel (bottom left), no model or group of models clearly dominates across the entire evaluation period. Our models exhibit strong performance in the early part of the sample but perform worse than the benchmark thereafter, ultimately converging towards a ratio close to, but still below, one. On the contrary, professional forecasts show poor performance

Figure 5: Cumulative RMSPE ratios for different models (colored lines) and metals (panel) at horizon 12.



relative to the RW-D benchmark at the beginning of the evaluation period but improve over time – as reflected by a declining cumulative ratio. A notable exception to these patterns is the forecast pooling strategies, which maintains a cumulative ratio below one for most of the sample. For zinc (bottom right), at the 18-month horizon, the best-performing strategies are the forecast pooling approaches, which achieve forecast error reductions of approximately 20% relative to individual models over the whole sample. Professional forecasters deliver forecasting performance comparable to that of forecast pooling over the entire evaluation period except the 2022–2023 period.

Figure 6: Cumulative RMSPE ratios for different models (colored lines) and metals (panel) at horizon 18.



Overall, this time-varying forecast evaluation leads to two key takeaways. First, forecast pooling is significantly more effective at longer horizons than at shorter ones, although it still yields cumulative ratios well below one even at shorter horizons. Second, at longer horizons, forecast pooling consistently emerges as the best forecasting strategy across the entire evaluation period, with the partial exception of copper.

6 Conclusion

This paper develops a real-time forecasting framework for monthly real prices of key metals, addressing a gap in the forecasting literature where metal markets have received considerably less systematic attention than energy markets. By constructing a real-time database of macroeconomic predictors and evaluating a broad set of forecasting models under true real-time information constraints, the paper clarifies which predictors and model classes offer genuine forecasting gains for aluminum, copper, nickel, and zinc. Our results deliver three main messages.

First, real-time macroeconomic information substantially improves forecasting accuracy at medium-term horizons, whereas short-term dynamics remain extremely difficult to predict. The dominant predictors are demand-side indicators from the primary metals manufacturing sector—especially new orders and capacity utilisation – which consistently outperform other macro-financial variables. This provides clear evidence that medium-term metal price dynamics reflect broader industrial activity rather than idiosyncratic noise.

Second, forecast performance is heterogeneous across metals. Copper and aluminum exhibit robust medium-term predictability; zinc displays moderate improvements driven by inventories and energy-transition variables; and nickel remains intrinsically difficult to forecast, with only modest gains from factor-augmented models. These differences highlight structural variation across metal markets that forecasting models must accommodate.

Third, while individual models based on metals-sector indicators often perform best at specific horizons, model uncertainty is pervasive. Forecast pooling, implemented via the Model Confidence Set, delivers accuracy gains for nearly all metals, particularly over longer horizons, where no single model dominates. A notable exception arises for copper, where the predictive performance of new orders is sufficiently strong that pooling may on the contrary not be necessary.

These results contribute to the forecasting literature in two ways. Methodologically, the paper demonstrates the value of constructing forecasting models using real-time data, an approach rarely applied to metal markets. Substantively, it provides evidence that medium-term metal prices are forecastable using publicly available macroeconomic indicators, offering

a transparent alternative to survey and futures-based expectations, both of which perform consistently worse in our evaluation.

The analysis also highlights several limitations. The availability of metal-specific monthly indicators remains limited; the forecasting environment is strongly affected by geopolitical disruptions and structural breaks; and interdependencies across metals are only partially captured by the models considered.

