EUROPEAN CENTRAL BANK
EUROSYSTEM

Working Paper Series

Jean-Charles Bricongne, Joao Macalos,
Baptiste Meunier, Julia Milis, Thomas Pical

Can satellites predict oil demand?

No 3198

Abstract

We investigate whether satellite observations of nitrogen dioxide (NO_2) – a short-lived pollutant primarily emitted by fossil fuel combustion – can improve the forecasting of oil demand. After retrieving, cleaning, and aggregating daily satellite data, we integrate NO_2 into a range of forecasting models. Across a panel of advanced and emerging economies, we find that including NO_2 significantly enhances nowcasting accuracy relative to benchmark models based on autoregressive terms and traditional predictors such as industrial activity, prices, weather, and vehicle registrations. Accuracy gains are particularly strong during crisis episodes but remain present in more stable times. Non-linear models, especially neural networks, yield the largest improvements, highlighting the non-linear link between energy demand and pollution. By offering a timely, globally consistent, and freely available proxy, satellite-based NO_2 data provide a valuable new tool for real-time monitoring of oil demand.

Keywords: big data, satellite data, nowcasting, machine learning, energy consumption

JEL classification: C51, C81, E23, E37

Non-technical summary

Recent events (e.g. Covid-19 pandemic, Russian war against Ukraine, US trade policies, geopolitical tensions in the Middle East) highlighted the importance of high-frequency indicators for monitoring economic activity in real time. This is notably the case for oil consumption – where official statistics are published with a 3-month delay – leaving policymakers, firms, and households without timely information in periods when conditions change rapidly.

In this paper, we explore whether satellite-based measurements of nitrogen dioxide (NO_2), a pollutant primarily emitted by the combustion of fossil fuels, can be used to improve the real-time monitoring of global oil consumption. Compared with other high-frequency indicators, satellite data present four unique advantages: 1) timeliness as data are available daily, within hours of capture; 2) global coverage as satellites cover all economies uniformly, including those with scarce or unreliable statistics; 3) high granularity as measurements are available at a high spatial resolution, allowing for both national and regional analysis; and 4) free access as satellite data can be obtained at no cost, unlike many proprietary data sources.

Satellite data are however not immediately suitable for economic analysis. They must be retrieved, cleaned, and adjusted as raw NO_2 readings can be distorted by cloud cover, snow, or solar angles – implying that a large share of observations is typically missing. To address this, we apply standard filtering procedures in which we leverage spatiotemporal averaging and combine near real-time data with more robust satellite products. This provides smooth satellite-based NO_2 pollution indices for major cities globally.

We then test whether NO_2 data can enhance the forecasting of monthly oil consumption across 10 advanced and emerging economies, including the United States, China, India, Japan, South Korea, Australia, and European countries – and that represent together around 60% of world GDP and NO_2 emissions. We compare models with and without NO_2, using both simple autoregressive benchmarks and richer specifications that include

standard predictors used in the literature such as industrial production, energy prices, weather variables, and number of cars in circulation. We explore the predictive power of NO_2 using a linear OLS model as well as non-linear techniques – namely support vector regressions and neural networks which have been shown to perform well in the oil forecasting literature.

Our results show that including satellite-based NO_2 data improves the accuracy of oil demand nowcasts in a broad and robust way. Across countries, models with NO_2 data outperform autoregressive and benchmark models. On average, forecasting errors decline by 25% compared to simple AR models and by around 20% compared to more sophisticated benchmarks. While accuracy gains are particularly pronounced during the pandemic, when oil demand shifted abruptly, improvements remain significant outside of the pandemic. Accuracy gains from including NO_2 as a predictor for oil demand remain when using machine and deep learning models. Finally, improvements are notable also when comparing NO_2 with other high-frequency indicators used in the literature, such as Google Mobility indices or Google Trends.

By providing a globally consistent, timely, and freely available proxy for oil consumption, satellite-based NO_2 data offer a valuable new tool for real-time monitoring of oil consumption. This could be especially important for emerging and developing economies where official statistics are scarce or published with long delays.

Introduction

The Covid-19 pandemic brought renewed attention to high-frequency data, as traditional statistics proved too delayed to capture the sudden collapse in activity. Weekly, daily, or even hourly data were mobilized to assess in real-time the impact of lockdowns, and in doing so, challenged conventional approaches to economic monitoring. Among these, satellite-based measures of air pollution stood up, as striking images of plummeting emissions during pandemic lockdowns circulated widely in the media (see for example China in **Figure 1**).

Figure 1. Tropospheric NO_2 pollution over China: January 2020 vs. February 2020

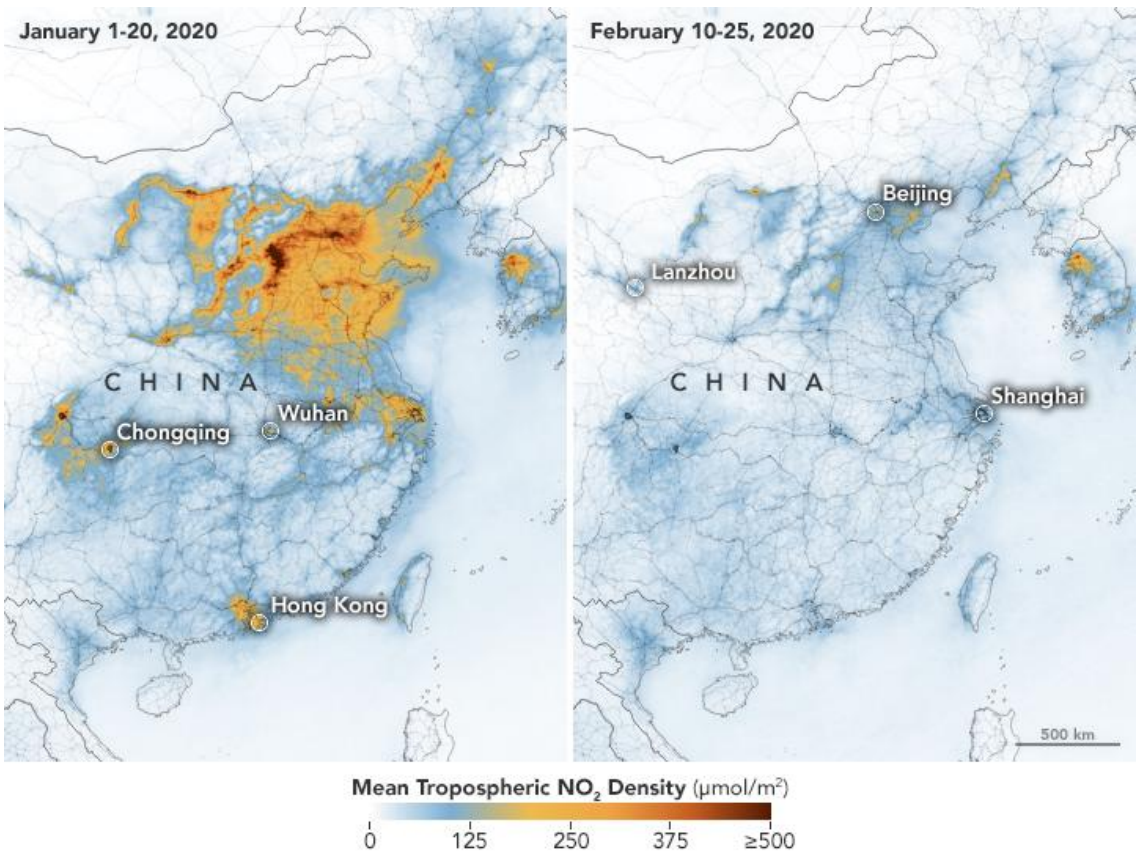

Source: Earth observatory.

Against this background, this paper investigates whether satellite data for tropospheric pollution can improve the forecasting of oil demand. This variable is key for both macroeconomic and energy analysis, yet timely information is scarce: official oil consumption data are often published with months of delay, leaving policymakers and

market participants without timely insights. Satellite-based measurements of NO_2, available in near real-time, offer a promising proxy to make up for the long publication delays of official statistics. The pollutant is primarily emitted by fossil fuel combustion and has a short lifetime in the atmosphere – making it well-suited for monitoring oil consumption. ESA satellites provide near-real-time measurements of NO_2 pollution, with information released within few hours of the capture, enabling a timely tracking of oil consumption dynamics. In addition, satellite data offer a global coverage with uniform quality, high spatial granularity, and free accessibility – providing key advantages over traditional statistics and even alternative high-frequency indicators.

We test whether NO_2 satellite data can enhance the forecasting performance of monthly oil consumption across 10 advanced and emerging economies, including the United States, China, India, Japan, South Korea, Australia, and European countries – and that represent together around 60% of world GDP and NO_2 emissions. We compare models with and without NO_2, using both simple autoregressive benchmarks and richer specifications that include standard predictors used in the literature such as industrial production, energy prices, weather variables, and number of cars in circulation. We explore the predictive power of NO_2 using a linear OLS model as well as non-linear techniques – namely support vector regressions and neural networks which have been shown to perform well in the oil forecasting literature.

We show that satellite-based NO_2 pollution enhances the real-time nowcasting of oil demand across a set of advanced and emerging economies. Compared with standard linear benchmarks, models incorporating NO_2 data deliver significant accuracy gains – not only relative to simple autoregressive benchmarks but also to richer models including abovementioned established predictors. Importantly, while accuracy gains are especially large during the pandemic, improvements in forecast performance persist even in more stable periods. Moving to non-linear frameworks consistent with more recent studies on predicting – namely support vector regressions and neural networks – we find that accuracy gains remain robust when using machine and deep learning models that capture non-linear dynamics. Finally, we find that using satellite-based NO_2

pollution yields more accurate forecasts than incorporating other high-frequency indicators used in the literature such as Google Mobility indices and Google Trends.

Our contribution to the literature is threefold. First, we extend the growing literature on satellite data in economics (e.g. Henderson et al., 2012; Donaldson and Storeygard, 2016) by exploring the potential of NO_2 pollution – which have been used very little in economics (see Ezran et al., 2023). More broadly, while a vast literature has shown how economic conditions influence air pollution (e.g. Boersma and Castellanos, 2012; Du and Xie, 2017; Le et al., 2020; Baimatova et al., 2020), very few have explored the reverse relationship where air pollution becomes an early indicator for economic conditions. Second, we add to the body of research on oil consumption forecasting by demonstrating the usefulness of satellite-based indicators on NO_2 pollution as a novel real-time proxy for energy demand. We demonstrate its value-added alongside more established predictors such as industrial activity, energy prices, weather conditions, and number of vehicles – while also highlighting its potential relative to alternative high-frequency indicators used in most recent analyses (e.g. Yu et al., 2019; Dey and Das, 2022). Third, we contribute to the debate on high-frequency data (see Ferrara et al., 2020; Carvalho et al., 2021; Li et al., 2021; Bricongne et al., 2023) by quantifying the benefits of daily satellite observations, highlighting their role in improving short-term forecasting – especially during crisis episodes.

The rest of the paper is organized as follows: **Section 1** reviews related literature, **Section 2** explains the choice of relying on NO_2 from satellite data – a source largely uncharted in the literature, **Section 3** details data processing and aggregation, while **Section 4** compares nowcasting performances across different linear and non-linear models.

Section 1: Literature review

Our paper first contributes to the growing literature that leverages satellite data for economic analysis – see Donaldson and Storeygard (2016) for a review. Following the seminal contribution of Henderson et al. (2012), satellite data on night-light intensity has been widely employed as a proxy for economic development, particularly in data-scarce countries (Ebener et al., 2005; Keola et al., 2015; Jean et al., 2016; Pinkovskiy and Sala-i-Martin, 2016). More recently, this approach has been applied to monitor the impact of specific shocks, such as India's demonetization (Chodorow-Reich et al., 2020) and the Covid-19 pandemic (Beyer et al., 2021). Methodological refinements have also emerged, with for instance Goldblatt et al. (2019) showing that daytime reflectance outperforms night lights in predicting enterprise counts, employment, and expenditure. However, limitations of both night and daytime light measures have been documented, notably in high-density areas and in advanced economies where their correlation with economic activity weakens (Sutton et al., 2007; Chen and Nordhaus, 2010; Tanaka and Keola, 2017; World Bank, 2017; Hu and Yao, 2019). In response, researchers have explored alternative satellite-based indicators such as land and forest cover (Lobell, 2013), infrared satellite images (d'Aspremont et al., 2025), or satellite-derived measures of floor space (Egger et al., 2023). Satellite data on air pollution – specifically NO_2 – has also gained prominence with Bricongne et al. (2021) and Ezran et al. (2023) demonstrating its usefulness for nowcasting industrial production and real GDP, respectively. Building on this strand of research, this paper is the first, to the best of our knowledge, to investigate the use of satellite-based NO_2 to predict oil demand.

