



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

February 2024



Working Paper

02.2024

What drives the European carbon market? Macroeconomic factors and forecasts

**Andrea Bastianin, Elisabetta Mirto
Yan Qin, Luca Rossini**

What drives the European carbon market? Macroeconomic factors and forecasts

Andrea Bastianin (University of Milan and Fondazione Eni Enrico Mattei); **Elisabetta Mirto** (University of Milan); **Yan Qin** (London Stock Exchange Group); **Luca Rossini** (University of Milan and Fondazione Eni Enrico Mattei)

Summary

Putting a price on carbon – with taxes or developing carbon markets – is a widely used policy measure to achieve the target of net-zero emissions by 2050. This paper tackles the issue of producing point, direction-of-change, and density forecasts for the monthly real price of carbon within the EU Emissions Trading Scheme (EU ETS). We aim to uncover supply- and demand-side forces that can contribute to improving the prediction accuracy of models at short- and medium-term horizons. We show that a simple Bayesian Vector Autoregressive (BVAR) model, augmented with either one or two factors capturing a set of predictors affecting the price of carbon, provides substantial accuracy gains over a wide set of benchmark forecasts, including survey expectations and forecasts made available by data providers. We extend the study to verified emissions and demonstrate that, in this case, adding stochastic volatility can further improve the forecasting performance of a single-factor BVAR model. We rely on emissions and price forecasts to build market monitoring tools that track demand and price pressure in the EU ETS market. Our results are relevant for policymakers and market practitioners interested in quantifying the desired and unintended macroeconomic effects of monitoring the carbon market dynamics.

Keywords: Bayesian inference, Carbon prices, Climate Changes, EU ETS, Forecasting.

JEL classification: C11, C32, C53, Q02, Q50

Address for correspondence:

Andrea Bastianin
Program Director, RAw Materials for the Energy Transition (RAMET)
Fondazione Eni Enrico Mattei (FEEM)
Corso Magenta 63, 20123 Milan (Italy)
e-mail: andrea.bastianin@feem.it

What drives the European carbon market? Macroeconomic factors and forecasts*

Andrea Bastianin[†] Elisabetta Mirto[‡] Yan Qin[§] Luca Rossini[¶]

February 7, 2024

Abstract: Putting a price on carbon – with taxes or developing carbon markets – is a widely used policy measure to achieve the target of net-zero emissions by 2050. This paper tackles the issue of producing point, direction-of-change, and density forecasts for the monthly real price of carbon within the EU Emissions Trading Scheme (EU ETS). We aim to uncover supply- and demand-side forces that can contribute to improving the prediction accuracy of models at short- and medium-term horizons. We show that a simple Bayesian Vector Autoregressive (BVAR) model, augmented with either one or two factors capturing a set of predictors affecting the price of carbon, provides substantial accuracy gains over a wide set of benchmark forecasts, including survey expectations and forecasts made available by data providers. We extend the study to verified emissions and demonstrate that, in this case, adding stochastic volatility can further improve the forecasting performance of a single-factor BVAR model. We rely on emissions and price forecasts to build market monitoring tools that track demand and price pressure in the EU ETS market. Our results are relevant for policymakers and market practitioners interested in quantifying the desired and unintended macroeconomic effects of monitoring the carbon market dynamics.

Key Words: Bayesian inference; Carbon prices; Climate Changes; EU ETS; Forecasting.

JEL Codes: C11; C32; C53; Q02; Q50.

*The authors gratefully acknowledge the participants at the 4th IWEEE in Bolzano for their useful feedback. Luca Rossini acknowledges financial support from the Italian Ministry MIUR under the PRIN-PNRR project Mapping and Pricing of Methane Emissions from the European Electricity Sector (MAP-of-MeLEES) (grant P2022H483A). This research used the Computational resources provided by the Core Facility INDACO, which is a project of High-Performance Computing at the University of Milan.

[†]University of Milan, Italy and Fondazione Eni Enrico Mattei (FEEM). andrea.bastianin@unimi.it

[‡]University of Milan, Italy. elisabetta.mirto@unimi.it

[§]London Stock Exchange Group, Norway. yan.qin@lseg.com

[¶]University of Milan, Italy and Fondazione Eni Enrico Mattei (FEEM). luca.rossini@unimi.it

1 Introduction

Climate change is one of the greatest challenges currently addressed by governments, central banks, and other national and supranational regulators. The size and complexity of the call to fight climate change are evident, even when – setting aside its environmental consequences – we focus solely on its direct socio-economic impacts (Carleton and Hsiang, 2016). The literature has highlighted a multitude of effects and transmission channels that connect the environmental influences of climate change to various economic aggregates (Dell et al., 2014; Ciccarelli and Marotta, 2024). Moreover, when designing mitigation and adaptation policies, a trade-off emerges between the stringency of measures and the undesired outcomes that these might induce (Fullerton and Muehlegger, 2019; Hsiang et al., 2019).

Among the policy measures implemented to achieve the target of net-zero emissions by 2050 – as laid down in the 2015 Paris Agreement – carbon taxes and carbon pricing have lately received a lot of attention from academics and policymakers interested in quantifying the desired and unintended macroeconomic effects of putting a price on carbon (see e.g. European Central Bank, 2021a; Parry et al., 2021; Pan et al., 2023). Given this literature and the need to incorporate climate and carbon market modules into macroeconomic models used by central banks, having access to reliable short- and medium-term forecasts of the price of carbon is becoming increasingly important (European Central Bank, 2021b; NGFS, 2022).

This paper focuses on producing point, direction-of-change, and density forecasts of the real price of carbon at a monthly sampling frequency in the world’s largest carbon market: the EU Emission Trading Scheme (EU ETS). Our research pursues two intertwined targets. First, we aim to uncover supply- and demand-side forces that can contribute to improving the prediction accuracy of models at short- and medium-term horizons. The second main objective of the paper is to highlight which methodological choices have the potential to improve the forecast accuracy of the estimated models.

Since the inception of the EU ETS in 2005, the research on carbon pricing has addressed a variety of issues (see Chevallier, 2012; Pan et al., 2023, for a survey). In particular, Bjørnland et al. (2023) and Känzig (2023) use macroeconomic and financial data, along with microdata to evaluate the impact of carbon price shocks. Recently Känzig and Konradt (2023) and

Moessner (2022) focus on European countries and the inflationary effects in a set of OECD countries, respectively.

Our paper directly contributes to the literature dealing with forecasting the price of carbon and identifying its determinants such as weather, energy prices, macroeconomic and financial conditions (see e.g. Chevallier, 2011a,b; Koop and Tole, 2013; Lei et al., 2022; Tan et al., 2022). Whether the results that emerge from this strand of the literature on carbon price forecasting – mostly based on daily and weekly data – can be directly applied to building monthly or quarterly econometric models through appropriate time aggregation of variables has to be empirically determined.

Differently from the literature, we focus on the real monthly price of carbon, which enables us to provide results that are relevant to the dialogue about the macroeconomic consequences of carbon pricing. Indeed, our analysis can be closely related to Chevallier (2011b), which evaluate the role of energy prices and business cycle movements in driving the price of carbon, and to Bjørnland et al. (2023), which use a Bayesian structural vector autoregressive (VAR) model with endogenous variables (i.e. real price of carbon, verified emissions, and industrial production) to identify demand and supply shocks driving the emissions and the real price of carbon. Moreover, while most of the papers dealing with carbon price forecasting rely on frequentist methods, we use Bayesian techniques to readily incorporate Stochastic Volatility (SV) dynamics into the models and to generate density forecasts.¹ Reliance on SV models is another novelty of our paper; in fact, several previous analyses have focused either on conditionally homoskedastic models or have modeled conditional volatility dynamics using Generalized AutoRegressive Conditional Heteroskedasticity models.

After having assessed a suite of benchmark univariate time series models, we consider small-scale Vector Autoregressive (VAR) models with endogenous variables capturing forces affecting the real price of carbon. First, we focus on the model put forth by Bjørnland et al. (2023) as a starting point. Next, we extend their VAR model to include factors capturing the influence of multiple predictors that affect the EU ETS real price.

¹One notable exception is Koop and Tole (2013) that also rely on Bayesian methods but consider daily data.

While we initially compare models based on their ability to deliver accurate point forecasts, we subsequently extend these results in several directions. The first extension evaluates different specifications for their ability to yield accurate sign forecasts and prediction densities. Directional accuracy is mostly relevant at short-term forecast horizons, while accurate prediction densities are useful for assessing the uncertainty surrounding point forecasts and for quantifying the probability of extreme price movements both in the short- and medium-run. Next, we examine the role of time-varying volatility in improving forecast accuracy (Clark and Ravazzolo, 2015; Chan, 2023) and assess the presence of forecast instabilities (Rossi, 2021).