References

- Alquist, R., Bhattarai, S., and Coibion, O. (2020). Commodity-price comovement and global economic activity. *Journal of Monetary Economics*, 112:41–56.
- Alquist, R. and Kilian, L. (2010). What do we learn from the price of crude oil futures? *Journal of Applied Econometrics*, 25(4):539–573.
- Alquist, R., Kilian, L., and Vigfusson, R. J. (2013). Forecasting the price of oil. In Elliott, G. and Timmermann, A., editors, *Handbook of Economic Forecasting*, volume 2A, chapter 8, pages 427–507. Elsevier.
- Ang, A., Bekaert, G., and Wei, M. (2007). Do macro variables, asset markets, or surveys forecast inflation better? *Journal of Monetary Economics*, 54(4):1163–1212.
- Bastianin, A., Casoli, C., and Galeotti, M. (2023). The connectedness of energy transition metals. *Energy Economics*, 128:107183.
- Bastianin, A., Galeotti, M., and Manera, M. (2014). Forecasting the oil–gasoline price relationship: Do asymmetries help? *Energy Economics*, 46:S44–S56.
- Bates, J. M. and Granger, C. W. (1969). The combination of forecasts. *Journal of the Operational Research Society*, 20(4):451–468.
- Baumeister, C., Huber, F., Lee, T. K., and Ravazzolo, F. (2025). Forecasting natural gas prices in real time. *Journal of Applied Econometrics*, forthcoming.
- Baumeister, C. and Kilian, L. (2012). Real-time forecasts of the real price of oil. *Journal of Business & Economic Statistics*, 30(2):326–336.
- Baumeister, C. and Kilian, L. (2015). Forecasting the real price of oil in a changing world: A forecast combination approach. *Journal of Business & Economic Statistics*, 33(3):338–351.
- Baumeister, C., Kilian, L., and Lee, T. K. (2017). Inside the crystal ball: new approaches to predicting the gasoline price at the pump. *Journal of Applied Econometrics*, 32(2):275–295.
- Baumeister, C., Verduzco-Bustos, G., and Ohnsorge, F. (2022). Pandemic, war, recession: Drivers of aluminum and copper prices. Special Focus, World Bank Commodity Markets Outlook.
- Boer, L., Pescatori, A., and Stuermer, M. (2024). Energy transition metals: bottleneck for net-zero emissions? *Journal of the European Economic Association*, 22(1):200–229.

- Carriero, A., Clark, T., Marcellino, M., and Mertens, E. (2024). Addressing COVID-19 Outliers in BVARs with Stochastic Volatility. *Review of Economics and Statistics*, 106(5):1403–1417.
- Carriero, A., Clark, T. E., and Marcellino, M. (2015). Bayesian vars: Specification choices and forecast accuracy. *Journal of Applied Econometrics*, 30(1):46–73.
- Carriero, A., Clark, T. E., and Marcellino, M. (2019). Large bayesian vector autoregressions with stochastic volatility and non-conjugate priors. *Journal of Econometrics*, 212(1):137–154.
- Chan, J. C. (2023). Comparing stochastic volatility specifications for large Bayesian VARs. *Journal of Econometrics*, 235(2):1419–1446.
- Chan, J. C. C. (2020). Large bayesian vars: A flexible Kronecker error covariance structure. *Journal of Business & Economic Statistics*, 38(1):68–79.
- Chen, Y.-C., Rogoff, K. S., and Rossi, B. (2010). Can exchange rates forecast commodity prices? *The Quarterly Journal of Economics*, 125(3):1145–1194.
- Clark, T. E. and McCracken, M. W. (2015). Nested forecast model comparisons: A new approach to testing equal accuracy. *Journal of Econometrics*, 186(1):160–177.
- Considine, J., Galkin, P., Hatipoglu, E., and Al Dayel, A. (2023). The effects of a shock to critical minerals prices on the world oil price and inflation. *Energy Economics*, 127B(106934).
- Croushore, D. (2010). An evaluation of inflation forecasts from surveys using real-time data. *The B.E. Journal of Macroeconomics*, 10(1).
- Deaton, A. and Laroque, G. (1992). On the behaviour of commodity prices. *The Review of Economic Studies*, 59(1):1–23.
- Deaton, A. and Laroque, G. (1996). Competitive storage and commodity price dynamics. *Journal of Political Economy*, 104(5):896–923.
- Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):253–263.
- Dovern, J., Fritsche, U., and Slacalek, J. (2012). Disagreement among forecasters in G7 countries. *Review of Economics and Statistics*, 94(4):1081–1096.
- Duarte, A. M., Gaglianone, W. P., de Carvalho Guillén, O. T., and Issler, J. V. (2021). Commodity prices and global economic activity: A derived-demand approach. *Energy Economics*, 96:105120.