Our paper also extends the literature on forecasting oil consumption – a topic of enduring relevance in both economics and energy research – by demonstrating the merits of satellite-based NO_2 pollution as a predictor.[1] Historically, energy demand prediction has relied primarily on linear models for their interpretability and ease of implementation (Mackay and Probert, 1994; Harris et al., 2018; Boamah, 2021). Over

[1] An extensive summary of the literature on forecasting energy consumption can be found in **Table A1** in **Appendix A**.

time, ordinary least squares (OLS) models were supplanted by more sophisticated time-series techniques such as autoregressive distributed lag (ARDL) models (e.g. Dilaver and Hunt, 2011; Fatima et al., 2019; He and Lin, 2018). Many of these approaches have incorporated Mixed Data Sampling (MIDAS) methods (He and Lin, 2018) or adjustments for trend in the context of long-term forecasting.[2] Other commonly used techniques include ARIMA and moving average models (Edigar and Akar, 2007; Rehman et al., 2017; Xu and Wang, 2010), Bayesian linear regressions (Chai et al., 2012) and vector autoregressions (Crompton and Wu, 2005). However, despite their widespread use, linear models face limitations when dealing with many predictors as they are prone to overfitting. This had long posed a significant challenge given the broad set of determinants for oil consumption such as real GDP (e.g. Mackay and Probert, 1994; Crompton and Wu, 2005), trade (e.g. Assareh et al., 2010), consumer price inflation (e.g. Zhu, 2023), weather conditions[3] (e.g. Fatima et al., 2019), population (e.g. Harris et al., 2018; Al-Fattah, 2021), oil prices (e.g. Aldabbagh et al., 2024) and number of vehicles (e.g. Behrang et al., 2011; Chai et al., 2012). Our paper extends this literature by proposing to use satellite-based NO_2 pollution as another predictor for oil consumption. This variable offers a dual advantage: it significantly improves short-term forecasting accuracy while also being available across all countries.

Our paper also delves into the more recent literature on forecasting oil consumption, by exploring machine and deep learning models. Although traditional econometric approaches remain widely used, recent research on oil demand forecasting increasingly relies on machine learning and deep learning methods. These approaches are particularly well-suited to capturing the non-linear and multifactorial nature of oil demand. To address the "curse of dimensionality" (Bellman, 1957) – arising when the

2 Long-term forecasts of energy demand often rely on the concept of "*underline energy demand trend*" (UEDT) introduced by Hunt et al. (2003), which captures technical changes in oil reserves as well as exogenous factors (Dilaver and Hunt, 2011; Fatima et al., 2019). In addition, some linear models have been refined to account for seasonality and trend (Romero-Gelvez et al., 2020).

3 Liu et al. (2021) finds that temperature variables are particularly important drivers for short- and medium-term energy demand forecasts. When temperatures are lower than a neutral threshold (often around 18 degrees Celsius), energy consumption increases for heating but when temperatures are higher, energy consumption increases for air conditioning. In line with the former, INSEE (2010) reports that a 1-degree decline in temperature can raise energy consumption by up to 2%.

number of potential predictors exceeds the number of observations – the literature has first employed feature selection methods such as genetic algorithms (Assareh et al., 2010), gravitational search algorithms (Behrang et al., 2011), and social spider optimisation (Alkhammash et al., 2022). Beyond feature selection, a growing body of work has focused on machine and deep learning methods capable of accommodating numerous predictors and of capturing complex, non-linear relationships in oil consumption. Among these, support vector regressions (Romero-Gelvez et al., 2020; Zhu, 2023) and neural networks (Liu et al., 2016; Yu et al., 2019; Al-Fattah, 2021; Huang et al., 2021; Aldabbagh et al., 2024) consistently emerge as the best-performing model classes, especially in short-term energy forecasting where high-dimensional input data is common (Liu et al., 2021). Our paper contributes to this strand of literature by demonstrating that the accuracy gains from incorporating NO_2 pollution into oil demand forecasting models remain, even when using these non-linear methods.

Our paper finally adds to the growing literature on alternative and high-frequency data. Following the Covid-19 pandemic, such data have been increasingly used for economic analysis – for example credit card transactions (Carvalho et al., 2021), online housing listings (Bricongne et al., 2023), and social media data (Li et al., 2021).[4] Although high-frequency data are often noisy and their contribution to forecast performance remains disputed (Ahnert and Bier, 2001; INSEE, 2020), their timeliness can lead them to outperform traditional data in nowcasting applications (Ferrara et al., 2020; Fosten and Nandi, 2025). In that respect, our work is closely related to Yu et al. (2019) who use Google Trends data to nowcast global oil consumption by incorporating search intensity for keywords such as "*oil consumption*", "*oil price*", and "*oil reserves*" into support vector regressions and neural networks. Our paper differs mainly in the indicator considered, as well as the broader country scope. Our study benefits from the fact that satellite-based data offer a consistent cross-country coverage, whereas quality and availability of Google Trends vary substantially across geographies depending on Google's market

[4] See Chetty et al. (2020) and Bricongne et al. (2020) for reviews of alternative data for the US and France during the Covid-19 pandemic, respectively.

share – notably, Google Trends are not available for China.[5] Moreover, Google trends mainly capture household behaviour, which may not fully reflect demand from other agents such as businesses and public administration.

[5] It can be added that Google Trends data are prone to substantial revisions by design given the scaling of the series between 0 and 100 (maximum). On the contrary, satellite-based pollution data are not subject to revisions.

Section 2: Why satellite data on NO_2?

Compared to other data sources, satellite data offer four key advantages: timeliness, global coverage, high granularity, and free accessibility. First, satellite data are available daily and released within hours of capture, thereby making it a suitable candidate to make up for substantial publication delays of official statistics. For instance, energy consumption statistics from the US Energy Information Administration (EIA) – one of the few institutions publishing cross-country data – are released around two months after the end of the month of interest (**Figure 2**). During this period, little to no public information is available. While such lags might be less problematic in "normal" times, where economic conditions remain broadly stable, this can become critical when sudden shocks occur and require a real-time assessment of oil demand fluctuations. In this context, the timeliness of daily satellite data is particularly valuable. Second, satellite data provide consistent global coverage, including in regions where official statistics are scarce or entirely absent, notably emerging and developing economies. The uniform coverage of satellite data also reduces potential biases that otherwise arise from country-specific data collection or reporting practices. Third, the spatial granularity of satellite data – available at resolutions of 5×3.5 km^2 areas – makes it possible to capture events at regional level whereas official statistics are typically reported only at national level. Finally, satellite data are freely available, further reinforcing their attractiveness as a complementary source for economic analysis.

Figure 2. Timeline for indicators of oil consumption

Sources: US EIA and authors.

While satellites measure concentration of several pollutants, we focus on NO_2 as the most suitable indicator for tracking oil consumption in real-time. There are four main reasons for this choice. First, NO_2 has a relatively short lifetime in the troposphere

compared with other pollutants (Lamsal et al., 2011), which makes it especially relevant for near-real-time monitoring. [6] Second, NO_2 is primarily emitted through the combustion of fossil energies, directly linked to energy use. Third, NO_2 acts as a precursor for other pollutants (Henneman et al., 2017), making it particularly valuable as an "early" indicator. Finally, empirical evidence confirms a robust relationship between energy consumption and NO_2 emissions (e.g. Cui et al., 2020; Hu and Guo, 2020; Abbas et al., 2022). This link has already been exploited during crisis episodes such as the Great Financial Crisis in 2008-2009 (Boersma and Castellanos, 2012; de Ruyter de Wilt et al., 2012; Russell et al., 2012; Du and Xie, 2017) and the Covid-19 pandemic (Le et al., 2020; Tobias et al., 2020; Baimatova et al., 2020; Keola and Hayakawa, 2021).

[6] NO_2 typically stays in the troposphere for only a few hours to up to one day – much shorter than other pollutants such as fine particles ($PM_{2.5}$ or PM_{10}) which persist for several days to weeks, carbon monoxide (CO; around one month), and methane (CH_4; around 10 years). This short lifetime reflects the high reactivity of NO_2 to chemical processes in the troposphere, notably its transformation into ozone and other secondary pollutants (Vincent et al., 2010). As a result, NO_2 does not travel over long distances, making it particularly relevant for capturing local activity.

Section 3: Data retrieval, cleaning, and aggregation

Our data are derived from satellite observations of tropospheric pollution collected by the TROPOMI instrument on-board the Sentinel-5P satellite, operated by the European Spatial Agency (ESA).[7] In technical terms, TROPOMI is a nadir-viewing imaging spectrometer, meaning it observes the atmosphere by looking straight down rather than diagonally. Tropospheric pollution is measured using passive remote sensing techniques that detect solar radiation reflected and radiated by the Earth's surface and atmosphere. Sentinel-5P follows a Sun-synchronous orbit, passing over each location daily at approximately the same local solar time (around 13:35), which ensures consistent mid-day observations across the globe.[8] As the TROPOMI instrument was started only recently, our dataset begins in July 2018.

We apply several steps to retrieve, clean, and aggregate the data, ensuring that we obtain a dataset that is balanced and is aggregated at a level comparable to oil consumption statistics (i.e. country level). These steps can be broadly separated into *pre-processing* in **Section 3.1** (i.e. general adjustments on raw data to make them more tractable for further usages) and *post-processing* in **Section 3.2** (i.e. specific aggregations and corrections on the pre-processed data to make them fit for oil nowcasting).

7 Earth atmosphere has a series of layers. Moving upward from ground level, these layers are the *troposphere* (0 to around 15 km), stratosphere, mesosphere, thermosphere and exosphere.

8 Other sources of NO_2 pollution data are available, but our choice of using TROPOMI is motivated by its global coverage, uniqueness of the sensor, and enhanced precision. Compared with ground-based monitoring stations which provide limited and location-dependent coverage determined by the (arbitrary) placement of sensors, TROPOMI offers global observations of uniform quality. Its single sensor also alleviates risks of idiosyncratic errors that can arise from ground-based instruments (e.g. due to malfunction of some specific sensors). Our choice is supported by the geophysical literature that advocates for the use of remote sensing over ground-based observations for large-scale studies (Remer et al., 2005). Finally, TROPOMI data have been shown to be more precise than alternative satellite measurements such as NASA's OMI, which in recent years has suffered from various instrument malfunctions (Griffin et al., 2019; Wang et al., 2020).

3.1. Pre-processing phase: data retrieval and initial cleaning

The retrieval process is automatized on a deviant server that automatically collects Sentinel-5P satellite data from the Copernicus Browser platform and checks the presence of the relevant observations. Among the variables provided by the Sentinel-5P dataset, we focus on the "*NO_2 tropospheric column concentration*". We use the second version (V2) of the dataset which offers NO_2 data in both "*offline*" (OFFL) and "*near-real time*" (NRTI) designs.[9] Between the two, the trade-off lies between timeliness and robustness. NRTI data are available only within a couple of hours of image acquisition while OFFL data are typically available within one week. This is because of differences in the weather data employed for the weather-adjustment of the satellite observations: NRTI relies on weather forecast, whereas OFFL incorporates *ex post* meteorological data – therefore more accurate to perform weather-normalization.[10] In addition, OFFL data has the advantage of correcting for potential retrieval delays from the satellite (Verhoelst et al., 2021). We fetch both OFFL and NRTI datasets and combine them in the post-processing phase (see **Section 3.2**) to benefit from both the timeliness of NRTI and the robustness of OFFL. Upon retrieval, the algorithm checks automatically for the consistency of entering files – notably whether NO_2 data are available – and flattens the data for further use.

The raw data are processed into Level 3 (L3) swaths, a gridded data cube with constant latitude and longitude coordinates, that can be more readily exploited for analysis. Copernicus provides the data as Level 2 (L2) swaths, consisting of pixels of 5x3.5 km^2 with centroid coordinates (longitude and latitude) and a quality index (between 0 and 1). Following ESA recommendations, we exclude all observations with a quality below 0.75, as snow cover, cloud presence, or variations in solar zenith and viewing angles can significantly alter TROPOMI estimates (ESA, 2020; Wang et al., 2020). Next, observations

9 NO_2 data are also available in a "*reprocessed*" (RPRO) design, but only available until June 2022.

10 Weather-normalization is a critical step when working with atmospheric pollution as pollution levels are very sensitive to meteorological conditions, which not only affect the chemical process of pollutant formation but also influence human behaviours that generate emissions (Rao and Zurbenko, 1994). The influence of weather – especially wind and temperature – has been found to often exceed the effect of policy interventions or economic events (Anh et al., 1997; Alix-Garcia and Millimet, 2021). In this paper, we rely on ESA-provided data that are already weather-normalized, which alleviates the need for additional corrections such as the steps performed in Bricongne et al. (2021).

are allocated at the municipal level. Starting from the geo-coordinates of each city centre of interest, we collect NO_2 measurement within a 40 km radius – broadly consistent with ranges reported in atmospheric science literature and supported by empirical tests and measurements (**Figure A1** in **Appendix A**).[11] This approach ensures consistency across cities by relying on an equal amount of input data, while also accounting for possible movements of NO_2 in the troposphere under windy conditions.

Overall, this pre-processing step (**Figure 3**) condenses the daily raw data (1-2 Gb across various datasets) into a single file of manageable size (10-20 Mb), facilitating subsequent analysis.

[11] The choice of a 40-km radius reflects a trade-off between including a sufficiently large set of valid observations for statistical robustness and ensuring that the signal remains representative of local urban activity. To validate our choice, we also tested the utilization of 20-km radius around city centres. The results using a smaller radius were consistently worst. As shown in **Figure A1** in **Appendix A**, a 20-km radius is insufficient to capture all the NO_2 dynamics around the major metropolitan areas used in this study. While we could have used grid-search techniques to test all possible radius between 10 and 50-km, we rejected this approach as it would have been computationally intensive and prone to overfitting the data. This radius is also in line with the atmospheric science literature, which shows that short-lived pollutants such as NO_2 typically disperse over few kilometres only. For example, Lamsal et al. (2011) and Russell et al. (2012) report dispersion ranges of 20-50 km, while Boersma et al. (2009) and de Ruyter de Wilt et al. (2012) adopt similar buffer sizes in urban air-quality analyses. Against this background, a 40-km radius provides a reasonable balance between coverage and representativeness.