We show that a simple VAR model, augmented with either one or two factors capturing key predictors of the price of carbon, provides substantial accuracy gains over a wide set of benchmark forecasts, including survey expectations and forecasts made available by data providers. We extend the forecasting study to verified emissions and demonstrate that, in this case, SV can further improve the performance of a single-factor VAR model. Lastly, we show how model-based forecasts can be used to build market monitoring tools that track demand and price pressure in the EU ETS market.

The rest of the paper is organized as follows. Section 2 offers an overview of the EU ETS; data and details of the forecasting exercise are illustrated in Section 3. Results for real price and verified emissions are presented in Sections 4 and 5, respectively. Section 6 concludes.

2 The EU Emissions Trading System

The EU ETS is a cap and trade system that started in 2005 intending to reduce carbon emissions. In this system, the maximum quantity of emissions (the cap) is set through unit permits (*European Unit Allowances*, EUA), which allow the owner of the permit to produce 1 ton of CO₂ or an equivalent quantity of other greenhouse gases. The European Commission sets a yearly cap on the total greenhouse gas emissions that can be produced by actors participating in the system. Since the aim is to decrease emissions over time, every year the cap is lower than the year before, and consequently, the maximum allocation of EUA is

reduced.

The EU ETS system is a financial market where actors can acquire EUA on the primary market through an auctioning system, and trade derivatives on the secondary market. A certain amount of permits is originally granted for free each year according to the needs of specific sector emissions, although the remaining amount of available allowances is allocated on the primary market through uniform price auctions with single rounds and sealed bids, conducted daily by the *European Energy Exchange* (EEX). Since EUA have been classified as financial instruments, the associated derivatives - such as spot, futures, options, and forward contracts - can be traded on secondary markets, both on exchange and over the counter. While auctions take place on the EEX, trading takes place also on the Intercontinental Exchange (ICE).

EUA are handed out to the market through a system of benchmark-based allocation or auctions. If emissions at the end of the year result to be lower than the cap set for each installation participating in the market, permits can be traded among actors for an economic value to be determined on the secondary market. In case emissions exceed the threshold, sanctions are applied to economic agents participating in the market.

Actors participating in the EU ETS market entail industries belonging to high emissions sectors: electricity and heat generation, energy-intensive industry sectors (including oil refineries, steel works, and production of iron, aluminum, metals, paper, etc.), aviation within the European Economic Area and, starting from 2024, maritime transport. Participation in the EU ETS is mandatory for companies in the covered sectors, however, for some of the sectors only production plants bigger than a threshold size are included.


Historically, the EU ETS evolution went through four phases. To meet the objectives set by the European Commission in terms of emissions reduction, each phase aims to reduce the number of EUA granted to each participating sector. The cap can be lowered by setting a decreasing number of allowances to be allocated each year or by establishing a yearly linear reduction factor, e.g. a linear reduction of 1.74% and 2.2% of the baseline 2008-2012 emissions have been set respectively from 2013 and 2021 onwards, with no end date, resulting in a year-on-year reduction by up to 43 million allowances (International Carbon Action Partnership, 2023).

The pilot phase (2005-2007) aimed to verify rules, regulations, emission detection systems, as well as the regulatory framework. In this phase, the allocation system was *grandfathering*: all EUA were allocated freely to industries, up to the cap set for each regulated sector. The second phase, which lasted from 2008 to 2012, was characterized by the introduction of the allocation of permits employing an auctioning system. In this phase, roughly 2 up to 5% of the total permits were allocated through auctions. This share increased to reach 54% in the third phase, which lasted from 2013 to 2020, and it includes more sectors and gases. Lastly, the fourth phase (2021-2030) has the aim of reducing net emissions by at least 55% by 2030 compared to 1990, as set in the European Climate Law, by further lowering the cap and targeting the *carbon leakage* phenomenon.

In July 2021, the European Commission adopted a series of legislative proposals regarding EU ETS aimed at increasing the pace of emissions cuts. These include, among others, covering more sectors and gases, gradually lowering the number of emission allowances each year, and reinforcing the Market Stability Reserve (MSR), which aims at reducing the surplus of allowances in the market. An excessive allowances surplus would lead to lower carbon prices, rendering, therefore, the ETS system less effective, by decreasing incentives of the economic actors participating in the market to lower emissions. The MSR is automatically applied when the total number of allowances on the market exceeds a certain threshold. In Phase IV, the free EUA allocation system was granted a ten-year extension, and specific measures have been taken for sectors exposed to a higher risk of *carbon leaking*.

The way allowance allocation has evolved throughout the years is described in Figure 1. In Phase II, auctions were introduced, but between 95% and 97% of allowances were still distributed for free. This changed completely during Phase III, where auctions became the predominant allocation method for most sectors, covering between 40% and 55% of the total allocated permits. Phase III also needed to deal with the aftermath of an excessive surplus of unused EUA in Phase II. The number of allocated EUA decreased in 2014, increasing the price of allowances and a reduction in the number of verified emissions and unused EUA (Bjørnland et al., 2023). In Phase IV, auctions remain the main allocation method on the primary market.

Figure 1: Allocated and verified emissions for all stationary installations in EU-27 countries (excluding the aviation sector).



Notes: authors' elaboration of data from <https://www.eea.europa.eu>.

3 Data and methods

Following Bjørnland et al. (2023), we focus on forecasting the end-of-month price of the one-month ahead futures contract traded on ICE that represents the most closely watched series by practitioners. Moreover, we deflate the nominal futures price by using the Euro area harmonized index of consumer prices. The left panel of Figure 2 shows the real price series jointly with shaded grey areas representing the recessions in the Euro Area.

As for the predictors, we follow the approach of Boivin and Ng (2006) and Baumeister et al. (2022). Therefore, instead of collecting a large number of series, we carefully select 21 predictors that capture demand and supply-side forces driving the price of carbon. More precisely, we concentrate on variables within the following categories:

- *Economic activity* (8 series): we collect data on aggregate industrial production (IP) for the EU-19 area, as well as indices for sectors covered by the EU ETS (i.e. electricity, gas, steam, and air conditioning supply, basic metals, manufacture of paper and its products, coke and refined petroleum products, chemical products, non-metallic min-

eral products) from Eurostat. Moreover, we consider the Euro Stoxx 50 stock price index, sourced from Refinitiv Eikon.

- *Energy prices* (7 series): we consider the prices of Brent crude oil, TTF natural gas (front-month and front-year), ARA API-2 coal (front-year), German power price (front-year), clean dark, and clean spark spreads (front month). These variables are sourced from Refinitiv Eikon.
- *Technical indicators* (3 series): we select some of the variables that practitioners use to track the functioning of the EU ETS market (Marcu et al., 2023). The auction coverage ratio, defined as the total number of bids in an auction divided by the number of available EUA, is a proxy for the actual auction demand relative to supply on the primary market. As a rule of thumb, a value greater (lower) than two indicates a high (low) auction demand relative to supply. The auction clearing price, and a volatility proxy based on the monthly auction price range. These variables are sourced from the EUA Primary Market Auction Report maintained by EEX.
- *Weather conditions* (2 series): temperature and precipitation anomalies for EU-19 countries are constructed as differences from long-term moving averages using data sourced from the Weather for Energy Tracker maintained by the International Energy Agency (IEA).


Moreover, we also collect information on verified emissions for all stationary installations in six sectors covered by the EU ETS² in the EU-19 countries. These data, available on an annual basis, refer to the actual amount of greenhouse gas emissions produced by a company or entity, as reported in its emissions report and verified by an accredited verifier by 31 March of the following year.³ Following Bjørnland et al. (2023) and Känzig (2023), we temporally disaggregate annual data with the Chow and Lin (1971) method. Annual and interpolated monthly verified emissions are shown in the right panel of Figure 2.⁴

²We do not include the aviation sector because it is covered by the EU ETS since 2012, while we interpolate data on verified emissions using IP indices starting from 2006.

³Based on verified emissions, companies must surrender a corresponding number of emission allowances by the end of April of that year. If a company's verified emissions exceed the number of allowances it holds, it may need to purchase additional allowances from the market to cover the excess emissions.

⁴More precisely, we consider the IP indices and emissions for all sectors covered by the EU ETS (except

Figure 2: Real EU ETS price (left), actual and interpolated verified emissions for EU-19 countries (right) from March 2005 to September 2023.



Notes: Interpolated emissions are obtained using IP indices of six sectors covered by the EU ETS. Shaded areas represent recessions in the Euro Area, as determined by the CEPR-EABCN (<https://eabcn.org>).

The forecasting exercise is based on data spanning from June 2012 to September 2023, comprising a total of 136 monthly observations. The start date of the sample is dictated by the availability of data on auctions, which are particularly relevant for explaining the allocation mechanism of emission allowances during Phases III and IV of the EU ETS (see Figure 1). Forecasts are generated using an expanding window approach: each time a new forecast is produced, the estimation sample is updated by adding a new observation. The first estimation sample ends in December 2017, and the last forecast is issued in September 2022. We consider forecast horizons of one month up to one year ahead. The forecast evaluation sample is the same for all forecast horizons and consists of 58 observations spanning from December 2018 to September 2023.