- Ferrari, D., Ravazzolo, F., and Vespignani, J. (2021). Forecasting energy commodity prices: A large global dataset sparse approach. *Energy Economics*, 98:105268.
- Foroni, C., Ravazzolo, F., and Rossini, L. (2023). Are low frequency macroeconomic variables important for high frequency electricity prices? *Economic Modelling*, 120:106160.
- Gargano, A. and Timmermann, A. (2014). Forecasting commodity price indexes using macroeconomic and financial predictors. *International Journal of Forecasting*, 30(3):825–843.
- Gianfreda, A., Ravazzolo, F., and Rossini, L. (2020). Comparing the forecasting performances of linear models for electricity prices with high res penetration. *International Journal of Forecasting*, 36(3):974–986.
- Hansen, P. R., Lunde, A., and Nason, J. M. (2011). The model confidence set. *Econometrica*, 79(2):453–497.
- Hauzenberger, N., Pfarrhofer, M., and Rossini, L. (2025). Sparse time-varying parameter vecms with an application to modeling electricity prices. *International Journal of Forecasting*, 41(1):361–376.
- IEA (2025a). Global critical minerals outlook 2025. IEA, Paris. Available online at: <https://www.iea.org/reports/global-critical-minerals-outlook-2025>.
- IEA (2025b). The role of traceability in critical mineral supply chains. IEA, Paris. Available online at: <https://www.iea.org/reports/the-role-of-traceability-in-critical-mineral-supply-chains>.
- Issler, J. V., Rodrigues, C., and Burjack, R. (2014). Using common features to understand the behavior of metal-commodity prices and forecast them at different horizons. *Journal of International Money and Finance*, 42:310–335.
- Kilian, L. and Murphy, D. P. (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, 29(3):454–478.
- Marcellino, M., Stock, J. H., and Watson, M. W. (2006). A comparison of direct and iterated multistep ar methods for forecasting macroeconomic time series. *Journal of Econometrics*, 135(1):499–526.
- McCracken, M. W. and Ng, S. (2016). Fred-md: A monthly database for macroeconomic research. *Journal of Business & Economic Statistics*, 34(4):574–589.
- Pierdzioch, C., Rülke, J.-C., and Stadtmann, G. (2013). Forecasting metal prices: Do forecasters herd? *Journal of Banking & Finance*, 37(1):150–158.

- Ravazzolo, F. and Rothman, P. (2013). Oil and u.s. gdp: A real-time out-of-sample examination. *Journal of Money, Credit and Banking*, 45(2-3):449–463.
- Ruben, A., Bjørnland, H. C., and Eliassen, P. (2025). Forecasting oil and natural gas prices: A model combination approach. CAMA Working Papers 2025-54, Centre for Applied Macroeconomic Analysis, Crawford School of Public Policy, The Australian National University.
- Stock, J. H. and Watson, M. W. (2002). Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association*, 97(460):1167–1179.
- Stuermer, M. (2017). Industrialization and the demand for mineral commodities. *Journal of International Money and Finance*, 76:16–27.
- Wallis, K. F. (1986). Forecasting with an econometric model: The ‘ragged edge’ problem. *Journal of Forecasting*, 5(1):1–13.
- Watkins, C. and McAleer, M. (2004). Econometric modelling of non-ferrous metal prices. *Journal of Economic Surveys*, 18(5):651–701.
- West, K. D. and Wong, K. F. (2014). A factor model for co-movements of commodity prices. *Journal of International Money and Finance*, 42:289–309.
- Wittenstein, M. and Stuermer, M. (2023). Why we need international data sharing on critical minerals. World Economic Forum. Available online at: <https://www.weforum.org/stories/2023/12/why-we-need-international-data-sharing-on-critical-minerals/>.