Figure 3. Data pre-processing

.

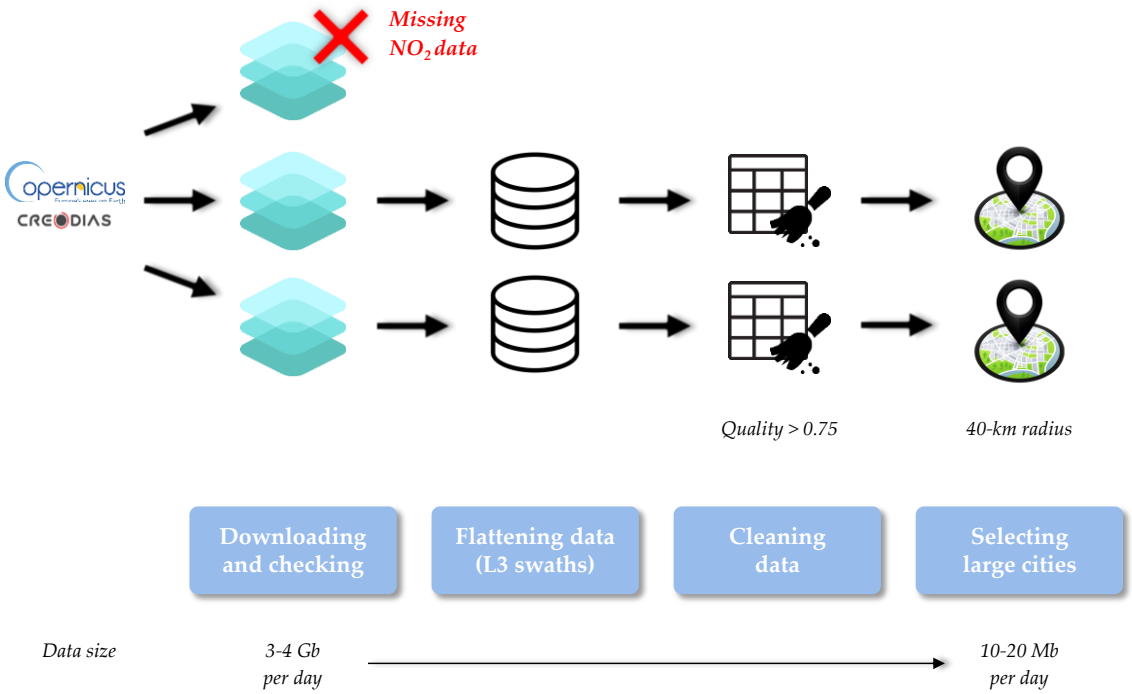

Source: authors.

3.2. Post-processing phase: dataset combination and aggregation

A key challenge when working with satellite data is the prevalence of missing values, with 10 to 40% of the observations rendered unusable by cloud cover, snow, or unfavourable viewing conditions depending on the country.[12] This well-documented issue has motivated a wide range of methods to address it. The geo-statistics literature emphasizes that the most accurate approaches exploit both spatial and temporal correlations, with spatiotemporal "kriging" regarded as the state of the art (Laslett, 1994; Tadic et al., 2017; Peng et al., 2020; Shao et al., 2020; Yang and Hua, 2018).[13] Related

[12] In addition, a key limitation of passive satellite observations is their dependence on sunlight to take measurements. This is why remote sensing generally performs less well for countries near the poles, such as Canada or Russia, where low luminosity and solar angle reduce retrieval accuracy (van Geffen et al., 2022).

[13] Another class of popular methods are "gap-fill" algorithms (Weiss et al., 2014). While spatiotemporal kriging and gap-fill algorithms aim to reconstruct missing values in environmental datasets, their methodological foundations differ. Kriging is a geostatistical approach that explicitly models spatial

contributions employ econometric methods ranging from linear extrapolation (Zhang et al., 2017) to machine and deep learning (Bi et al., 2019; Chi et al., 2020; Fouladgar and Främling, 2020). However, these sophisticated approaches often involve substantial computational costs, which make them poorly suited to high-frequency, large-coverage datasets – where methods like "kriging" become intractable due to the need to invert large covariance matrices (Kianan et al., 2021).[14] Instead, most studies using high frequency (daily or intra-daily) data with large coverage – as in our paper – adopt simpler techniques such as linear extrapolation or local averages, which have been shown to provide sufficient accuracy (e.g. Noor et al., 2006; Hirabayashi and Kroll, 2017). Against this background, we employ a local-averaging approach, striking a practical balance between accuracy and computational feasibility. Our method incorporates both spatial and temporal dimensions:

- **Spatial dimension**: we reduce the data to a single observation per municipality per day by computing the unweighted average of all valid measurements within the city radius. We implement this on both NRTI and OFFL datasets.

- **Temporal dimension**: we smooth the resulting series using an unweighted 28-day moving average. The choice of a 28-day window accounts for weekly seasonality – being a multiple of 7. This allows us to track monthly dynamics NO_2 dynamics while reducing short-term noise.

We combine the NRTI with the OFFL datasets to leverage their respective strengths, incorporating the timeliness of NRTI while retaining the robustness of OFFL. To construct the smoothed series, we apply a 28-day moving average, such that any observation Y_t at time t is the average of the previous 28 observations. Since OFFL data are available with a publication lag of 8 days, the computation follows equation (1):

and temporal correlation structures, yielding both predictions and associated uncertainty estimates. Gap-fill algorithms, by contrast, encompass a wide range of heuristic or data-driven methods that are often computationally lighter and more scalable, but typically do not quantify uncertainty.

14 While a "lattice kriging" has been developed to overcome this issue over large datasets (Nychka et al., 2015), it remains important when dealing with world data at daily frequency. In addition, the same technical difficulty is faced with "gap-fill" algorithms.

(1) $$Y_t = \frac{(\sum_{t}^{t-7} NRTI + \sum_{t-8}^{t-27} OFFL)}{N_{NRTI} + N_{OFFL}}$$

Where N_{NRTI} denotes the number of non-missing NRTI values between t and $t-7$, and N_{OFFL} the number of non-missing OFFL values between $t-8$ and $t-27$. The main advantage of our approach is that it exploits the complementarity of missing patterns across the two datasets (**Figure 4**) since outages typically affect them at different times. Finally, this also ensures that the moving average series remains continuous even when one source dataset temporarily fails.

Finally, we aggregate city-level observations to the country level. To ensure that meaningful variation in NO_2 pollution is not diluted by broadly stable pollution over sparsely populated areas, we restrict the sample to major urban centres, selecting for each country a maximum of 25 cities with more than 275,000 inhabitants.[15] At this stage, any remaining missing values at city-level are extrapolated to obtain a balanced dataset. The resulting city-level series are then aggregated into a country-level indicator by aggregating the NO_2 concentration across the cities.

Figure 4. Differing patterns of missing data between NRTI and OFFL datasets

Panel a) Share of missing data in NRTI dataset

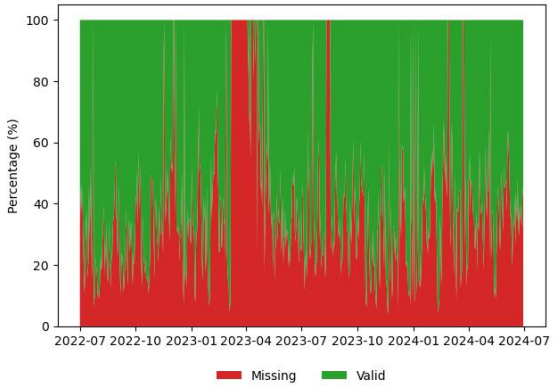

Panel b) Share of missing data in OFFL dataset

Sources: ESA and authors.
Note: both panels refer to China.

15 The thresholds of 275,000 inhabitants and 25 cities aim at striking a balance between ensuring sufficient representativeness of urban activity and maintaining comparability across countries of different sizes, while avoiding undue influence from smaller, less representative localities. Our empirical tests have shown little benefit from adding more than 25 cities per country, while the computation costs increase linearly with the inclusion of new cities - this further justified our choice as a trade-off between including many cities and keeping the processing algorithm tractable.

Section 4: Nowcasting oil demand

4.1. Comparing against benchmark linear models

We now assess whether satellite-based NO_2 pollution data can improve the real-time forecasting of oil demand. To this end, we evaluate the out-of-sample performances of nowcasting models that incorporate NO_2 pollution data, using two benchmarks. The first one is a straightforward autoregressive (AR) model as in equation (2) which we compare with an augmented specification that adds our satellite-based index of NO_2 pollution alongside the AR term. The second one is a more sophisticated benchmark that includes the predictors commonly used in the oil-demand forecasting literature, which we then compare to the same specification augmented with our NO_2 pollution index. More specifically, this second (more sophisticated) benchmark predicts oil consumption (y_t) using an AR term together with variables for industrial activity (IP_t), energy prices (P_t), number of vehicles (V_t), and weather conditions (W_t) as formalized in equation (3). The specification accounts for publication delays faced by real-time forecasters by taking lagged values when appropriate, as denoted by subscripts k in equation (3).[16] In our baseline specification, we use industrial production as the variable IP_t following Dilaver and Hunt (2011), the Producer Price Index (PPI) as variable P_t following Aldabbagh et al. (2024), new car registrations as variable V_t following Behrang et al. (2011) and Chai et

[16] Oil consumption is generally published with at least a two-month delay, so we set k_y to 2. The same lag applies to most macroeconomic variables (e.g. industrial production). By contrast, Purchasing Managers' Index (PMI) enters with only a one-month lag ($k_{ip} = 1$) given its timeliness while weather variables and oil prices enter contemporaneously ($k_w, k_p = 0$) as they are available in quasi real-time. All variables are expressed in log differences – except for PMIs and weather variables which enter directly in levels. Further details on data are available in **Table A2** in **Appendix A**.

al. (2012),[17] and temperatures as variable W_t following INSEE (2010) and Liu et al. (2021).[18]

$$y_t = \beta_0 + \beta_1 \cdot y_{t-k_y} + \varepsilon_t \quad (2)$$

$$y_t = \beta_0 + \beta_1 \cdot y_{t-k_y} + \beta_2 \cdot IP_{t-k_{ip}} + \beta_3 \cdot P_{t-k_p} + \beta_3 \cdot V_{t-k_v} + \beta_4 \cdot W_{t-k_w} + \varepsilon_t \quad (3)$$

We evaluate forecasting performance by comparing out-of-sample root mean squared errors (RMSE) across models, using a recursive real-time nowcasting exercise that replicates the information set that would have been available to a real-time forecaster. Specifically, the model is estimated (in-sample) up to month $t-1$, and oil consumption is then nowcasted (out-of-sample) for month t. The exercise starts in July 2018, with the first out-of-sample nowcast in January 2020 and the last one in December 2024. We conduct this procedure separately for each of the 10 advanced and emerging economies in our sample (Australia, China, France, India, Italy, Japan, South Korea, Spain, United Kingdom, United States) for which oil consumption and control variables are available at monthly frequency – and that represent together around 60% of world GDP and NO_2 emissions.[19] To reconcile the frequency mismatch between *daily* NO_2 pollution and *monthly* oil consumption, we use monthly average of NO_2 pollution – computed as in **Section 3**.[20] On average, models incorporating satellite-based NO_2 pollution outperform

[17] We use new car registrations (i.e. the *flow*) rather than the *stock* of vehicles as up-to-date monthly data on the latter are unavailable. The series cover both electric and thermal vehicles, since comparable cross-country data distinguishing between the two was not available to us. This implies that in countries with high penetration of electric vehicles, new registrations might overstate oil consumption. However, over our sample period (2018-2024), electric vehicles accounted for only 8% of global car sales (IEA, 2025). In addition, empirically NO_2 pollution still significantly improves oil-demand forecasting in countries with above-average penetration of electric vehicles, such as China (where 17% of car sales were electric vehicles over our sample).

[18] To preserve parsimony given the limited timespan of our dataset, we include only one variable per category in the benchmark specification. Since different proxies have been used per category in the literature (see **Table A1** in **Appendix A**), we also tested alternative benchmark specifications. The results, reported in **Figure A2** in **Appendix A**, are broadly consistent across specifications and confirm higher forecasting accuracy when satellite-based NO_2 pollution is included as a predictor.

[19] For other countries, monthly series for oil consumption are not available.

[20] We also examined the time dimension of *daily* satellite-based NO_2 pollution by running *daily* nowcasts of monthly oil consumption. Specifically, for each day d of month t, we estimate the model (in-sample) up to month $t-1$ and then nowcast (out-of-sample) oil demand in month t using the information set available on day d – that is, daily NO_2 observations up to day $d-1$. Results are reported in **Figure A3** in **Appendix A** and show that accuracy improves progressively as more information becomes

benchmark models, notably during crisis episodes. **Figure 5** presents results averaged across countries and shows that models including daily NO_2 satellite-based pollution perform better than a naïve AR model (panel a) and then a more sophisticated benchmark model (panel b). Over the full sample (2020-2024), adding NO_2 pollution as a predictor improves forecasting accuracy by 26% on average relative to the naïve AR model and by 22% relative to the multi-variable benchmark model. These averages, however, mask some cross-country heterogeneity with gains up to 60% in some cases (e.g. China, Spain) but closer to 5% in other (e.g. Japan, South Korea). Detailed results in **Table A3** in **Appendix A** show these accuracy gains are nevertheless significant for most countries, based on a Clark and West (2007) test.