We denote the level of the real price of carbon in month t as R_t and the log price as $r_t = \log R_t$. Models are estimated using the first difference of the logarithm of the real price of carbon, Δr_t , and forecasts are constructed iteratively from the estimated models and converted into levels as follows:

$$\hat{R}_{t+h|t} = \exp \left(r_t + \sum_{\ell=1}^h \Delta \hat{r}_{t+\ell|t} \right),$$

aviation) and construct an emission-weighted IP index to interpolate annual verified emissions. The MATLAB library provided by Quilis (2013) is used to implement the Chow-Lin method. Since verified emissions for 2023 were not available, we used the 2022 value to interpolate data for 2023.

where $\Delta\hat{r}_{t+\ell|t}$ is the ℓ -step ahead forecasted value. The evaluation of point forecasts relies on the relative Root Mean Squared Forecast Error (RMSFE) that represents the ratio of the RMSFE of a model to the RMSFE of the benchmark, such as the Random Walk (RW). Therefore, a relative RMSFE lower than unity is taken as evidence that a certain model is more accurate than the benchmark.

Sign, or direction-of-change forecasts, are defined as: $\text{sign}\left(\hat{R}_{t+h|t} - R_t\right)$, where $\text{sign}(x)$ equals -1 if $x < 0$, 0 if $x = 0$, and 1 if $x > 0$. We use the Success Ratio (SR), defined as the proportion of correctly predicted signs, to gauge the directional accuracy. A SR greater than 0.5 indicates a gain in accuracy relative to the RW model that implies a no-change forecast.

As a further measure of forecasting, we rely on the quantile-based continuous ranked probability score (qCRPS) of Gneiting and Ranjan (2011), which is a density forecasting measure denoted by

$$\widehat{QS}_t = \frac{1}{J-1} \sum_{j=1}^{J-1} \widehat{QS}_t^{\alpha_j} = \frac{1}{J-1} \sum_{j=1}^{J-1} 2 \left[\mathbb{I}\left(R_{t+h} \leq \hat{q}_{t+h|t}^{\alpha_j}\right) - \alpha_j \right] \times \left(\hat{q}_{t+h|t}^{\alpha_j} - R_{t+h} \right), \quad (1)$$

where $\mathbb{I}(\cdot)$ denotes the indicator function and $\hat{q}_{t+h|t}^{\alpha_j}$ is the h -step ahead quantile forecast for R_{t+h} at level $\alpha_j = j/J$ with $J = 20$, which corresponds to $\alpha_j = 0.05, 0.10, \dots, 0.95$. We construct weighted versions of the qCRPS, where the weights are selected to emphasize specific regions, such as the center or one of the tails of the distribution:

$$\widehat{wQS}_t = \frac{1}{J-1} \sum_{j=1}^{J-1} \nu(\alpha_j) \widehat{QS}_t^{\alpha_j}, \quad \text{with } \nu(\alpha_j) = \begin{cases} \alpha_j(1 - \alpha_j) & \text{(center),} \\ \alpha_j^2 & \text{(right tail),} \\ (1 - \alpha_j)^2 & \text{(left tail).} \end{cases} \quad (2)$$

In a pairwise comparison, the model with the lowest score is ranked as the most accurate.

4 Results

4.1 Univariate time series models

Most previous papers rely on daily or weekly nominal data, therefore in the case of monthly real carbon prices, it is not clear which benchmark to use. In the literature concerning the forecasting of commodity prices, particularly for crude oil, the simple Random Walk (RW) model, or no-change forecast, is commonly used as a benchmark. In this initial phase of the analysis, we focus on point and sign forecasts, while we will consider density forecasts at a later stage. Results are summarized in Table 1, which shows the relative RMFSE and the SR of different models for selected forecast horizons of $h = 1, 2, 3, 6, 9, 12$ months ahead.⁵

In the short run, including a drift in the RW model (RWD) provides modest gains in point forecast accuracy, while for forecasts ranging from one-quarter up to one year ahead, accuracy is comparable to that of the RW model. Coupled with an SR below 0.5 at the 12-month horizon, this implies that the RWD model does not qualify as a more accurate benchmark for point and sign forecasting. If Autoregressive (AR) and/or Moving Average (MA) components are included, models perform poorly at horizons 1 and 2, but they lead to RMSFE reductions at horizons from one quarter up to one year that can be as large as 16.98%. See columns 3 and 4 of Table 1. Looking at the success ratio (panel *b*), ARIMA models have some directional accuracy, especially at longer horizons.

We conclude with Autoregressive, $AR(p)$, models specified for the log first difference estimated with Bayesian natural conjugate prior and denoted as BAR (columns 5-8). We consider fixed lag orders $p = 1, 3$, and 12, as well as lag order selection based on the Akaike Information Criterion (AIC). In this case, we set the maximum lag order to 12 and select p each time a new forecast is issued. Indeed, BAR models for Δr_t display directional accuracy and lead to RMSFE reductions that reach 17% at horizon 12 in the case of a simple BAR(1). Moreover, it is worth pointing out that fixed and small lag orders are preferable to either setting $p = 12$ or selecting p recursively with the AIC. To sum up, given that at horizons 1 and 2, not even BAR models for Δr_t outperform the RW, we cannot definitively discard it

⁵In the Supplement, we display results for all forecast horizons from 1 to 12 months ahead and for additional univariate time series models.

Table 1: Relative RMSFE and Success Ratio of univariate time series models

(a) Relative RMSFE							
h				First-difference, Δr_t			
	RWD	ARIMA(1,1,1)	ARIMA(0,1,1)	BAR(1)	BAR(3)	BAR(12)	BAR(aic)
1	<u>0.993</u>	1.069	1.071	1.043	1.050	1.046	1.047
2	<u>0.999</u>	1.010	1.002	1.003	1.003	1.013	1.004
3	1.0000	0.998	0.993	0.990	0.988	0.998	0.992
6	1.0000	0.905	<u>0.901</u>	0.903	0.903	0.923	0.911
9	1.0000	0.856	<u>0.854</u>	0.855	0.858	0.864	0.858
12	1.0000	0.830	0.831	<u>0.830</u>	0.834	0.848	0.834

(b) Success Ratio							
h				First-difference, Δr_t			
	RWD	ARIMA(1,1,1)	ARIMA(0,1,1)	BAR(1)	BAR(3)	BAR(12)	BAR(aic)
1	<u>0.603</u>	0.535	0.535	0.517	0.535	0.517	0.517
2	<u>0.621</u>	0.517	0.517	0.517	0.517	<u>0.621</u>	0.517
3	<u>0.638</u>	0.535	0.569	0.603	0.621	<u>0.638</u>	0.603
6	<u>0.776</u>	0.759	0.759	<u>0.776</u>	<u>0.776</u>	<u>0.776</u>	<u>0.776</u>
9	0.569	<u>0.828</u>	<u>0.828</u>	<u>0.828</u>	<u>0.828</u>	<u>0.828</u>	<u>0.828</u>
12	0.431	<u>0.897</u>	<u>0.897</u>	<u>0.897</u>	<u>0.897</u>	<u>0.897</u>	<u>0.897</u>

Notes: Panel (a) shows the ratio of RMSFE of model m over the RW model. Values below one suggest superior forecast performance of model m to the RW (in bold). * denotes that the null hypothesis of the Diebold-Mariano test is rejected at the 90% (95%) confidence level. Panel (b) reports the success ratios; entries in bold suggest that the model can accurately predict the direction of change over 50% of the time. * indicates that the p-value of the Pesaran and Timmermann (2009) test of the null of no directional accuracy is below 0.1, hence providing evidence of statistical significance at the 10% significance level. Models yielding the lowest RMSFE or highest SR are underlined.

as a benchmark.

In Table 1, we have used the Diebold and Mariano (1995) test, as modified by Coroneo and Iacone (2020), to verify the statistical significance of RMSFE reductions, and the Pesaran and Timmermann (2009) test to assess the statistical significance of SR. Possibly due to the small size of the evaluation sample or because of instabilities in forecasting performance, we are never able to provide evidence of statistically significant improvements over the RW. We address the presence of forecast instabilities in Sections 4.4 and 4.5.

4.2 VAR models of the EU ETS carbon market

To link the real price of carbon to its determinants, a natural starting point is a small-scale VAR(p) model of the EU ETS market:

$$\mathbf{y}_t = \mathbf{a} + \sum_{j=1}^p \mathbf{A}_j \mathbf{y}_{t-j} + \mathbf{u}_t, \quad (3)$$

where \mathbf{a} is a $n \times 1$ vector of intercepts, \mathbf{A}_j are $n \times n$ matrices of coefficients for $j = 1, \dots, p$, and \mathbf{u}_t is an $n \times 1$ vector of zero-mean innovations Normally distributed with covariance matrix Σ . Following Bjørnland et al. (2023), we consider a baseline VAR specification in which \mathbf{y}_t is a 3×1 vector including Δr_t (or alternatively, r_t), the first difference of the logarithm of interpolated verified emissions, $emis_t$, and first difference of the logarithm of aggregate industrial production for EU-19 countries, Δip_t .