Our Working Papers are available on the Internet at the following address:

<https://www.feem.it/pubblicazioni/feem-working-papers/>

“NOTE DI LAVORO” PUBLISHED IN 2025

1. M. A. Marini, S. Nocito, Climate Activism Favors Pro-environmental Consumption
2. J. A. Fuinhas, A. Javed, D. Sciulli, E. Valentini, Skill-Biased Employment and the Stringency of Environmental Regulations in European Countries
3. A. Stringhi, S. Gil-Gallen, A. Albertazzi, The Enemy of My Enemy
4. A. Bastianin, X. Li, L. Shamsudin, Forecasting the Volatility of Energy Transition Metals
5. A. Bruni, Green Investment in the EU and the US: Markup Insights
6. M. Castellini, C. D'Alpaos, F. Fontini, M. Moretto, Optimal investment in an energy storage system
7. L. Mauro, F. Pigliaru, G. Carmeci, Government Size, Civic Capital and Economic Performance: An O-ring approach
8. A. Bastanin, C. F. Del Bo, L. Shamsudin, The geography of mining and its environmental impact in Europe
9. E. Cavallotti, I. Colantone, P. Stanig, F. Vona, Green Collars at the Voting Booth: Material Interest and Environmental Voting
10. L. Ciambezi, M. Guerini, M. Napoletano, A. Roventini, Accounting for the Multiple Sources of Inflation: an Agent-Based Model Investigation
11. L. Fontanelli, M. Guerini, R. Miniaci, A. Secchi, Predictive AI and productivity growth dynamics: evidence from French firms
12. A. Drigo, An Empirical Analysis of Environmental and Climate Inequalities across Italian census tracts
13. F.F. Frattini, F. Vona, F. Bontadini, I. Colantone, The Local Job Multipliers of Green Industrialization
14. A. Abatemarco, R. Dell'Anno, E. Lagomarsino, Measuring equity in environmental care: methodology and an application to air pollution
15. F. F. Frattini, Political Participation and Competition in Concurrent Elections: Evidence from Italy
16. M. Bonacina, M. Demir, A. Sileo, A. Zanoni, What Hinders Electric Vehicle Diffusion? Insights from a Neural Network Approach
17. C. A. Bollino, M. Galeotti, Native-borns and migrants do not contribute equally to domestic CO₂ emissions
18. D. Legrenzi, E. Ciola, D. Bazzana, Adaptation to climate-induced macrofinancial risks: top-down and bottom-up solutions
19. M. Bonacina, R. C. Tokong, Is Italy on Track? A Data-Driven Forecast for Road Transport Decarbonisation by 2030
20. B. P. Bosco, C. F. Bosco, P. Maranzano, Income Tax Treatment and Labour Supply in a multi-level hierarchical Difference-in-Differences model
21. S. Ciucci, The effect of a lottery on collusion sustainability
22. G. Benini, E. Enstad, A. A. Mersha, L. Rossini, Shadow Price Signals in the Steel Sector: From Efficiency Gaps to Policy Maps
23. C. Casoli, M. Manera, L. Pedini, D. Valenti, “It’s not the heat, it’s the humidity!” New Climate Indices for Europe with a Multilevel Factor Model
24. J. Myftiu, D. Crotti, E. Maggi, Road safety in urban sustainable cycling tourism. An application of the Protection Motivation Theory
25. F. Ravazzolo, L. Rossini, A. Viselli, Modeling European electricity market integration during turbulent times
26. S. Ciucci, The tax rebate trade-off
27. J. Kitsnik, Ethical Considerations of the Wealth Gap Sentiments in Europe
28. G. Marin, A. T. Oo, Sectoral exposure to heat: heterogeneous impacts of extreme heat on workplace accidents in Italy
29. C. Falco, V. Raimondi, Hydropower dams, deforestation, and land use change: evidence from Brazil
30. F. Zilia, P. Nota, A. Olper, Weather Shocks and Sectoral Labour Reallocation: Evidence from the European Regions
31. M. Guerini, G. Marin, F. Vona, Easing Financial Constraints Reduce Carbon Emissions? Evidence from a Large Sample of French Companies
32. A. Bellocchi, C. Lodi, G. Marin, G. Travaglini, M. Zavalloni, Floods, Public Budgets and Fiscal Resilience: Evidence from Italian Municipalities
33. F. Rubbi, Ecosystem services demand in social-ecological systems: towards a typology of assessment approaches

Fondazione Eni Enrico Mattei

Corso Magenta 63, Milano – Italia

Tel. +39 02 403 36934

E-mail: letter@feem.it

www.feem.it