To better understand whether the forecasting gains from NO_2 pollution are concentrated in periods of heightened volatility, we then separate the analysis between crisis and non-crisis times. In line with the literature suggesting that high-frequency data enhance forecasting performance primarily during crisis episodes (e.g. Jardet and Meunier, 2022), we split the sample between the pandemic year (2020) and the subsequent period (2021-2024). In line with the literature, our results indicate substantially larger gains during the pandemic, with models including NO_2 pollution outperforming the naïve AR model by more than 30% on average. Nevertheless, accuracy gains persist in more stable periods, with improvements of 10-15% over the benchmarks between 2021 and 2024. Detailed results in **Table A4** in **Appendix A** show that accuracy gains remain statistically significant for several countries (e.g. United States, China, Spain) even outside of crisis times.

available, consistent with intuition. Nevertheless, the exercise shows that even on the first day of the month, including NO_2 pollution in the model yields accuracy gains.

Figure 5. Out-of-sample RMSE of OLS models incorporating satellite-based NO_2 pollution – relative to standard benchmarks

(average across countries)

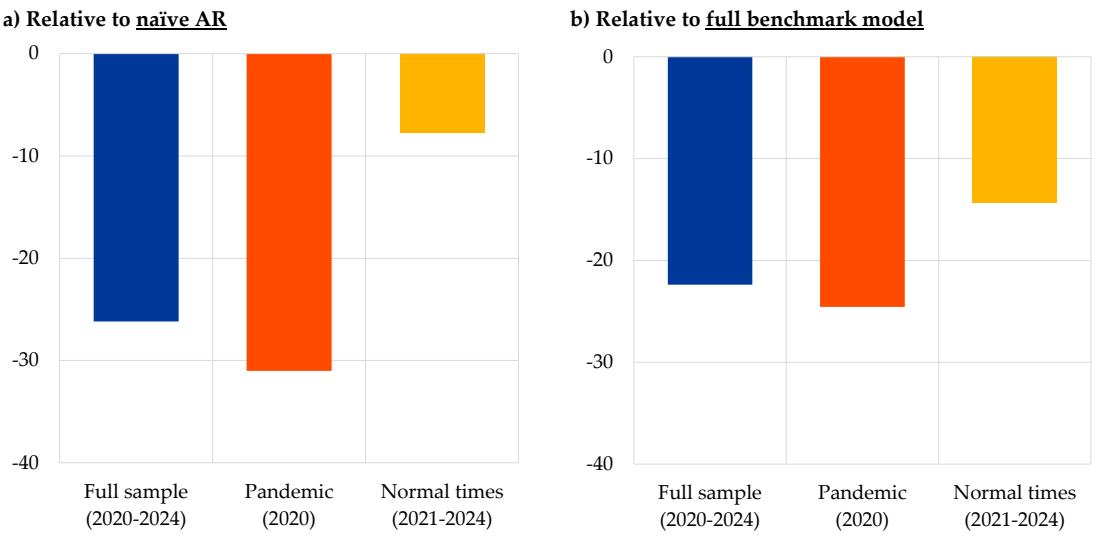

Sources: Haver Analytics, US Energy Information Administration, International Energy Agency, OECD, Bloomberg and authors.
Notes: Countries in the sample are Australia, China, France, India, Italy, Japan, South Korea, Spain, the United Kingdom and the United States. The full benchmark model incorporates an AR term, a variable for industrial activity (industrial production), a variable for prices (PPI), a variable for the number of vehicles (new car registrations), and a variable for weather (temperature). Results are presented relative to the benchmark: a negative value indicates an over-performance of the model with satellite-based NO_2 pollution. Results by country are available in ***Tables A3 and A4*** *in* ***Appendix A****.*

We finally benchmark our satellite-based NO_2 pollution index against alternative high-frequency indicators that have been proposed in the literature to nowcast oil demand, namely Google Mobility indices (Dey and Das, 2022), Google Trends (Yu et al., 2019) and the Oxford Stringency indices (Aldabbagh et al., 2024).[21] As mentioned in **Section 2**, a key advantage of the satellite index is its global coverage, whereas Google-based indicators are unavailable for China due to the platform's ban. Moreover, since Google indices rely on user activity, their quality depends on Google's market penetration in the country, which can be limited or biased toward specific populations segments. In addition, the Oxford Stringency index (Hale et al., 2021), while a series with global coverage, was a one-off effort and stops shortly after the pandemics – just like the Google

[21] All high-frequency indicators enter with no lag as they are available in quasi real-time. Google Trends are expressed in first differences; Google Mobility and Oxford Stringency indices (which both measure deviation to a normal situation) both enter directly in levels.

Mobility indices.[22] Results are reported in **Figure 6**.[23] On average, models including satellite-based NO_2 pollution outperform moderately those using Google Mobility and Stringency indices (by 1%) but more markedly Google Trends (by 13%). But when distinguishing between pandemic and post-pandemic periods, a sharper pattern emerges. During the pandemic, forecasting performance is somewhat better for Google Mobility and Oxford Stringency indices – whereas in normal times, the NO_2 index outperforms significantly all benchmarks by 10-20%. Overall, these findings suggest that while pandemic-specific indicators (Google mobility and Oxford Stringency indices) can be relevant during the Covid-19 period, their predictive power diminishes strongly in normal times. By contrast, the NO_2 pollution index retains its relevance outside crisis, as evidenced by its outperformance relative to both alternative high-frequency indicators (**Figure 6**) and the standard benchmarks (**Figure 5**).

22 The extrapolation of missing observations for the Oxford Stringency index and Google Mobility indices (before and after the pandemic) is made by assuming no restrictions prior to the Covid-19 period, and a gradual return to no restriction after the last observation (October and December 2022 for Google Mobility and Oxford Stringency, respectively).

23 Following Yu et al. (2019), we use Google Trends for the keyword *"oil consumption"* which we translate into local languages of the countries in our sample. We focus on the keyword "*oil consumption*" as in Yu et al. (2019) who find this keyword to yield better forecasting improvements over other keywords tested (e.g. "*oil prices*", "*oil inventories*").

Figure 6. Out-of-sample RMSE of OLS models incorporating satellite-based NO_2 pollution – relative to alternative high-frequency indicators

(average across countries)

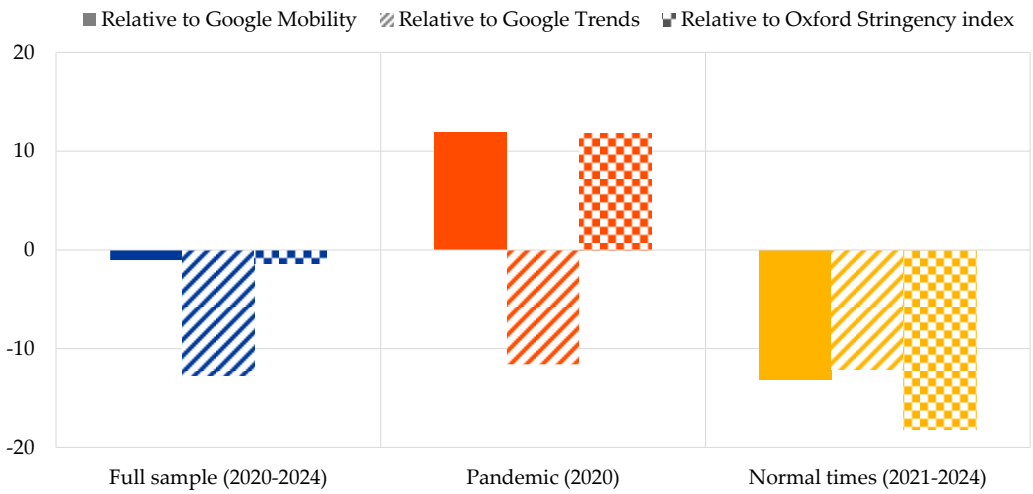

Sources: Haver Analytics, US Energy Information Administration, International Energy Agency, OECD, Bloomberg and authors.
*Notes: Countries in the sample are Australia, France, India, Italy, Japan, South Korea, Spain the United Kingdom and the United States. China is excluded as Google indices are not available. The benchmark model incorporates an AR term, a variable for industrial activity (industrial production), a variable for prices (PPI), a variable for the number of vehicles (new car registrations), and a variable for weather (temperature). Results are presented relative to the benchmark: a negative value indicates an outperformance of the model with satellite-based NO_2 pollution. Results by country are available in **Table A5** in **Appendix A**.*

4.2. Exploring non-linear models

We now turn to non-linear models – specifically Support Vector Regressions (SVR) and artificial neural networks which have consistently shown strong predictive performance in prior studies (Yu et al., 2019; Romero-Gelvez et al., 2020; Huang et al., 2021; Zhu, 2023; Aldabbagh et al., 2024) – by extending the OLS analysis of **Section 4.1** to these specifications.

The SVR model adapts the Support Vector Machine (SVM) framework of Cortes and Vapnik (1995), originally developed for binary classification. The underlying idea is to project input data into a higher dimensional space - called "*feature space*"- where linear separation becomes feasible (**Figure 7,** panel a). In regression form, SVR follows a similar approach and estimates the relationship between explanatory and dependent variables using a "*kernel*" function that governs the transformation of input data. Predictions allow for a margin of error around the fitted function: observations lying within this margin

are ignored, while those outside – the *"support vectors"* – determine the fit (**Figure 7**, panel b). Formally, SVR predicts oil demand y_t as in equation (4):

$$y_t = \sum_{i \epsilon T} (a_i - a_i^*)\, K(X_i, X_t) + \beta_0 \tag{4}$$

Where X denotes the vector of explanatory variables, i the training period ($T = 0, .., t - 1$), and t the predicted month. The kernel function $K(X_i, X_t)$ compares current predictors (X_t) with support vectors, assigning higher weights (a_i, a_i^*) to more similar observations.[24] Equation (4) can also be written as in equation (5):

$$y_t = \omega \cdot \Phi(X_t) + \beta_0 \tag{5}$$

Where Φ is the mapping function into the feature space, for which ω (weights of the support vectors) and β_0 (constant) are obtained by solving the minimization problem in equation (6):

$$\min_{\omega, \beta_0, \xi_i, \xi_i^*} \frac{1}{2}\|\omega\|^2 + C \sum_{i \epsilon T} (\xi_i + \xi_i^*) \tag{6}$$

Subject to:

$$\begin{cases} y_i - \omega \cdot \Phi(X_i) - \beta_0 \leq \varepsilon + \xi_i \\ \omega \cdot \Phi(X_i) + \beta_0 - y_i \leq \varepsilon + \xi_i^* \\ \xi_i, \xi_i^* \geq 0 \end{cases}$$

Where ξ_i and ξ_i^* are slack variables measuring deviations outside the error band ε (the *"tube width"* or error-insensitive loss function), and C is a regularization parameter controlling the trade-off between model complexity and training error. A wider tube width ε allows more volatility around the kernel and, thus, fewer support vectors.[25]

[24] Specifically, weights are allocated such that $(a_i - a_i^*) \neq 0$ for support vectors, and that higher weights are allocated to observations i for which X_i are more similar to X_t.

[25] Following Yu et al. (2019), we use a radial basis function (RBF) kernel formally defined as $K(x_i, x) = \exp(-\gamma\|x_i - x\|^2)$. We tune hyper-parameters by conducting a grid search on regularization parameter C, the epsilon-insensitive loss ε, and the RBF kernel parameter γ. Performance is evaluated by cross-validation.

Figure 7. Support vectors

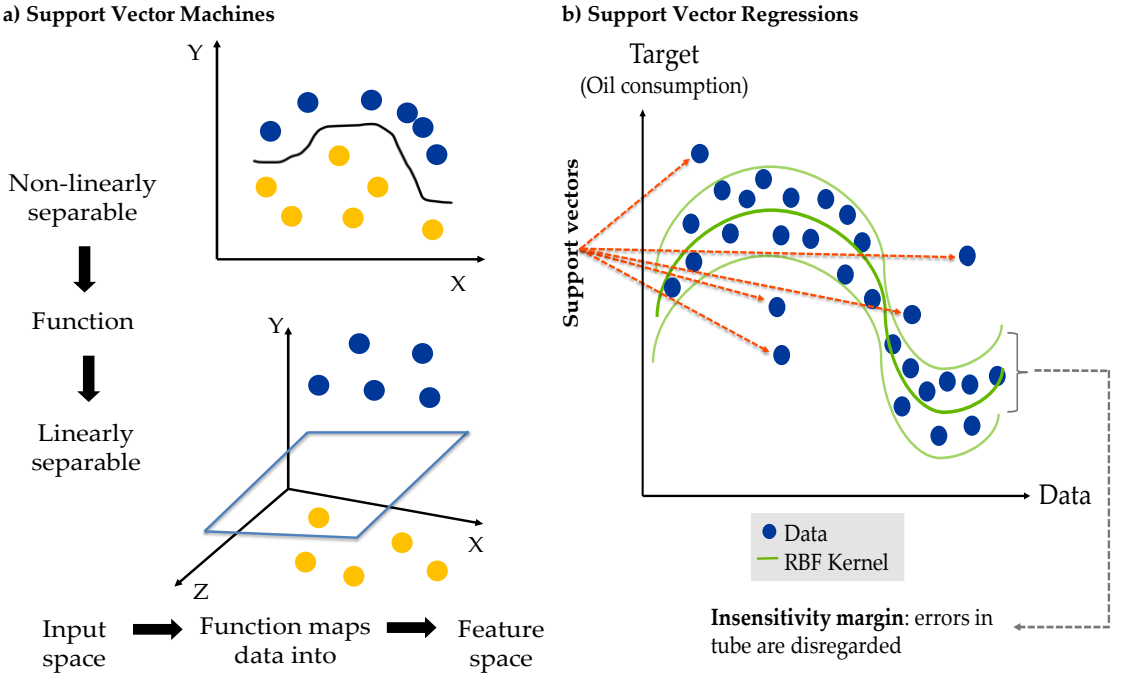

Sources: Palaogu (2023) and authors.
Note: Panel b) represents a radial basis function kernel. Other popular kernels include linear, polynomial, or sigmoid.