The baseline model does not include many predictors tracked by practitioners, such as energy prices, technical indicators related to the auctioning of EU ETS allowances, and weather anomalies that might affect the demand for electricity and gas. Differently from Bjørnland et al. (2023), we consider also the Euro Stoxx stock price index as an additional proxy of real economic activity, as well as a wider set of IP indices. In particular, we add IP indices for the main sectors covered by the EU ETS to better capture demand-side pressures affecting the real price of carbon.


Given the relatively small size of our sample of data, including the set of 21 predictors described in Section 3 in the model would not be feasible. Similarly, we want to avoid sparse model representations implied by the selection of only a handful of the potentially relevant predictors. Thus, we consider a Factor-augmented VAR model (FAVAR), where we estimate a modified version of the baseline VAR specification, in which we replace the aggregate IP index for EU-19 with up to three factors extracted from the set of 21 predictors.

In detail, we pool the information from the 21 predictors based on principal component analysis.⁶ We decided to focus on the first three factors since they account for 48% of the variance of the 21-time series. The percentages of total variance explained by the first, second, and third factors are 22%, 17.8%, and 8.6%, respectively.

Figure 3 presents the R^2 from regressing a factor on an individual predictor using data from June 2012 to December 2017. Broadly speaking, the first factor loads mostly on energy prices. The second-factor loads on a combination of a few IP indices and the cover ratio,

⁶Before the analysis, all variables have been transformed to induce stationarity and then standardized. IP indices, the Euro Stoxx price index, energy prices, and the auction price are transformed into monthly growth rates (first differences of logarithms). Clean dark and clean spark spreads are first-differenced, while the monthly auction price range is log-transformed. We do not apply any transformation to the auction cover ratio, temperature, and precipitation anomalies. A preliminary screening for outliers was also carried out. However, given that only a few extreme observations (i.e. observations exceeding 20 times the interquartile range from the median) were detected during the COVID pandemic, we decided to keep them in the sample.

Figure 3: R^2 between factors and individual predictors obtained from regressing a factor on an individual predictor using data from June 2012 to December 2017.



while the third loads heavily on the two spreads, precipitation anomalies, and the remaining IP indices.

As shown in the coming sections, the forecasting exercise will show that a BVAR(1) model augmented with one factor predicts both the real price of carbon and verified emissions adequately and better than alternative models. Therefore, it is useful to characterize the variables underlying the dynamics of the first factor. Figure 4 indicates that the first factor tracks the European business cycle as well as the evolution of energy prices and is highly correlated with both IP growth and returns of TTF natural gas. Indeed, during the COVID-19 recession, the first factor is negative and well below its mean, while it becomes positive in the recovery phase and for most of the first semester of 2022 when natural gas prices rise due to the Russian invasion of Ukraine.

Figure 5 displays the evolution of the first factor over time, along with bars representing


Figure 4: Correlation between Factor 1 and industrial production (top) and growth rate of the price of TTF natural gas (bottom) from December 2017 to September 2023.



Notes: we depict centered 3-month moving average of Factor 1 and predictors (scaled to have same variance as Factor 1).

the contribution of different predictors grouped by class (i.e. economic activity, energy prices, technical indicators, and weather anomalies). Interestingly, the drivers underlying the movements of the first factor change over time, but are always mostly related to economic activity proxies and energy prices. Measures of real economic activity account for the largest share of the downward movement of the factor in 2020 when European economies were frozen during lockdowns. The upward movement of the factor in 2022 and its successive decline in 2023 are attributed to the pressures related to the war in Ukraine and the subsequent easing of conditions in energy markets.

Figure 5: Contribution of predictors (grouped by class) to Factor 1, represented as bars, along with the evolution of Factor 1 (green line) from December 2017 to September 2023.



Notes: we depict the centered 3-month moving average of both Factor 1 and the predictors' contributions to it.

4.3 Can VAR models forecast the real price of carbon?

In this section, we focus on Bayesian techniques (BVAR or BFAVAR) with Minnesota prior to forecast the real price of carbon. Overall, both for baseline BVAR models and factor-augmented specifications, setting the lag order to 1 or 3 yields the most accurate point and direction-of-change forecasts. Indeed, inspecting the leftmost column of Table 2, the baseline BVAR model with $p = 12$ – as set by Bjørnland et al. (2023) – is regularly outperformed by specifications with $p = 1$ and $p = 3$.⁷

Table 2 shows that at horizons up to 2 months, BVAR and BFAVAR models are as accurate as no-change point forecasts. However, as the forecast horizon grows, BVAR and BFAVAR models lead to sizable accuracy gains both for point and sign forecasts. For instance, at a one-year horizon, the BVAR(1) and factor-augmented BVAR(1) model result in a reduction in RMSFE in the range 17.8-18.2% and can accurately predict the direction of

⁷In the Supplement, we have estimated the same models with the real price of carbon log-levels, and point forecasts were always less accurate than growth rates estimation. As for the lag length of BVAR models, we consider pre-specified values of $p = 1, 3, 12$ along with a data-driven approach that sets p by minimizing the AIC at each forecast origin, selecting among BVAR(p) models up to order 12.

price movements 89.7% of the time.

Table 2: Relative RMSFE and Success Ratio for Bayesian Vector Autoregressive and Factor models.

(a) Relative RMSFE									
h	BVAR(1)	BVAR(3)	BVAR(12)	1 Factor		2 Factors		3 Factors	
				BVAR(1)	BVAR(3)	BVAR(1)	BVAR(3)	BVAR(1)	BVAR(3)
1	1.057	1.075	1.074	1.062	1.073	1.039	1.077	1.062	1.114
2	1.006	1.019	1.028	1.010	1.028	1.024	1.102	1.066	1.154
3	<u>0.995</u>	0.996	1.002	1.002	1.019	0.999	1.052	1.015	1.062
4	0.953	0.952	0.970	<u>0.947</u>	0.953	0.963	0.960	0.976	0.975
5	0.965	0.964	0.984	<u>0.957</u>	0.960	0.970	0.971	0.985	0.983
6	0.915	0.916	0.937	<u>0.906</u>	0.916	0.921	0.919	0.935	0.927
7	0.897	0.899	0.917	<u>0.886</u>	0.899	0.902	0.894	0.916	0.899
8	0.844	0.852	0.858	<u>0.831</u>	0.847	0.850	0.856	0.867	0.863
9	0.857	0.865	0.873	<u>0.851</u>	0.872	0.856	0.901	0.863	0.894
10	0.825	0.836	0.838	<u>0.824</u>	0.850	0.830	0.887	0.833	0.881
11	<u>0.822</u>	0.837	0.843	0.824	0.849	0.828	0.886	0.831	0.883
12	<u>0.818</u>	0.834	0.848	0.822	0.849	0.828	0.892	0.834	0.891

(b) Success Ratio									
h	BVAR(1)	BVAR(3)	BVAR(12)	1 Factor		2 Factors		3 Factors	
				BVAR(1)	BVAR(3)	BVAR(1)	BVAR(3)	BVAR(1)	BVAR(3)
1	0.552	<u>0.569</u>	0.517	0.552	0.500	<u>0.569</u>	<u>0.569</u>	0.552	<u>0.569</u>
2	0.535	0.552	<u>0.569</u>	0.535	<u>0.569</u>	0.552	<u>0.569</u>	0.552	0.535
3	0.603	0.621	0.621	0.621	0.603	0.621	<u>0.655</u>*	0.603	0.603
4	<u>0.707</u>	0.672	0.672	0.690	0.655	<u>0.707</u>	0.690	0.690	0.672
5	<u>0.672</u>	<u>0.672</u>	0.655	<u>0.672</u>	0.638	<u>0.672</u>	0.655	<u>0.672</u>	<u>0.672</u>
6	<u>0.776</u>	0.759	0.759	<u>0.776</u>	0.741	<u>0.776</u>	0.724	<u>0.776</u>	0.759
7	<u>0.793</u>	0.776	<u>0.793</u>	<u>0.793</u>	0.759	<u>0.793</u>	0.759	<u>0.793</u>	0.776
8	<u>0.845</u>	0.828	<u>0.845</u>	<u>0.845</u>	0.810	<u>0.845</u>	0.810	<u>0.845</u>	0.828
9	<u>0.828</u>	0.810	<u>0.828</u>	<u>0.828</u>	0.793	<u>0.828</u>	0.793	<u>0.828</u>	0.810
10	<u>0.862</u>	0.828	<u>0.862</u>	<u>0.862</u>	0.828	<u>0.862</u>	0.810	<u>0.862</u>	0.828
11	<u>0.914</u>	0.879	<u>0.914</u>	<u>0.914</u>	0.879	<u>0.914</u>	0.862	<u>0.914</u>	0.874
12	<u>0.897</u>	0.879	<u>0.897</u>	<u>0.897</u>	0.862	<u>0.897</u>	0.845	<u>0.897</u>	0.862

Notes: see notes to Table 1

At intermediate forecast horizons (i.e., 4 up to 10 months ahead), the BVAR(1) model augmented with a single factor regularly yields the most accurate point forecasts, and its performance tends to improve as the horizon grows. Increasing the number of factors above one generally leads to less accurate point forecasts. In particular, adding the third factor does not seem to provide any advantage to models with one or two factors. The model with two factors, on the other hand, seems to be slightly better than the single-factor specification at shorter forecast horizons.

4.4 Are there instabilities in forecasting performance?

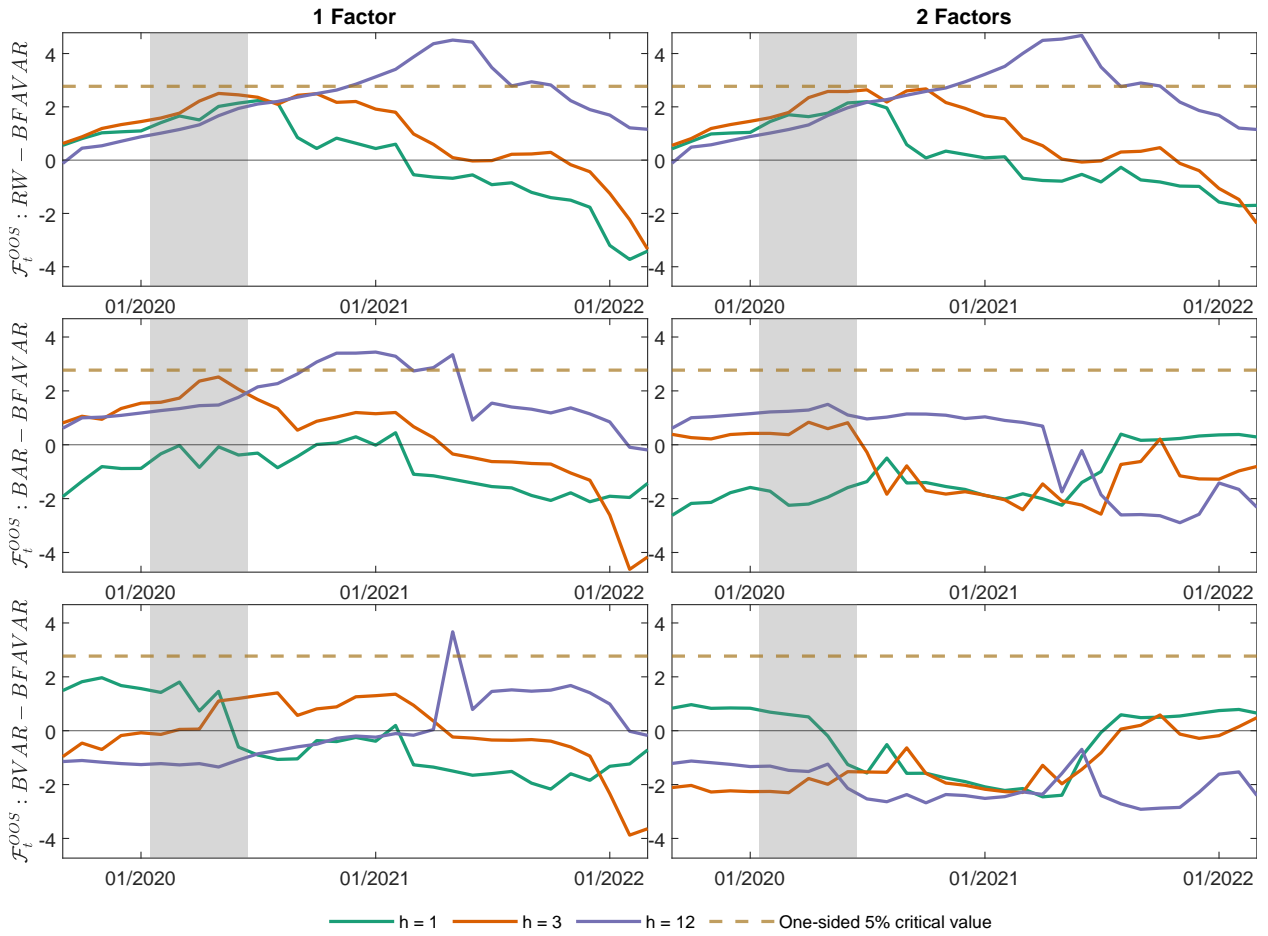
There is widespread evidence that the relative forecasting performance of models changes over time due to parameter instability, shocks with time-varying volatilities, and changes in the variance of the predictors (Giacomini and Rossi, 2010; Rossi, 2021). In such cases, averaging the difference in forecasting performance over the full evaluation sample, as we did in the previous sections, results in a loss of information that might lead to standard tests of predictive ability to conclude that two competing models are equally accurate.

In Figure 6, we investigate the existence of instabilities in forecasting performance. We do so by implementing the fluctuation test statistic, \mathcal{F}_t^{OOS} , of Giacomini and Rossi (2010) considering a centered moving average over a 19-month window.⁸ The test is based on the (standardized) difference between the MSFE of a benchmark model and of the factor-augmented BVAR(1) model with one or two factors. Positive values of the test statistic indicate that the BFAVAR model has a lower MSFE than the benchmark. All tests are one-sided: the null hypothesis is the factor VAR model has the same MSFE as the benchmark, while the alternative is that the former is more accurate than the latter. The dashed line indicates the 5% critical value, $CV_{0.05}$, and the null hypothesis is rejected when $\max \mathcal{F}_t^{OOS} > CV_{0.05}$.

In our comparison, we focus on one month, one quarter, and one year ahead forecast horizons. To raise the bar of forecast evaluation, we do not solely focus on the RW model. Instead, we compare the performance of factor models against the BAR(1) and the BVAR(1) models, which, as shown in previous sections, appear to be more competitive benchmarks than the RW model.

⁸The test has a nonstandard distribution, and the critical values provided by Giacomini and Rossi (2010) depend on the ratio between the size of the window used to compute the moving average and the number of observations in the evaluation sample. In the Supplement, we show that the results are robust to changes in the size of the window.

Figure 6: Fluctuation test statistic for a BFAVAR with 1 factor (left) and 2 factors (right) against different benchmarks for forecast horizons $h = 1, 3, 12$ months from September 2019 to December 2022.



Notes: the fluctuation test statistic, \mathcal{F}_t^{OOS} , of Giacomini and Rossi (2010), is calculated with a 19-month centered rolling window. Positive values indicate that the BFAVAR model is better than the benchmark. All tests are one-sided, with the null hypothesis being that the BFAVAR(1) model has the same MSFE as the benchmark; the alternative is that the BFAVAR(1) model forecasts better than the benchmark. The dashed line indicates the one-sided 5% critical value, $CV_{0.05}$. The null hypothesis is rejected when $\max \mathcal{F}_t^{OOS} > CV_{0.05}$.

Several interesting results emerge from Figure 6. First, regardless of the benchmark model, the relative forecasting performance of BVAR specifications – augmented with either one or two factors – changes over time and tends to deteriorate at the end of 2022. Indeed, especially at shorter forecast horizons, the test statistic often becomes negative in this period. Second, the model with one factor (left panel) is usually preferable to the model with two factors (right panel), in that the latter rarely beats the benchmarks. Third, at a forecast horizon of one month, the single-factor model performs better than the RW model until the beginning of 2021; while it never beats the BAR(1). However, at longer forecast horizons and especially for $h = 12$, the single-factor model performs better than the benchmarks

during most of the evaluation sample. Moreover, for $h = 12$, the fluctuation test leads to a rejection of the null hypothesis of indistinguishable forecasting performance; in all cases, the test statistic lies above the 5% critical value in late 2020 and in part of 2021.

Thus the single-factor BVAR(1) model, condensing information of a broad set of predictors related to the EU ETS, is a promising specification for forecasting the magnitude as well as the direction-of-change of the real price of carbon at longer horizons. At horizons shorter than a quarter, VAR models do not offer a real advantage over simple univariate specifications, especially in point forecasting.

4.5 Density forecasts and the role of stochastic volatility

To gauge the uncertainty associated with point forecasts and following evidence that modeling SV improves density forecasts of macroeconomic aggregates (Clark and Ravazzolo, 2015; Chan, 2023), we evaluate density forecasting for future values of the real price of carbon.