Our results indicate that incorporating satellite-based NO_2 pollution data improves forecast accuracy – confirming the results obtained with an OLS specification. **Figure 8** reports average RMSE gains across countries and shows that SVR models including NO_2 consistently outperform those without it. Compared with a naïve AR model (panel a), incorporating satellite-based NO_2 increases forecasting accuracy by about 3% over the full sample, with larger gains (around 5%) during the pandemic. The more sophisticated benchmark (panel b) also benefits from including NO_2 data, with average gains of 3% – and slightly stronger improvements outside the Covid period. As with the OLS results, country-level results vary: Spain records RMSE gains of up to 10%, while France and Italy exhibit more modest improvements of 2-4%. Overall, these findings confirm that the predictive gains from NO_2 extend beyond linear models, strengthening their robustness across different forecasting frameworks. Nevertheless, the larger gains observed with OLS likely reflect the model's more limited ability to capture complex relationships when predictors are omitted. By contrast, SVR already flexibly captures

non-linearities and interactions among the benchmark variables, so the marginal contribution of adding NO_2 pollution would be smaller.

Figure 8. Out-of-sample RMSE of SVR models incorporating satellite-based NO_2 pollution – relative to benchmarks

(average across countries)

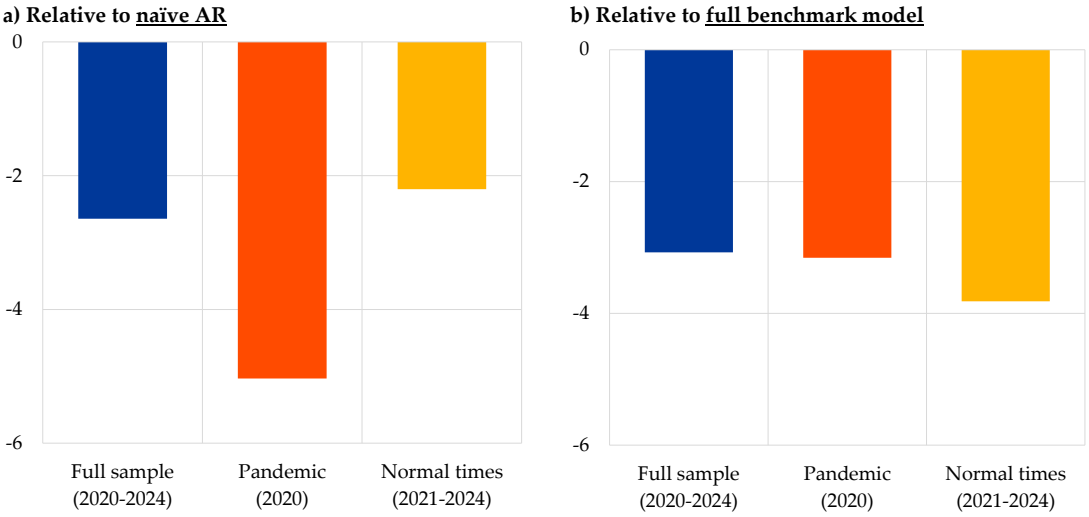

Sources: Haver Analytics, US Energy Information Administration, International Energy Agency, OECD, Bloomberg and authors.
*Notes: Countries in the sample are Australia, China, France, India, Italy, Japan, South Korea, Spain the United Kingdom and the United States. The benchmark model incorporates an AR term, a variable for industrial activity (industrial production), a variable for prices (PPI), a variable for the number of vehicles (new car registrations), and a variable for weather (temperature). Results are presented relative to the benchmark: a negative value indicates an over-performance of the model with satellite-based NO_2 pollution. Results by country can be found in **Table A6** in **Appendix A**.*

We then compare the performances of linear versus non-linear models for predicting oil consumption. **Figure 9** reports the out-of-sample RMSEs of SVR models, relative to the same specifications in OLS. On average across countries, SVRs outperform OLS by up to 55% over the full sample, confirming prior findings that machine learning models, particularly SVRs, outperform linear models in forecasting oil demand (Yu et al., 2019; Romero-Gelvez et al., 2020; Zhu, 2023). Gains are especially pronounced during the pandemic, whereas in normal times SVRs deliver a similar accuracy as OLS. This pattern is in line with the literature on asymmetric forecast performances, which point to more complex models performing better during downturns but being of comparable accuracy with naïve models in normal times (e.g. Chauvet and Potter, 2013; Siliverstovs and Wochner, 2021). At the country level, accuracy gains from SVRs are largest in China (around 75%) but are more modest in South Korea (around 3%).

Figure 9. Out-of-sample RMSE of SVR model compared to OLS

(average across countries)

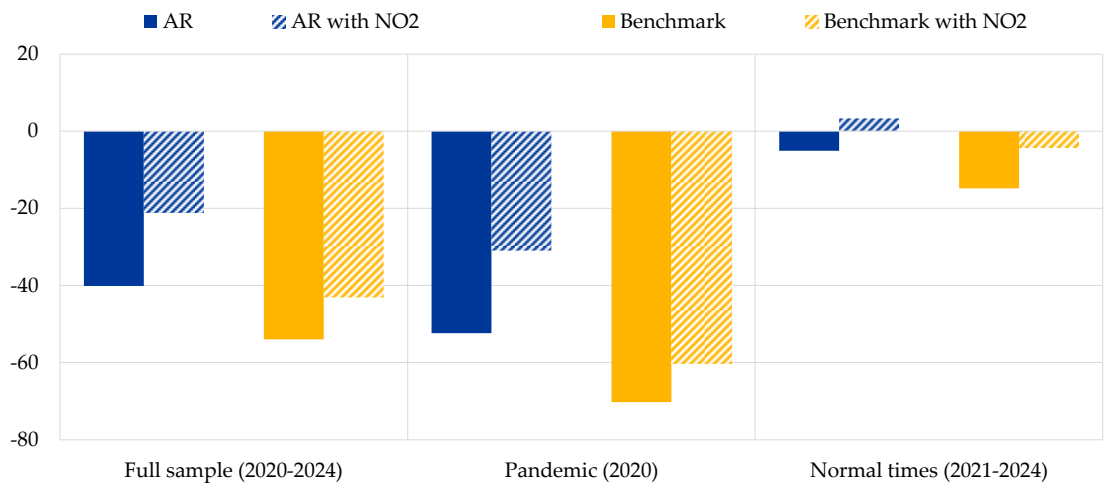

Sources: Haver Analytics, US Energy Information Administration, International Energy Agency, OECD, Bloomberg and authors.
Notes: Countries in the sample are Australia, China, France, India, Italy, Japan, South Korea, Spain the United Kingdom and the United States. The benchmark model incorporates an AR term, a variable for industrial activity (industrial production), a variable for prices (PPI), a variable for the number of vehicles (new car registrations), and a variable for weather (temperature). Results are presented relative to the OLS model: a negative value indicates an over-performance of the SVR model compared to OLS. Results by country can be found in ***Table A7*** *in* ***Appendix A****.*

In a second exploration of non-linear models, we employ Artificial Neural Networks (ANN) to nowcast oil demand. Compared with SVRs, ANNs offer greater flexibility in capturing complex, non-linear relationships, making them well suited to noisy datasets. Neural networks consist of *"layers"* with *"neurons"*: each neuron applies weights to its inputs and transforms them via a (generally non-linear) *"activation function"*, before passing the output forward. This process is run across multiple interconnected layers: information flows from the *"input layer"* (receiving raw predictors) through *"hidden layers"* to the *"output layer"* that delivers the forecasted value (**Figure 10, panel a**). This layered architecture enables ANNs to approximate complex functions by progressively extracting relevant information and non-linearities. However, the large number of parameters increases the risk of overfitting, especially for small datasets as in our paper. To mitigate this, we adopt a parsimonious network with few hidden layers, we also constrain the number of *"epochs"* (training iterations over in-sample data) and incorporate *"dropout layers"* (that randomly deactivate a subset of neurons during training). Finally, given the sensitivity of ANNs to initialization, we follow Woloszko

(2024) and average predictions across an ensemble of five neural networks, each initialized with different random seeds.[26]

Figure 10. Neural networks

a) General architecture **b) ELU activation function**

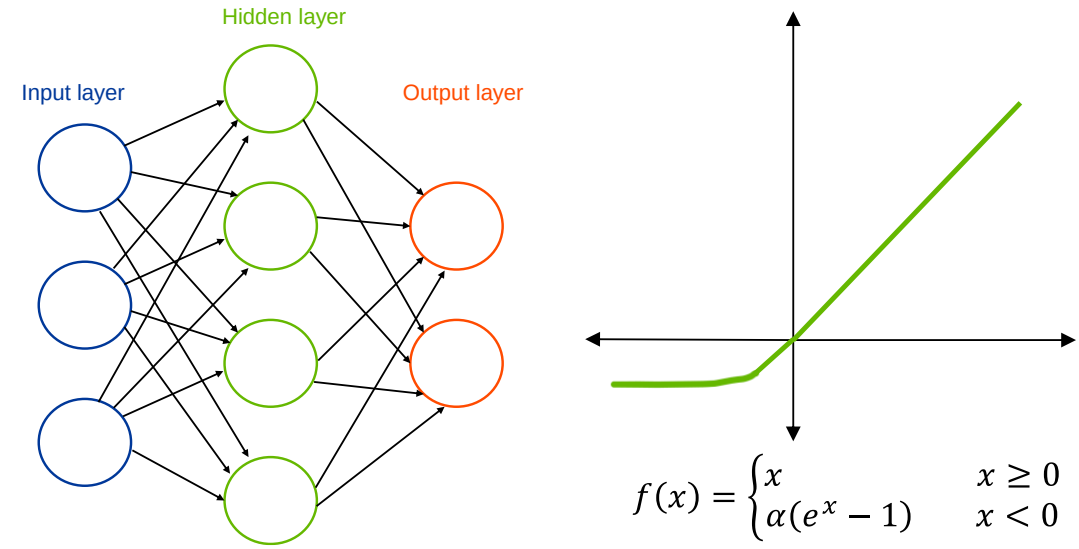

Source: authors.

In line with the results of OLS and SVR, our findings indicate that satellite-based NO_2 pollution data adds to forecast accuracy, even in the context of neural networks. While the improvements in accuracy are around 3% in crisis times, they remain more modest (below 1%) in normal periods, resembling our findings for OLS. Satellite data improves performance more significantly in the US and UK (3-3.5%), while its impact is more limited in China and Spain (see **Table A8 in Appendix A**).

Using ANNs improves accuracy compared with both OLS and SVR models, with gains of up to 80%. **Figure 11** shows the relative accuracy of ANNs to OLS and SVR models,

[26] We set hyperparameters through a grid search over the number of epochs, layers, neurons, and activation functions. In the end, we set the number of epochs to 20, number of layers to 2, number of neurons to 420 (on the first layer, then decreasing by a factor 2 on each layer), activation function to *Exponential Linear Unit* (ELU), and no L2 regularization. The ELU is linear for positive values and asymptotically approaching a bound for negative values (**Figure 10, panel b**) with this property facilitating activations closer to a zero mean, which enhances the convergence speed during training and has the potential to improve forecast accuracy. We also set batch size (number of observations used by the algorithm before updating its parameters) to 1 considering the small dataset and use the stochastic gradient descent of Adam optimizer (Kingma and Ba, 2017) as standard in the literature. Finally, we tested for other hyper-parameters whose importance was however found to be less crucial such as weight initializers (we use Xavier uniform in Pytorch) and L2 regularization.

where negative values indicate superior performance of neural networks. On average, ANNs improve accuracy by about 50% relative to OLS and 9% relative to SVR. Relative to OLS, neural networks consistently enhance forecasting performance, with gains up to 80% for China and around 30% for Spain and the United Kingdom. Relative to SVRs, ANNs also deliver accuracy gains in most countries, notably South Korea (by 32%), though SVRs slightly outperforms in the United States, Australia, and more substantially in China. These results support the hypothesis that the relationship between oil demand and satellite-based NO_2 data is non-linear, and that neural networks are particularly well-suited to capture it. This is consistent with previous studies highlighting the potential of deep learning for oil demand forecasting (Al-Fattah, 2021; Yu et al., 2019; Huang et al., 2021), and more generally in nowcasting applications (Hopp, 2022; Woloszko, 2024, d'Aspremont et al., 2025).