We also decided to include in BVAR models a Choleski multivariate SV process (Chan, 2023, for details):

$$\mathbf{u}_t \sim N(\mathbf{0}, \boldsymbol{\Sigma}_t), \quad \boldsymbol{\Sigma}_t^{-1} = \mathbf{B}'_0 \mathbf{D}_t^{-1} \mathbf{B}_0, \quad (4)$$

$$h_{i,t} = \mu_i + \phi_i (h_{i,t-1} - \mu_i) + \varepsilon_{i,t}, \quad \varepsilon_{i,t} \sim N(0, \sigma_\varepsilon^2), \quad (5)$$

where \mathbf{B}_0 is a lower triangular matrix of size $n \times n$ with ones along the main diagonal and $\mathbf{D}_t = \text{diag}(e^{h_{1,t}}, \dots, e^{h_{n,t}})$, and Equation (5) specifies an independent autoregressive volatility process for each variable in the model for $t = 2, \dots, T$ with initial condition $h_{i,1} \sim N(\mu_i, \sigma_i^2 / (1 - \phi_i^2))$.

To evaluate the density performance, we rely on qCRPS (Gneiting and Ranjan, 2011), where the model with the lowest qCRPS is ranked as the most accurate.⁹ In Panel (a) of Table 3, we focus on the center of the distribution, while the accuracy in forecasting the right and left tails of the distribution, which are of interest to assess the probability of extreme price movements, is evaluated in Panel (b) and (c), respectively.

⁹Note that, contrary to the previous tables, these numbers are not ratios to a benchmark but are expressed in the same scale as real prices.

Table 3: Quantile-weighted Continuous Ranked Probability Score (qCRPS).

(a) Quantile-weighted CRPS: center

h	BVAR(1)	1 Factor	2 Factors	BVAR-SV(1)	1 Factor	2 Factors
		BVAR(1)	BVAR(1)		BVAR-SV(1)	BVAR-SV(1)
1	0.729	0.721	<u>0.692</u>	0.712	0.714	0.704
2	0.978	<u>0.966</u>	0.968	0.977	0.974	0.983
3	1.231	1.226	<u>1.200</u>	1.257	1.252	1.253
4	1.322	<u>1.310</u>	1.318	1.349	1.341	1.355
5	1.625	<u>1.611</u>	1.623	1.648	1.639	1.659
6	1.670	<u>1.649</u>	1.663	1.701	1.689	1.712
7	1.790	<u>1.761</u>	1.782	1.804	1.793	1.819
8	1.860	<u>1.831</u>	1.861	1.886	1.877	1.905
9	2.121	<u>2.102</u>	2.109	2.144	2.139	2.161
10	2.180	<u>2.169</u>	2.185	2.245	2.248	2.277
11	2.325	<u>2.325</u>	2.336	2.407	2.404	2.432
12	<u>2.454</u>	2.472	2.480	2.587	2.585	2.612

(b) Quantile-weighted CRPS: right tail


h	BVAR(1)	1 Factor	2 Factors	BVAR-SV(1)	1 Factor	2 Factors
		BVAR(1)	BVAR(1)		BVAR-SV(1)	BVAR-SV(1)
1	1.286	1.217	<u>1.124</u>	1.237	1.220	1.198
2	1.733	1.676	<u>1.663</u>	1.681	1.664	1.678
3	2.207	2.131	<u>2.088</u>	2.166	2.139	2.170
4	2.392	<u>2.315</u>	2.327	2.358	2.332	2.389
5	2.915	2.835	2.847	2.845	<u>2.815</u>	2.880
6	3.111	3.042	3.064	3.036	<u>3.015</u>	3.084
7	3.324	3.219	3.268	3.235	<u>3.208</u>	3.299
8	3.515	<u>3.394</u>	3.452	3.462	3.437	3.544
9	3.984	<u>3.890</u>	3.917	3.914	3.900	4.008
10	4.304	<u>4.230</u>	4.264	4.302	4.293	4.417
11	4.630	<u>4.576</u>	4.596	4.638	4.634	4.757
12	5.015	<u>4.9996</u>	5.008	5.090	5.098	5.222

(b) Quantile-weighted CRPS: left tail

h	BVAR(1)	1 Factor	2 Factors	BVAR-SV(1)	1 Factor	2 Factors
		BVAR(1)	BVAR(1)		BVAR-SV(1)	BVAR-SV(1)
1	1.270	1.282	<u>1.235</u>	1.278	1.288	1.268
2	1.599	1.605	<u>1.588</u>	1.677	1.671	1.680
3	1.897	1.914	<u>1.834</u>	2.046	2.026	2.001
4	<u>1.943</u>	1.958	1.953	2.116	2.091	2.097
5	<u>2.404</u>	2.412	2.410	2.606	2.572	2.598
6	2.303	<u>2.292</u>	2.306	2.506	2.464	2.491
7	2.471	<u>2.467</u>	2.471	2.642	2.609	2.629
8	2.534	<u>2.506</u>	2.562	2.689	2.656	2.686
9	2.901	2.888	<u>2.863</u>	3.082	3.052	3.048
10	2.806	<u>2.803</u>	2.811	3.021	2.998	3.001
11	<u>2.946</u>	2.960	2.963	3.179	3.155	3.165
12	<u>3.007</u>	3.029	3.053	3.289	3.266	3.282

Notes: The best forecasts, associated with the lowest scores, are underlined.

Figure 7: Fluctuation test statistic for the right tail for BFAVAR with 1 factor (left) and 2 factors (right). September 2019 to December 2022.



With very few exceptions, homoskedastic factor-augmented BVAR models – with either one or two factors – are more accurate at forecasting the center and both tails than the alternative specifications. In the case of the real price of carbon, incorporating SV does not seem to offer any sizable advantage on any horizon. Interestingly, the homoskedastic BVAR(1) model augmented with two factors is more accurate than any other specifications at horizons up to a quarter-ahead. This result stands in contrast to what is observed for point and sign forecasts, for which factor-augmented models do not yield accuracy gains at shorter horizons. Moreover, it emphasizes that, for short-term density forecasts, the second factor, which captures predictors beyond energy prices and economic activity that are the main drivers of the first factor mostly, might be relevant.¹⁰

We complement these results with evidence of instabilities in the (right) tail forecasting. When evaluating the single-factor model, against the BAR(1) and BVAR(1) models, results in the first column of Figure 7 allow rejecting the null of equal accuracy in right tail fore-

¹⁰See the Supplement for results concerning the inclusion of SV in BVAR models.

casting at all forecast horizons. Therefore, the relative forecasting performance is not stable, but in several months, the single-factor model offers accuracy gains. As for the specification with two factors, we can see that the null is rejected at all horizons when the benchmark is the BAR(1) model and at horizons 1 and 12 when the benchmark is the BVAR(1) model. Results concerning the center of the distribution largely mimic what is observed for point forecasts: the single-factor model does better than the benchmarks at horizons of one quarter and one year.¹¹

5 Expert forecasts, verified emissions and market monitoring

5.1 Expert forecasts of the nominal price of carbon


As a further step, we provide a qualitative comparison of the single-factor BVAR(1) point forecasts against those issued by the Carbon Team at the London Stock Exchange Group (LSEG; formerly Refinitiv), and its survey forecasts. In both cases, LSEG provides forecasts expressed in current euros; thus, we need to transform our real price forecasts into nominal terms by producing inflation forecasts from an Unobserved Component SV model and using them to get nominal price forecasts.¹²

LSEG Carbon Team’s forecasts. The Carbon Team at LSEG produces forecasts of nominal EU ETS price with an irregular cadence for the period 2014-2023 ranging from 3 to 6 times per year, where the forecast horizon can be either the current year or several years in the future. One challenge of working with these data is that they are “fixed event forecasts”, while our models produce “fixed horizon forecasts”. The characteristic of “fixed event forecasts” is that the forecast horizon changes as the forecast origin moves forward. At each forecast origin, the LSEG team produces forecasts for the current year, $f_{t+k|t}^{FE}$, and for the next year, $f_{t+k+12|t}^{FE}$ where $k = 1, \dots, 12$ represents the number of months until the end of the year (e.g.

¹¹In the Supplement, we consider the center and the left tail as well as robustness checks involving changes in the implementation of the fluctuation test. We have considered if dropping verified emissions from BVAR and BFAVAR models could alter our main conclusions, but the relative ranking of models is not affected.

¹²Given the short time span of our evaluation sample, using different inflation forecasts, such as RW forecasts, does not alter the results.

Figure 8: Nominal EU ETS price and forecasts.



Notes: Nominal EU ETS price (red line), LSEG (blue dots), RW (dashed black line), BFAVAR(1) one-year-ahead forecasts (dotted blue line). The bars represent the difference between the RMSFE of the LSEG forecast and RW (black) or BFAVAR(1) (blue). Positive values indicate that LSEG is less accurate than the alternative forecast.

$k = 12$ in January and $k = 1$ in December). To approximate one year ahead fixed horizon forecasts, $f_{t+12|t}^{FH}$, using LSEG's fixed event forecasts, we follow Dovert et al. (2012):


$$f_{t+12|t}^{FH} = \frac{k}{12} f_{t+k|t}^{FE} + \frac{12-k}{12} f_{t+k+12|t}^{FE}, \quad (6)$$

where weights are proportional to the degree of overlap of the two fixed event forecasts.¹³ We obtain a set of 39 one-year-ahead forecasts irregularly spaced over the period January 2015 - May 2023; only 19 of these forecasts overlap with those in our evaluation period spanning from December 2018 to September 2023.

Figure 8 shows that LSEG and RW forecasts are remarkably similar; indeed, the correlation of the respective forecast errors is 0.95 over the period January 2015 - May 2023. For the period when the single-factor BVAR(1) model forecasts overlap with those from LSEG, the model-based RMSFE is smaller than LSEG's RMSFE in 15 cases out of 19. All in all,

¹³A fixed horizon forecast issued in January, $k = 12$, would therefore be equal to $f_{t+12|t}^{FH} = f_{t+12|t}^{FE}$.

Figure 9: Survey and model-based one-year-ahead density forecasts of the nominal EU ETS price from 2021 to 2023.



Notes: each year the survey provides a different set of price ranges across which respondents can choose. The title indicates the date of the forecast, while the data come from the survey of the previous year. The 2020 survey – leftmost plot – provided only three categories “< 25”, “About 25” and “> 25” Euro per tCO₂e.

while the short sample of data only allows for a qualitative comparison, these results show that the model-based forecasts are strikingly different from those issued by LSEG, which, on the other hand, are similar to those from a RW model.

Survey forecasts. Carbon Market Surveys run by LSEG each year (from 2020 to 2022) between February and April (in 2020 it was closed in March) capture the market sentiment of respondents from a multitude of countries who are mostly stakeholders with tangible and financial interests in carbon markets (e.g., traders, firms covered by an ETS). Participants who have answered the section of the survey concerning EU ETS price expectations are 60 in 2020, 119 in 2021, and 88 in 2022. Since also survey forecasts are of the fixed event type, we approximate fixed horizon one-year-ahead prediction densities derived from surveys and

compare them with prediction densities from the BFAVAR model and the realized price at the target date (see Figure 9).

Only in 2023 survey and model-based forecasts are aligned and centered around the realized price. In 2021 and 2022, both the survey and the single-factor BVAR model tend to under-forecast and hence largely miss the run-up in the European carbon price over this period. Nevertheless, we notice that the realized prices lie in the right tail of the one-year-ahead prediction density of the BFAVAR model. The fact that forecasting the price for March 2021 a year in advance was a hard task appears evident, noticing that the price lies outside the 95% BFAVAR prediction interval only in this case. The 2021 forecast is based on survey data collected in March 2020, at the beginning of the COVID-19 pandemic, when overall macroeconomic uncertainty was at a record level, especially in European countries. Similarly, also BFAVAR's one-year-ahead forecasts for 2021 and 2022 are based on macroeconomic data for the COVID-19 recession period, without making any adjustments to them. These factors should be kept in mind when evaluating the prediction accuracy for 2021 and 2022 in absolute terms.

5.2 Verified emissions forecasts

Forecasts of verified emissions are as crucial as price forecasts, and they have received essentially no attention in the academic literature. Entities participating in the EU ETS need reliable emission forecasts to plan emission reduction strategies and ensure compliance with regulatory requirements, while environmental agencies rely on these forecasts to assess the impact of emissions reduction initiatives. BVAR models considered in this paper might represent a valuable tool to produce verified emission forecasts.

Table 4 shows that, as far as point forecasts are concerned, the baseline BVAR(1)-SV model represents the most accurate model at all horizons above 1 month. Interestingly, since data on verified emissions are published once a year in April, the one-year-ahead forecast horizon appears particularly relevant. Indeed, the baseline BVAR(1)-SV model and the factor-augmented BFAVAR(1)-SV models yield RMSFE reductions over 7%. Moving to density forecasting, Table 5 highlights that the BVAR(1)-SV and the single-factor

BFAVAR(1)-SV models yield the most accurate forecasts of the center of the distribution of verified emissions.

Table 4: Relative RMSFE for verified emissions.

h	BAR(1)	BVAR(1)	1 Factor	2 Factors	BAR(1)-SV	BVAR(1)-SV	1 Factor	2 Factors
			BVAR(1)	BVAR(1)			BVAR(1)-SV	BVAR(1)-SV
1	1.102	1.241	1.182	1.237	1.005	1.019	1.029	1.018
2	1.055	1.230	1.166	1.204	0.999	0.997	1.002	0.998
3	1.046	1.252	1.174	1.218	0.998	0.986	0.989	0.988
4	1.042	1.268	1.175	1.226	0.997	0.975	0.979	0.978
5	1.033	1.247	1.154	1.204	0.996	0.963	0.965	0.965
6	1.029	1.238	1.141	1.193	0.996	0.956	0.958	0.957
7	1.023	1.209	1.117	1.166	0.995	0.948	0.950	0.950
8	1.021	1.195	1.104	1.153	0.996	0.944	0.946	0.946
9	1.018	1.173	1.086	1.133	0.996	0.939	0.942	0.941
10	1.014	1.149	1.071	1.112	0.996	0.935	0.938	0.938
11	1.014	1.143	1.065	1.106	0.996	0.930	0.934	0.933
12	1.012	1.129	1.053	1.093	0.996	0.926	0.929	0.929

Notes: see notes to Table 1

Table 5: Quantile-weighted Continuous Ranked Probability Score (qCRPS) for verified emissions.

(a) Quantile-weighted CRPS: center

h	BAR(1)	BVAR(1)	1 Factor	2 Factors	BAR(1)-SV	BVAR-SV(1)	1 Factor	2 Factors
			BVAR(1)	BVAR(1)			BVAR-SV(1)	BVAR-SV(1)
1	154.876	154.834	154.813	154.799	154.837	154.819	154.812	154.824
2	154.875	154.813	154.787	154.774	154.791	154.775	154.759	154.781
3	154.890	154.810	154.786	154.771	154.757	154.743	154.725	154.750
4	154.929	154.836	154.814	154.798	154.746	154.734	154.716	154.742
5	154.975	154.868	154.853	154.839	154.743	154.730	154.712	154.741
6	155.076	154.958	154.946	154.931	154.795	154.782	154.765	154.795
7	155.180	155.055	155.046	155.035	154.851	154.839	154.825	154.856
8	155.332	155.200	155.194	155.191	154.953	154.9380	154.928	154.966
9	155.424	155.284	155.282	155.301	154.997	154.983	154.974	155.020
10	155.583	155.442	155.441	155.451	155.109	155.093	155.091	155.134
11	155.698	155.551	155.550	155.568	155.175	155.156	155.160	155.207
12	155.792	155.632	155.633	155.658	155.226	155.201	155.208	155.257

Notes: see notes to Table 3

5.3 Market monitoring

Following Baumeister et al. (2022), we construct indices of *demand pressure*, *upward*, and *downward price pressure* for the EU ETS market. For approximating demand pressure, we take the difference between one-year and one-month-ahead verified emission forecasts from the BFAVAR(1)-SV model. A negative value of the proxy signals expectations of loosening market conditions over the next year.

Using the predictive densities delivered by Bayesian estimation of the single-factor BVAR(1) model, upward and downward price pressure indices are defined as follows:

$$PP_t^+ = \frac{1}{12} \sum_{h=1}^{12} \mathbb{I} \left[\hat{R}_{t+h|t} > \max(R_t, R_{t-1}, \dots, R_{t-11}) \right], \quad (7)$$


$$PP_t^- = \frac{1}{12} \sum_{h=1}^{12} \mathbb{I} \left[\hat{R}_{t+h|t} < \min(R_t, R_{t-1}, \dots, R_{t-11}) \right]. \quad (8)$$

These proxies estimate the probability that over the next 12 months, the real price of carbon is above (below) the maximum (minimum) value observed in the previous year.

The upper panel of Figure 10 displays the demand pressure proxy as a set of bars along with a line representing verified emissions. Note that, although verified emissions are available at an annual sampling frequency, BVAR models are estimated with monthly data; therefore, the demand pressure index allows monitoring expectations about the EU ETS market in real time each month. The index is always negative to track the long-term decline of verified emissions, interspersed with an increase in 2021 in the aftermath of the COVID-19 recession.