Figure 11. Out-of-sample RMSE of neural networks compared to OLS and SVR models

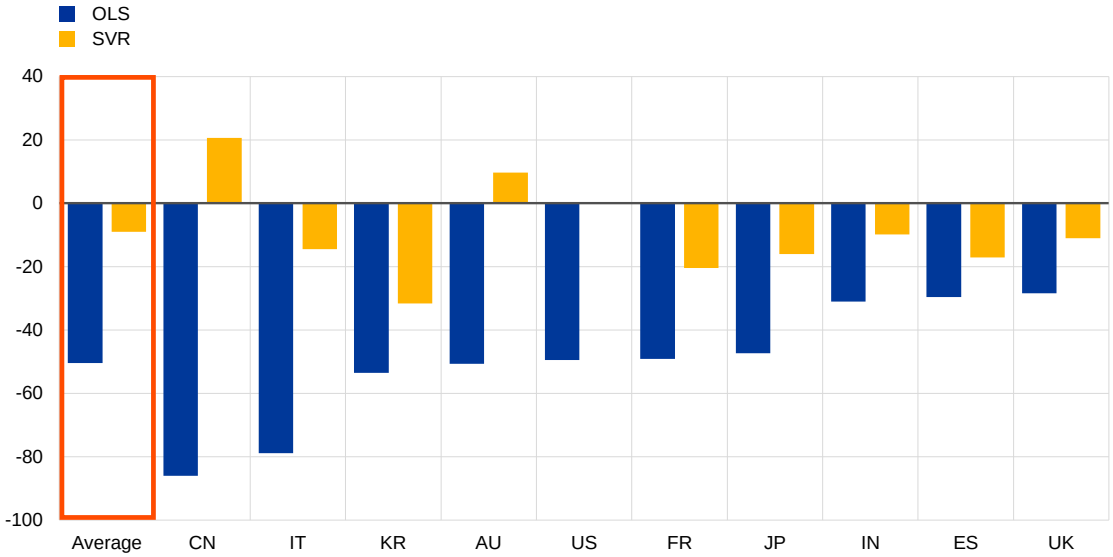

Sources: Haver Analytics, US Energy Information Administration, International Energy Agency, OECD, Bloomberg and authors.
Notes: Predictions of neural network are the average of an ensemble of 5 models initialized with different random weights. The y-axis represents the accuracy compared to the OLS and SVR set-up. A negative value indicates an over-performance of neural networks compared to OLS and SVR respectively. Numeric results by country can be found in ***Table A8*** *in* ***Appendix A****.*

Conclusion

This paper shows how satellite-based measurements of tropospheric NO_2 pollution can be transformed into an informative and timely indicator for oil demand. After retrieving, cleaning, and aggregating raw data, we construct a daily index that complements existing statistics, overcoming their publication delays and limited coverage. Our forecasting results demonstrate that incorporating NO_2 pollution significantly improves the accuracy of nowcasting models, both relative to simple autoregressive benchmarks and to richer models including more established predictors (industrial activity, energy prices, weather, and vehicle registrations). Gains are strongest during crisis periods such as the Covid-19 pandemic, but improvements remain in more stable times, highlighting the robustness of NO_2 as a real-time proxy for energy demand. The results also extend to non-linear models. Finally, our NO_2-satellite-based indicator appears informative across both advanced and emerging economies. This contrasts with night lights – a widely used satellite-based proxy for economic activity – whose relevance weakens significantly across advanced economies (Hu and Yao, 2019).

The broader implication is that satellite-based pollution data can help close information gaps in energy statistics. More generally, the findings illustrate the value of high-frequency satellite observations in economic monitoring, complementing traditional sources. Compared with traditional statistics and alternative data sources, satellite data have the key advantages of daily observations (vs. monthly for most official statistics), global coverage (vs. official or alternative indicators often limited to few countries and with heterogeneous quality depending on countries), granularity, and free access. Future research could extend this approach to other pollutants and satellite-based measures to develop a broader toolkit of real-time indicators for economic analysis. This approach might also be extended at infra-national level, given the granularity provided by satellite data.

References

Abbas, S., Ali, G., Qamer, F. M., and Irteza, S. M. (2022). Associations of air pollution concentrations and energy production dynamics in Pakistan during lockdown. *Environmental Scientific Pollution Research*, 29: 35036–35047.

Ahnert, H., and Bier, W. (2001). Trade-off between timeliness and accuracy. *Economisch Statistische Berichten*, No 4299.

Alix-Garcia, J., and Millimet, D. (2021). Remotely Incorrect? Accounting for Nonclassical Measurement in Satellite Data on Deforestation., *Unpublished manuscript.*

Aldabbagh, A.M., Economou, A., and Chariton, C. (2024). Forecasting global oil demand: Application of machine learning techniques. *The Oxford Institute for Energy Studies.*

Al-Fattah, S. M. (2021). Application of the artificial intelligence gannats model in forecasting crude oil demand for Saudi Arabia and China. *Journal of Petroleum Science and Engineering*, 200: 108368.

Alkhammash, E. H., Kamel, A. F., Al-Fattah, S. M., and Elshewey, A. M. (2022). Optimized multivariate adaptive regression splines for predicting crude oil demand in Saudi Arabia. *Discrete Dynamics in Nature and Society*.

Anh, V., Duc, H., and Azzi, M. (1997). Modelling Anthropogenic Trends in Air Quality Data. *Journal of Air & Waste Management Association*, 47: 66–71.

d'Aspremont, A., Ben Arous, S., Bricongne, J.-C., Lietti, B., and Meunier, B. (2025). Satellite turn "concrete": Tracking cement with satellite data and neural networks. *Journal of Econometrics*, 249(Part C): 105923.

Assareh, E., Behrang, M. A., Assari, M. R., and Ghanbarzadeh, A. (2010). Application of PSO (particle swarm optimization) and Ga (genetic algorithm) techniques on demand estimation of oil in Iran. *Energy*, 35(12): 5223–5229.

Baimatova, N., Bukenov, B., Ibragimova, O. P., Karaca, F., Kenessov, B., Kerimray, A., and Plotitsyn, P. (2020). Assessing air quality changes in large cities during COVID-19 lockdowns: The impacts of traffic-free urban conditions in Almaty, Kazakhstan. *Science of the Total Environment*, 730.

Bellman, R. (1957). *Dynamic Programming*. Princeton University Press.

Behrang, M.A., Assareh, E., Ghalambaz, M., Assari, M. R., and Noghrehabadi, A. R. (2011). Forecasting future oil demand in Iran using GSA (gravitational search algorithm). *Energy*, 36(9): 5649–5654.

Beyer, R., Franco-Bedoya, S., and Galdo, V. (2021). Examining the economic impact of COVID-19 in India through daily electricity consumption and nighttime light intensity. *World Development*, 140.

Bi, J., Belle, J., Wang, Y., Lyapustin, A., Wildani, A., and Liu, Y. (2019). Impacts of snow and cloud covers on satellite-derived PM2.5 levels. *Remote Sensing of Environment*, 221: 665–674.

Boamah, V. (2021). Forecasting the demand of oil in Ghana: A statistical approach. *Management Science and Business Decisions*, 1(1): 29–43.

Boersma, K., and Castellanos, P. (2012). Reductions in nitrogen oxides over Europe driven by environmental policy and economic recession. *Scientific Reports*, 2(265).

Bricongne, J.-C., Coffinet, J., Delbos, J.-B., Kaiser, V., Kien, J.-N., Kintzler, E., Lestrade, A., Meunier, B., Mouliom, M., and Nicolas, T. (2020). Tracking the economy during the Covid-19 pandemic: The contribution of high-frequency indicators. *Bulletin de la Banque de France*, 231.

Bricongne, J.-C., Meunier, B., and Pical, T. (2021). Can satellite data on air pollution predict industrial production? *Working papers*, No 847, Banque de France.

Bricongne, J.-C., Meunier, B., and Pouget, S. (2023). Web Scraping Housing Prices in Real-time: the Covid-19 Crisis in the UK, *Journal of Housing Economics*, 59(Part B): 101906.

Carvalho, V., Garcia, J., Hansen, S., Ortiz, Á., Rodrigo, T., Rodríguez Mora, J., and Ruiz, J. (2021). Tracking the COVID-19 Crisis with High-Resolution Transaction Data. *Royal Society Open Service*, 8: 210218.

Chai, J., Wang, S., Wang, S., and Guo, J. (2012). Demand forecast of petroleum product consumption in the Chinese transportation industry. *Energies*, 5(3): 577–598.

Chauvet, M., and Potter, S. (2013). Forecasting output. In Elliott, G., and Timmermann, A. (eds.), *Handbook of forecasting*, 2: 1–56. North Holland.

Chetty, R., Friedman, J., Hendren, N., Stepner, M., and the Opportunity Insights Team (2020). How Did COVID-19 and Stabilization Policies Affect Spending and Employment? A New Real-Time Economic Tracker Based on Private Sector Data. *NBER Working Paper*, No 27431, National Bureau of Economic Research.

Chi, Y., Wu, Z., Liao, K., and Ren, Y. (2020). Handling Missing Data in Large-Scale MODIS AOD Products Using a Two-Step Model. *Remote Sensing*, 12(22): 3786.

Chodorow-Reich G., Gopinath G., Mishra P., and Narayanan A. (2020). Cash and the Economy: Evidence from India's Demonetization. *The Quarterly Journal of Economics*, 135(1): 57-103.

Clark, T. E., and West, K. D. (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics*, 138(1): 291–311.

Crompton, P., and Wu, Y. (2005). Energy consumption in China: Past trends and future directions. *Energy Economics*, 27(1): 195–208.

Cui, Y., He, S., Jiang, L., Kong, H., and Zhou, H. (2020). Effects of the socio-economic influencing factors on SO2 pollution in Chinese cities: A spatial econometric analysis based on satellite observed data. *Journal of Environmental Management*, 268.

de Ruyter de Wildt, M., Eskes, H., and Boersma, K. (2012). The global economic cycle and satellite-derived NO2 trends over shipping lanes. *Geophysical Research Letters*, 39: 1–6.

Dey, A., and Das, K. (2022). How do mobility restrictions and social distancing during COVID-19 affect oil price? *Journal of Statistical Theory and Practice*, 16(22).

Dilaver, Z., and Hunt, L.C. (2011). Industrial Electricity Demand for Turkey: A structural time series analysis. *Energy Economics*, 33(3): 426–436.

Du, Y., and Xie, Z. (2017). Global financial crisis making a V-shaped fluctuation in NO2 pollution over the Yangtze River Delta. *Journal of Meteorological Research*, 31: 438-447.

Ebener, S., Murray, C., Tandon, A., and Elvidge, C. (2005). From wealth to health: modelling the distribution of income per capita at the subnational level using night-time light imagery. *International Journal of Health Geographics*, 4(5).

Ediger, V., and Akar, S. (2007). ARIMA forecasting of primary energy demand by fuel in Turkey. *Energy Policy*, 35(3): 1701–1708.

Egger, P., Rao, S.X., and Papini, S. (2023). Building floorspace in China: a dataset and learning pipeline, arXiv pre-print: 2303.02230. https://arxiv.org/abs/2303.02230.

European Spatial Agency (2020). Flawed estimates of the effects of lockdown measures on air quality derived from satellite observations. *Technical Report*

Ezran, I., Morris, S. D., Rama, M., and Riera-Crichton, D. (2023). Measuring Global Economic Activity Using Air Pollution. *Policy Research Working Papers*, No 10445, World Bank.

Fatima, T., Xia, E., and Ahad, M. (2019). Oil demand forecasting for China: A fresh evidence from structural time series analysis. *Environment, Development and Sustainability*, 21(3): 1205–1224.

Ferrara, L., Froidevaux, A., and Huynh, T.-L. (2020). Macroeconomic nowcasting in times of Covid-19 crisis: On the usefulness of alternative data. *Econbrowser.*

Fouladgar, N., and Främling, K. (2020). A Novel LSTM for Multivariate Time Series with Massive Missingness. *Sensors*, 20(10): 2832.

Fosten, J., and Nandi, S. (2025). Nowcasting U.S. state-level CO2 emissions and Energy Consumption. *International Journal of Forecasting*, 41(1): 20–30.

Griffin, D., Zhao, X., McLinden, C. A., Boersma, F., Bourassa, A., Dammers, E., Degenstein, D., Eskes, H., Fehr, L., Fioletov, V., Hayden, K., Kharol, S. K., Li, S.-M., Makar, P., Martin, R. V., Mihele, C., Mittermeier, R. L., Krotkov, N., Sneep, M., Lamsal, L. N., ter Linden, M., van Geffen, J., Veefkind, P., Wolde, M., and Zhao, X. (2019). High-Resolution Mapping of Nitrogen Dioxide With TROPOMI: First Results and Validation Over the Canadian Oil Sands. *Geophysical Research Letters*, 46: 1049–1060.

Goldblatt, R., Heilmann, K., and Vaizman, Y. (2019). Can Medium-Resolution Satellite Imagery Measure Economic Activity at Small Geographies? Evidence from Landsat in Vietnam. *The World Bank Economic Review*, 34(3): 635–653.

Hale, T., Angrist, N., Goldszmidt, R., Kira, B., Petherick, A., Phillips, T., Webster, S., Cameron-Blake, E., Hallas, L., Majumdar, S., and Tatlow, H. (2021). A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker). *Nature Human Behaviour*, 5: 528–538.

Harris, T., Devkota, J. P., Khanna, V., Eranki, P. L., and Landis, A. E. (2018). Logistic growth curve modeling of US energy production and consumption. *Renewable and Sustainable Energy Reviews*, 96: 46–57.

He, Y., and Lin, B. (2018). Forecasting China's total energy demand and its structure using ADL-MIDAS model. *Energy*, 151: 420–429.

Henderson, J., Storeygard, A., and Weil, D. (2012). Measuring Economic Growth from Outer Space. *American Economic Review*, 102(2): 994–1028.

Henneman, L., Liu, C., Hu, Y., Mulholland, J., and Russell, A. (2017). Air quality modeling for accountability research: Operational, dynamic, and diagnostic evaluation. *Atmospheric Environment*, 166: 551–565.

Hirabayashi, S., and Kroll, C. (2017). Single imputation method of missing air quality data for i-Tree Eco analyses in the conterminous United States. *unpublished manuscript*.

Hopp, D. (2022). Economic Nowcasting with Long Short-Term Memory Artificial Neural Networks (LSTM). *Journal of Official Statistics*, 38(3): 847–873.

Hu, Y., and Yao, J. (2019). Illuminating Economic Growth. *IMF Working Paper*, No 19/77, International Monetary Fund.

Hu, F., and Guo, Y. (2020). Impacts of electricity generation on air pollution: Evidence from data on Air Quality Index and six criteria pollutants. *SN Applied Sciences*, 3(1).

Huang, Y., Li, S., Wang, R., Zhao, Z., Huang, B., Wei, B., and Zhu, H. (2021). Forecasting oil demand with the development of comprehensive tourism. *Chemistry and Technology of Fuels and Oils*, 57(2): 299–310.

Hunt, L.C., Judge, G. and Ninomiya, Y. (2003). Modelling underlying energy demand trends. *Energy in a Competitive Market*.

INSEE (2010). Modélisation de la consommation ou de la production d'énergie en fonction de la température. *Note de conjuncture de l'INSEE.*

INSEE (2020). Les données « haute fréquence » sont surtout utiles à la prévision économique en période de crise brutale. *Note de conjoncture de l'INSEE.*

International Energy Agency (IEA) (2025). Global EV Outlook 2025.

Jardet, C., and Meunier, B. (2022). Nowcasting world GDP growth with high-frequency data. *Journal of Forecasting*, 41(6): 1181–1200.

Keola, S., Andersson, M., and Hall, O. (2015). Monitoring Economic Development from Space: Using Nighttime Light and Land Cover Data to Measure Economic Growth. *World Development*, 66: 322–334.

Keola, S., and Hayakawa, K. (2021). Do Lockdown Policies Reduce Economic and Social Activities? Evidence from NO2 Emissions. *The Developing Economies*, 59(2): 178–205.

Kianian, B., Liu, Y., and Chang, H. (2021). Imputing Satellite-Derived Aerosol Optical Depth Using a Multi-Resolution Spatial Model and Random Forest for PM2.5 Prediction. *Remote Sensing*, 13(1): 126.

Kingma, D., and Ba, J. (2017). *Adam: A Method for Stochastic Optimization*. arXiv pre-print: 1412.6980. https://arxiv.org/abs/1412.6980

Lamsal, L., Martin, R., Padmanabhan, A., van Donkelaar, A., Zhang, Q., Sioris, C., Chance, K., Kurosu, T., and Newchurch, M. (2011). Application of satellite observations for timely updates to global anthropogenic NOx emission inventories. *Geophysical Research Letters*, 38(5).

Laslett, G. (1994). Kriging and splines: An empirical comparison of their predictive performance in some applications. *Journal of the American Statistical Association*, 89: 391-400.

Le, T., Wang, Y., Liu, L., Yang, J., Yung, Y., Li, G., and Seinfeld, J. (2020). Unexpected air pollution with marked emission reductions during the COVID-19 outbreak in China. *Science*, 369(6504): 702–706.

Li, Y., Ahani, A., Zhan, H., Foley, K., Alshaabi, T., Linnell, K., Dodds, P. S., Danforth, C. M., and Fox, A. (2021). Blending search queries with social media data to improve forecasts of economic indicators. arXiv pre-print: 2107.06096. https://arxiv.org/abs/2107.06096.

Li, J., Wang, R., Wang, J., and Li, Y. (2018). Analysis and forecasting of the oil consumption in China based on combination models optimized by artificial intelligence algorithms. *Energy*, 144: 243–264.

Liu, J., Wang, S., Wei, N., Chen, X., Xie, H., and Wang, J. (2021). Natural Gas Consumption Forecasting: A discussion on forecasting history and future challenges. *Journal of Natural Gas Science and Engineering*, 90: 103930.

Liu, X., Moreno, B., and García, A.S. (2016). A grey neural network and input-output combined forecasting model. Primary energy consumption forecasts in Spanish economic sectors. *Energy*, 115(1): 1042–1054.

Lobell, D. B. (2013). The use of satellite data for crop yield gap analysis. *Field Crops Research*, 143(1): 56–64.

Mackay, R. M., and Probert, S. D. (1994). Modified logit-function demand model for predicting national crude-oil and natural-gas consumptions. *Applied Energy*, 49(1): 75–90.

Noor, N., Yahayn, A., Ralmi, N., and Al Bakri, A. M. (2006). The Replacement of Missing Values of Continuous Air Pollution Monitoring Data Using Mean Top Bottom Imputation Technique. *Journal of Engineering Research & Education*, 3: 96–105.

Nychka, D., Bandyopadhyay, S., Hammerling, D., Lindgren, F., and Sain, S. (2015). A multiresolution Gaussian process model for the analysis of large spatial datasets. *Journal of Computational and Graphical Statistics*, 24: 579–599.

Palaoglu, S. (2023). What is Kernel in Machine Learning? *Dida blog*.

Peng, X., Shen, H., Zhang, L., Zeng, C., Yang, G., and He, Z. (2020). Spatially continuous mapping of daily global ozone distribution (2004–2014) with the Aura OMI sensor. *Journal of Geophysical Research: Atmospheres*, 121(12): 702–722.

Pinkovskiy, M., and Sala-i-Martin, X. (2016). Lights, Camera . . . Income! Illuminating the National Accounts-Household Surveys Debate. *The Quarterly Journal of Economics*, 131(2): 579–631.

Rao, S., and Zurbenko, I. (1994). Detecting and tracking changes in ozone air quality. *Journal of Air & Waste Management Association*, 44: 1089–1092.

Rehman, S., Cai, Y., Fazal, R., Das Walasai, G., and Mirjat, N. M. (2017). An integrated modeling approach for forecasting long-term energy demand in Pakistan. *Energies*, 10(11): 1868.

Remer, L. A., Kaufman, Y. J., Tanré, D., Mattoo, S., Chu, D. A., Martins, J. V., Li, R.-R., Ichoku, C., Levy, R. C., Kleidman, R. G., Eck, T. F., Vermote, E., and Holben, B. N. (2005). The MODIS Aerosol Algorithm, Products, and Validation. *Journal of the Atmospheric Science*, 62: 947–973.

Romero-Gelvez, J. I., Villamizar, E. F., Garcia-Bedoya, O., and Herrera-Cuartas, J. A. (2020). Demand Forecasting for Inventory Management using Limited Data Sets: A Case Study from the Oil Industry. *ICAIW 2020: Workshops at the Third International Conference on Applied Informatics*.

Russell, A., Valin, L., and Cohen, R. (2012). Trends in OMI NO2 observations over the United States: effects of emission control technology and the economic recession. *Atmospheric Chemistry and Physics*, 12: 12197–12209.

Shao, Y., Ma, Z., Wang, J., and Bi, J. (2020). Estimating daily ground-level PM2.5 in China with random-forest-based spatiotemporal kriging. *Science of the Total Environment*, 740.

Siliverstovs, B., and Wochner, D. (2021). State-dependent evaluation of predictive ability. *Journal of Forecasting*, 40: 547–574.

Tadic, J., Qiu, X., Miller, S., and Michalak, A. (2017). Spatio-temporal approach to moving window block kriging of satellite data v1.0. *Geoscientific Model Development*, 10: 709–720.

Tobias, A., Carnerero, C., Reche, C., Massagué, J., Via, M., Minguillón, M. C., Alastuey, A., and Querol, X. (2020). Changes in air quality during the lockdown in Barcelona (Spain) one month into the SARS-CoV-2 epidemic. *Science of the Total Environment*, 726.

van Geffen, J., Eskes, H., Compernolle, S., Pinardi, G., Verhoelst, T., Lambert, J-C., Sneep, M., ter Linden, M., Ludewig, A., Boersma, K., and Veefkind, J. P. (2022). Sentinel-5P TROPOMI NO2 retrieval: impact of version v2.2 improvements and comparisons with OMI and ground-based data. *Atmospheric Measurement Techniques*, 15(7): 2037–2060.

Verhoelst, T., Compernolle, S., Pinardi, G., Lambert, J. C., Eskes, H. J., Eichmann, K. U., Fjæraa, A. M., Granville, J., Niemeijer, S., Cede, A., Tiefengraber, M., Hendrick, F.,

Pazmiño, A., Bais, A., Bazureau, A., Boersma, K. F., Bognar, K., Dehn, A., Donner, S., Elokhov, A., Gebetsberger, M., Goutail, F., Grutter de la Mora, M., Gruzdev, A., Gratsea, M., Hansen, G. H., Irie, H., Jepsen, N., Kanaya, Y., Karagkiozidis, D., Kivi, R., Kreher, K., Levelt, P. F., Liu, C., Müller, M., Navarro Comas, M., Piters, A. J. M., Pommereau, J-P., Portafaix, T., Prados-Roman, C., Puentedura, O., Querel, R., Remmers, J., Richter, A., Rimmer, J., Rivera Cárdenas, C., Saavedra de Miguel, L., Sinyakov, V. P., Stremme, W., Strong, K., Van Roozendael, M., Pepijn Veefkind, J., Wagner, T., Wittrock, F., Yela González, M., and Zehner, C. (2021). Ground-based validation of the Copernicus Sentinel-5p TROPOMI NO 2 measurements with the NDACC ZSL-DOAS, MAX-DOAS and Pandonia global networks. *Atmospheric Measurement Techniques*, 14(1): 481–510.

Weiss, D., Atkinson, P., Bhatt, S., Mappin, B., Hay, S., and Gething, P. (2014). An effective approach for gap-filling continental scale remotely sensed time-series. *ISPRS Journal of Photogrammetry and Remote Sensing*, 98: 106–118.

Wang, C., Wang, T., Wang, P., and Rakitin, V. (2020). Comparison and Validation of TROPOMI and OMI NO2 Observations over China. *Atmosphere*, 11(6): 636.

Wei, N., Li, C., Peng, X., Zeng, F., and Lu, X. (2019). Conventional models and artificial intelligence-based models for Energy Consumption Forecasting: A Review. *Journal of Petroleum Science and Engineering*, 181: 106187.

Woloszko, N. (2024). Nowcasting with panels and alternative data: The OECD weekly tracker. *International Journal of Forecasting,* 40(4): 1302–1335.

World Bank (2017). Growth Out of the Blue. *South Asia Economic Focus. Fall.*

Yang, J., and Hua, M. (2018). Filling the missing data gaps of daily MODIS AOD using spatiotemporal interpolation. *Science of the Total Environment*, 633.

Xiao, J., Sun, H., Hu, Y., and Xiao, Y. (2015). GMDH based auto-regressive model for China's energy consumption prediction. *2015 International Conference on Logistics, Informatics and Service Sciences (LISS)*.

Xu, G., and Wang, W. (2010). Forecasting China's natural gas consumption based on a combination model. *Journal of Natural Gas Chemistry*, 19(5): 493–496.

Yu, L., Zhao, Y., Tang, L., and Yang, Z. (2019). Online big data-driven oil consumption forecasting with Google trends. *International Journal of Forecasting*, 35(1): 213–223.

Zhang, T., Zeng, C., Gong, W., Wang, L., Sun, K., Shen, H., Zhu, Z., and Zhu, Z. (2017). Improving Spatial Coverage for Aqua MODIS AOD using NDVI-Based Multi-Temporal Regression Analysis. *Remote Sensing*, 9(4): 340.

Zhou Z., Zhou Y., and Ge X. (2018). Nitrogen Oxide Emission, Economic Growth and Urbanization in China: A Spatial Econometric Analysis. *IOP Conference Series Materials Science and Engineering*, 301.

Zhu, H. (2023). Oil demand forecasting in importing and exporting countries: AI-based analysis of endogenous and exogenous factors. *Sustainability*, 15(18): 13592.

Appendix A: Additional tables and charts

Table A1: Detailed literature review on oil demand forecasting

Reference	*Energy product*	*Country*	*Explanatory variables*	*Forecasting method*	*Time horizon*
Linear / time-series models					
Mackay and Probert (1994)	Crude oil and natural gas	UK	Share of fuel consumption to GDP, population, GDP per capita	Logit	Annual
Xu and Wang (2010)	Natural gas	China	n.a.[27]	Polynomial Curve and Moving Average Combination Projection (PCMACP) model	Annual
Dilaver and Hunt (2011)	Industrial electricity demand	Turkey	Industrial value added, real electricity prices, past industrial electricity consumption	ARDL with underlying energy demand trend (UEDT)	Annual
Rehman et al. (2017)	Electricity, oil, petroleum, natural gas, coal, LPG	Pakistan	Sector level demand of energy components, lagged values	ARIMA, Holt-Winters, Long-range energy alternatives planning system (LEAP)	Annual
Fatima et al. (2019)	Total and transport oil	China	GDP, oil price, oil reserve, underlying energy demand trend	ARDL with UEDT	Annual
Boamah (2021)	Oil components (petroleum, gasoline, distillate fuel oil, liquefied petroleum gas)	Ghana	Energy supply, oil consumption by sector	Linear regression, exponential smoothing	Annual

[27] Information not disclosed as the paper is not publicly available.