As shown by Bjørnland et al. (2023), the dynamics of verified emissions are elicited by three main forces: a supply shock and two distinct demand shocks. On the supply side, we have EU ETS regulation: the cap and trade scheme has progressively tightened the limit of total greenhouse gas emissions of covered entities. This is the first element factored in the negative demand pressure index. The two demand shocks are related to economic activity and transition demand. The first essentially depends on the business cycle and, therefore, can be associated with positive or negative changes in the demand pressure index. For instance, during the COVID pandemic, expectations of a future quick recovery associated

Figure 10: Demand and price pressure indices for the EU ETS market from March 2018 to August 2023.



Notes: we plot the backward 3-month moving average of price and demand pressure indices. Indices are aligned with the forecast origin.

with an increase in industrial activity are captured by a rebound of the demand pressure index. Much like the supply-side shock, the transition-demand shock – capturing, among other things, the increased usage of renewables – contributes to the steady decline in verified emissions and therefore keeps the pressure index in negative territory.

The two price pressure indices, along with the evolution of the real EU ETS prices,

are shown in the middle and bottom panels of Figure 10. Given that between 2018 and 2023, the real price has steadily increased – peaking at almost 120 Euros in early 2023 – the upward pressure index in the middle panel always signals expectations of soaring prices. Two exceptions are recorded in 2020 and in early 2023 when the downward price pressure index rises temporarily, capturing expectations of price decreases.

6 Conclusions

The fact that the EU ETS regulation is strengthening over time calls for even more analyses focusing on the macroeconomic effects of carbon price shocks. Indeed, the ECB already embeds technical assumptions on carbon pricing in its projections (European Central Bank, 2021b). Technical assumptions boil down to setting the trajectory for key variables entering the ECB’s macroeconomic models over the projection horizon and are derived in a variety of ways, including univariate and multivariate econometric models and using the price of futures contracts (European Central Bank, 2016).

In this paper, we have identified carbon price drivers and methodological choices that can directly inform projections and scenario analyses used to gauge the macroeconomic effect of carbon price shocks. Our results show that EU ETS prices and verified emissions can be forecasted with relatively simple BVAR models estimated with monthly data.

There are at least two aspects of our analysis that deserve further investigation. First, the use of time series sampled at different frequencies. Mixed Frequency and Mixed-data sampling (MIDAS) models could be used to further improve monthly forecasts relying on data sampled at a higher frequency, such as weather and financial variables. Moreover, given that several key predictors of real carbon prices – such as the level of verified emissions and some macroeconomic aggregates – are available only at the lower sampling frequency, it is also interesting to rely on reverse-MIDAS approaches to exploit these data (Feroni et al., 2023). Lastly, our results could be extended in the direction of real-time data to assess the impact of data revisions on monthly forecasts of the price of carbon.

Supplementary Materials

An online Supplement provides further forecasting results and different fluctuation tests for various models and rolling windows.

References

- Baumeister, C., Korobilis, D., and Lee, T. K. (2022). Energy markets and global economic conditions. *Review of Economics and Statistics*, 104(4):828–844.
- Bjørnland, H., Cross, J. L., and Kapfhammer, F. (2023). The drivers of emission reductions in the European carbon market. CAMP Working Paper Series 8/2023, BI Norwegian Business School.
- Boivin, J. and Ng, S. (2006). Are more data always better for factor analysis? *Journal of Econometrics*, 132(1):169–194.
- Carleton, T. A. and Hsiang, S. M. (2016). Social and economic impacts of climate. *Science*, 353(6304):aad9837.
- Chan, J. C. (2023). Comparing stochastic volatility specifications for large Bayesian VARs. *Journal of Econometrics*, 235(2):1419–1446.
- Chevallier, J. (2011a). Macroeconomics, finance, commodities: interactions with carbon markets in a data-rich model. *Economic Modelling*, 28(1-2):557–567.
- Chevallier, J. (2011b). A model of carbon price interactions with macroeconomic and energy dynamics. *Energy Economics*, 33(6):1295–1312.
- Chevallier, J. (2012). *Econometric analysis of carbon markets*. Springer Netherlands.
- Chow, G. C. and Lin, A.-L. (1971). Best linear unbiased interpolation, distribution, and extrapolation of time series by related series. *The Review of Economics and Statistics*, 53(4):372–375.

- Ciccarelli, M. and Marotta, F. (2024). Demand or supply? An empirical exploration of the effects of climate change on the macroeconomy. *Energy Economics*, 129:107163.
- Clark, T. E. and Ravazzolo, F. (2015). Macroeconomic forecasting performance under alternative specifications of time-varying volatility. *Journal of Applied Econometrics*, 30(4):551–575.
- Coroneo, L. and Iacone, F. (2020). Comparing predictive accuracy in small samples using fixed-smoothing asymptotics. *Journal of Applied Econometrics*, 35(4):391–409.
- Dell, M., Jones, B. F., and Olken, B. A. (2014). What do we learn from the weather? The new climate-economy literature. *Journal of Economic literature*, 52(3):740–798.
- Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(1):253.
- Dovern, J., Fritsche, U., and Slacalek, J. (2012). Disagreement among forecasters in G7 countries. *Review of Economics and Statistics*, 94(4):1081–1096.
- European Central Bank (2016). *A guide to the Eurosystem/ECB staff macroeconomic projection exercises*. Available at: <https://data.europa.eu/doi/10.2866/35667>.
- European Central Bank (2021a). Climate change and monetary policy in the Euro Area. Occasional Paper Series 271, European Central Bank.
- European Central Bank (2021b). ECB presents action plan to include climate change considerations in its monetary policy strategy. Technical report, Press Release. Available at: <https://www.ecb.europa.eu>.
- Froni, C., Ravazzolo, F., and Rossini, L. (2023). Are low frequency macroeconomic variables important for high frequency electricity prices? *Economic Modelling*, 120:106160.
- Fullerton, D. and Muehlegger, E. (2019). Who bears the economic burdens of environmental regulations? *Review of Environmental Economics and Policy*, 13(1):62–82.
- Giacomini, R. and Rossi, B. (2010). Forecast comparisons in unstable environments. *Journal of Applied Econometrics*, 25(4):595–620.

- Gneiting, T. and Ranjan, R. (2011). Comparing density forecasts using threshold- and quantile-weighted scoring rules. *Journal of Business & Economic Statistics*, 29(3):411–422.
- Hsiang, S., Oliva, P., and Walker, R. (2019). The distribution of environmental damages. *Review of Environmental Economics and Policy*, 13(1):83–103.
- International Carbon Action Partnership (2023). Emissions trading worldwide. *Status Report 2023*.
- Känzig, D. R. (2023). The unequal economic consequences of carbon pricing. Working Paper 31221, National Bureau of Economic Research.
- Känzig, D. R. and Konradt, M. (2023). Climate policy and the economy: evidence from Europe’s carbon pricing initiatives. Working Paper 31260, National Bureau of Economic Research.
- Koop, G. and Tole, L. (2013). Forecasting the European carbon market. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 176(3):723–741.
- Lei, H., Xue, M., and Liu, H. (2022). Probability distribution forecasting of carbon allowance prices: a hybrid model considering multiple influencing factors. *Energy Economics*, 113:106189.
- Marcu, A., Hernández, J. F. L., Romeo, G., Alberola, E., Faure, A., Obienu, C., Qin, B., O’Neill, M., Caneill, J. Y., and Schleicher, S. (2023). 2023 State of the EU ETS Report. Available at: <https://ercst.org/2023-state-of-the-eu-ets-report/>.
- Moessner, R. (2022). Effects of carbon pricing on inflation. Working Papers 9563, CESifo.
- NGFS (2022). NGFS scenarios for central banks and supervisors. NGFS publications, Central Banks and Supervisors Network for Greening the Financial System (NGFS). Available at: <https://www.ngfs.net/en>.
- Pan, J., Cross, J. L., and Zhang, B. (2023). A review of carbon emissions reduction mechanisms. Working paper.

- Parry, I., Black, S., and Roaf, J. (2021). Proposal for an international carbon price floor among large emitters. Staff Climate Notes 2021/001, International Monetary Fund.
- Pesaran, M. H. and Timmermann, A. (2009). Testing dependence among serially correlated multicategory variables. *Journal of the American Statistical Association*, 104(485):325–337.
- Quilis, E. M. (2013). Temporal disaggregation library. MATLAB Central File Exchange.
- Rossi, B. (2021). Forecasting in the presence of instabilities: How we know whether models predict well and how to improve them. *Journal of Economic Literature*, 59(4):1135–1190.
- Tan, X., Sirichand, K., Vivian, A., and Wang, X. (2022). Forecasting European carbon returns using dimension reduction techniques: commodity versus financial fundamentals. *International Journal of Forecasting*, 38(3):944–969.

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 2024

1. 2024, A. Sileo, M. Bonacina, The automotive industry: when regulated supply fails to meet demand. The Case of Italy



Fondazione Eni Enrico Mattei

Corso Magenta 63, Milano - Italia

Tel. +39 02 403 36934

E-mail: letter@feem.it

www.feem.it