Crompton and Wu (2005)	Coal, oil, gas, and hydroelectricity	China	Real fuel price, real GDP, population	BVAR	Annual
Ediger and Akar (2007)	Coal, oil, gas, and primary energy	Turkey	n.a.[28]	ARIMA	Annual
He and Lin (2018)	Coal, oil, gas, hydroelectricity, nuclear and wind power	China	GDP, value added	Autoregressive distributed lag with mixed data sampling (ADL-MIDAS)	Annual
Harris et al. (2018)	Crude oil, natural gas, wind and solar	US	Population, starting /ending annual production rate, midpoint of growth	Multi-cycle logistic model	Annual
Chai et al. (2012)	Oil products (diesel, petroleum), electricity	China	GDP per capita, number of civilian vehicles, level of urbanisation, turnover of passengers (freight, in 100m passenger kilometres)	Bayesian linear regression with Markov chain	Annual
Machine learning models					
Behrang et al. (2011)	Oil	Iran	Population, GDP, imports / exports, net exports, number of light-duty vehicles (LDVs) / heavy-duty vehicles (HDVs)	Gravitational search algorithm (GSA), genetic algorithm (GA), particle swarm optimisation (PSO) for parameter selection; linear estimation model	Annual
Romero-Gelvez et al. (2020)	Oil	Oil exporting countries	Endogenous and exogenous factors	SVR, PSO	Annual
Al-Fattah (2021)	Crude oil	Saudi Arabia, China	GDP, population, oil prices, gas prices, transport data	Artificial neural networks (ANN), GA	Annual

[28] Information not disclosed as the paper is not publicly available.

Alkhammash et al. (2022)	Crude oil	Saudi Arabia	GDP, population, brent prices, number of LDVs /HDVs	Optimized linear regression with multivariate adaptive regression splines (LR-MARS) based on social spider optimization (SSO) algorithm; Ridge and linear regression	Annual
Zhu (2023)	Oil	Argentina, Chile, and Brazil, United States, Korea, Malaysia, the Philippines, Singapore, China, India, Japan, Australia, New Zealand, Turkey, South Africa., Bahrain, Kuwait, Oman, Qatar, Saudi Arabia, the United Arab	Carbon emissions, income level, price level, energy price, GDP, population growth, urbanisation, trade liberalisation, inflation, foreign direct investment (FDI), financial development, Covid-19 dummies, energy sanction indicators	ARIMA, generalised linear model (GLM), SVR	Annual

		Emirates, Iraq, Iran, Algeria, Canada, Indonesia, Mexico, Nigeria, England, and Norway.			
Aldabbagh et al. (2024)	Total oil, LPG, Naptha, Gasoline, Jet fuel, Diesel, fuel oil	OECD (US, other Americas, Europe, APAC), non-OECD (China, India, other non-OECD), global	Global oil and commodity prices, population, GDP, CPI, disposable income, world trade index, China specific index, industrial production, PPI, investment, gross output, global industry indicator, stringency index, air passenger index	Gradient boosting, neural networks	Annual
Yu et al. (2019)	Oil	Global	Google trends for "oil consumption", "oil price", "oil inventory"	Linear and logistic regression, back propagation neural networks (BPNN), SVR, decision trees, extreme learning machine	Monthly
Assareh et al. (2010)	Oil	Iran	Population, GDP, import / export	Exponential regression with PSO	Annual
Liu et al. (2016)	Coal, crude oil, natural gas, renewable and nuclear energy	Spain	GDP	Grey forecasting with BPNN	Annual

Li et al. (2018)	Oil	China	Crude oil, diesel, heating oil	Quadratic exponential smoothing regression, generalised nonlinear model (with GA)	Annual
Huang et al. (2021)	Oil consumption in tourism industry	China	Variables relevant for tourism e.g. tourism revenue, total tourism expenditure factors	ANN, Layer Recurrent Neural Network	Annual

Table A2: *Data sources*

Variables	*Sources*
NO_2 pollution	QuantCube based on ESA satellite data
Oil demand	US Energy Information Administration, International Energy Agency, OECD, Bloomberg, national sources via Haver Analytics
PMI	S&P Global via Haver Analytics
CPI and PPI inflation	National statistical offices via Haver Analytics
Industrial production	National statistical offices via Haver Analytics
Heating Degree Days (HDD), Cooling Degree Days (CDD), temperature, humidity and precipitation	International Energy Agency, national sources via Haver Analytics, World Bank
New car registrations	National sources via Haver Analytics
Oil prices	Bloomberg
Google mobility	Google
Google Trends	Google
Oxford Stringency indices	Hale et al. (2021)

Table A3. Relative RMSE (out-of-sample) of models using satellite-based NO_2 pollution compared with benchmarks

	NO_2 model vs. AR model	*NO_2 model vs. Full benchmark*
Australia	-39.1% **	-28.5% *
China	-60.7% ***	-29.4% **
France	-1.2%	-14.0%*
India	-22.3% **	-16.2% *
Italy	-50.6% ***	-15.3%
Japan	-6.6%	-1.2%
South Korea	3.1%	-14.9%
Spain	-54.6% ***	-53.5% ***
United Kingdom	-1.0%*	-3.6%**
United States	-28.8% **	-36.3% ***
Average	**-26.2%**	**-22.4%**

*Notes: The benchmark model incorporates an AR term, a variable for industrial activity (industrial production), a variable for prices (PPI), a variable for the number of vehicles (new car registrations), and a variable for weather (temperature). Results are presented relative to the benchmark: a negative value indicates an over-performance of the model with satellite-based NO_2 pollution. ***, **, and * indicate that the outperformance in predictive accuracy is significant at respectively the 1%, 5%, and 10% levels, based on a Clark and West (2007) test. Test results are not available for the average.*

Table A4. Relative RMSE (out-of-sample) of models using satellite-based NO_2 pollution compared with benchmarks

	NO_2 model vs. AR model		*NO_2 model vs. Full benchmark*	
	Pandemic (2020)	*Normal times (2021-24)*	*Pandemic (2020)*	*Normal times (2021-24)*
Australia	-45.1% ***	-7.2%	-30.7% **	-12.2%
China	-65.1% ***	-11.1% *	-30.3% *	-20.2% *
France	-0.7%	-16.0%	-2.0%	-5.9%
India	-21.6% **	-22.0% *	-26.8% **	-33.2% **
Italy	-59.6% ***	-19.0% **	-9.4%	-20.3% *
Japan	-12.6%	-5.4%	3.1%	-3.3%
South Korea	-6.7%	-19.5% *	5.9%	3.0%
Spain	-59.2% ***	-58.5% ***	-23.2% *	-32.4% **
United Kingdom	-1.7%	-4.8%	1.3%	-2.0%
United States	-37.8% **	-39.2% **	-8.0%	-17.2%
Average	**-31.0%**	**-7.7%**	**-24.6%**	**-14.4%**

*Notes: The benchmark model incorporates an AR term, a variable for industrial activity (industrial production), a variable for prices (PPI), a variable for the number of vehicles (new car registrations), and a variable for weather (temperature). Results are presented relative to the benchmark: a negative value indicates an over-performance of the model with satellite-based NO_2 pollution. ***, **, and * indicate that the outperformance in predictive accuracy is significant at respectively the 1%, 5%, and 10% levels, based on a Clark and West (2007) test. Test results are not available for the average.*

Table A5. Relative RMSE (out-of-sample) of models using satellite-based NO_2 pollution compared with benchmarks using Google indices

	NO_2 model vs. Google mobility indices	*NO_2 model vs. Google Trends*	*NO_2 model vs. Oxford Stringency index*
Australia	0.6%	-24.6% *	8.9%
China	N.A.	N.A.	51.3%
France	25.2%	19.4%	-7.9%
India	-3.5%	-49.7% ***	4.9%
Italy	20.6%	-10.8%	22.7%
Japan	-35.6% **	5.0%	-7.8%
South Korea	-16.1%	-15.3%	-94.0% ***
Spain	-21.9% *	-4.5%	-22.2% *
United Kingdom	-26.9% *	-7.7%	-13.3%
United States	48.4%	-26.4% *	43.6%
Average	**-1.0%**	**-12.7%**	**-1.4%**

*Notes: China is excluded as Google indices are not available. The benchmark model incorporates an AR term, a variable for industrial activity (industrial production), a variable for prices (PPI), a variable for the number of vehicles (new car registrations), and a variable for weather (temperature). Results are presented relative to the benchmark: a negative value indicates an outperformance of the model with satellite-based NO2 pollution. ***, **, and * indicate that the outperformance in predictive accuracy is significant at respectively the 1%, 5%, and 10% levels, based on a Clark and West (2007) test. Test results are not available for the average.*

Table A6. Relative RMSE (out-of-sample) of SVR models using satellite-based NO_2 pollution compared with benchmarks

	NO_2 model vs. AR model	*NO_2 model vs. Full benchmark*
Australia	-4.1%***	-8.6%**
China	-0.9%*	-3.9%***
France	-1.7%*	-1.1%
India	-0.3%	0.5%
Italy	-4.1%***	-3.0%**
Japan	-0.1%	-2.9%**
South Korea	2.4%	2.3%
Spain	-10.1%***	-9.0%***
United Kingdom	1.2%	-0.0%
United States	-8.6%***	-4.9%***
Average	**-2.6%**	**-3.1%**

*Notes: The benchmark model incorporates an AR term, a variable for industrial activity (industrial production), a variable for prices (PPI), a variable for the number of vehicles (new car registrations), and a variable for weather (temperature). Results are presented relative to the benchmark: a negative value indicates an over-performance of the model with satellite-based NO_2 pollution. ***, **, and * indicate that the outperformance in predictive accuracy is significant at respectively the 1%, 5%, and 10% levels, based on a Clark and West (2007) test. Test results are not available for the average.*

Table A7. Relative RMSE (out-of-sample) of SVR models using satellite-based NO_2 pollution compared with OLS

	AR model	*AR model with NO_2*	*Benchmark model*	*Benchmark model with NO_2*
Australia	-56.1%	-30.9%	-64.8%	-55.0%
China	-73.9%	-34.4%	-91.5%	-88.4%
France	-17.0%*	17.4%*	-44.4%**	-36.1%**
India	-48.0%*	-33.3%	-36.2%**	-23.6%**
Italy	-51.8%	-6.6%*	-78.4%	-75.2%
Japan	-19.7%	-14.1%	-36.1%*	-37.2%**
South Korea	3.1%	2.4%	-43.4%	-31.9%*
Spain	-58.7%	-18.3%	-56.5%	-15.0%*
United Kingdom	-43.2%	-41.9%	-22.4%**	-19.5%**
United States	-35.8%*	-17.7%*	-66.2%	-49.5%
Average	**-40.1%**	**-21.2%**	**-54.0%**	**-43.1%**

*Notes: The benchmark model incorporates an AR term, a variable for industrial activity (industrial production), a variable for prices (PPI), a variable for the number of vehicles (new car registrations), and a variable for weather (temperature). Results are presented relative to the benchmark: a negative value indicates an over-performance of SVRs compared to OLS. ***, **, and * indicate that the outperformance in predictive accuracy is significant at respectively the 1%, 5%, and 10% levels, based on a Clark and West (2007) test. Test results are not available for the average.*

Table A8. Relative RMSE (out-of-sample) of neural networks using satellite-based NO_2 pollution compared with benchmark model without NO_2, OLS and SVR models

	ANN with NO_2 vs. ANN without NO_2	*ANN with NO_2 vs. OLS with NO_2*	*ANN with NO_2 vs. SVR with NO_2*
Australia	-0.6%	-50.7%*	-9.7%***
China	-0.09%	-86.0%	20.7%*
France	-2.6%*	-49.2%**	-20.4%***
India	-1.9%	-32.1%***	-9.8%*
Italy	-5.2%*	-78.8%	-14.5%***
Japan	0.8%	-47.4%**	-16.1%**
South Korea	-0.7%	-53.5%**	-31.6%***
Spain	0.2	-29.6%**	-17.2%***
United Kingdom	-3.2%**	-28.4%***	-11.0%**
United States	-3.4%**	-49.5%*	0.0%**
Average	**-1.7%**	**-50.4%**	**-9.0%**

*Notes: The benchmark model incorporates an AR term, a variable for industrial activity (industrial production), a variable for prices (PPI), a variable for the number of vehicles (new car registrations), a variable for weather (temperature) with satellite-based NO_2. Results are presented relative to the benchmark without NO_2, OLS and SVR models: a negative value indicates an over-performance of the neural network with NO_2. ***, **, and * indicate that the outperformance in predictive accuracy is significant at respectively the 1%, 5%, and 10% levels, based on a Clark and West (2007) test. Test results are not available for the average.*

Figure A1. NO_2 levels above major cities (January 20, 2025)

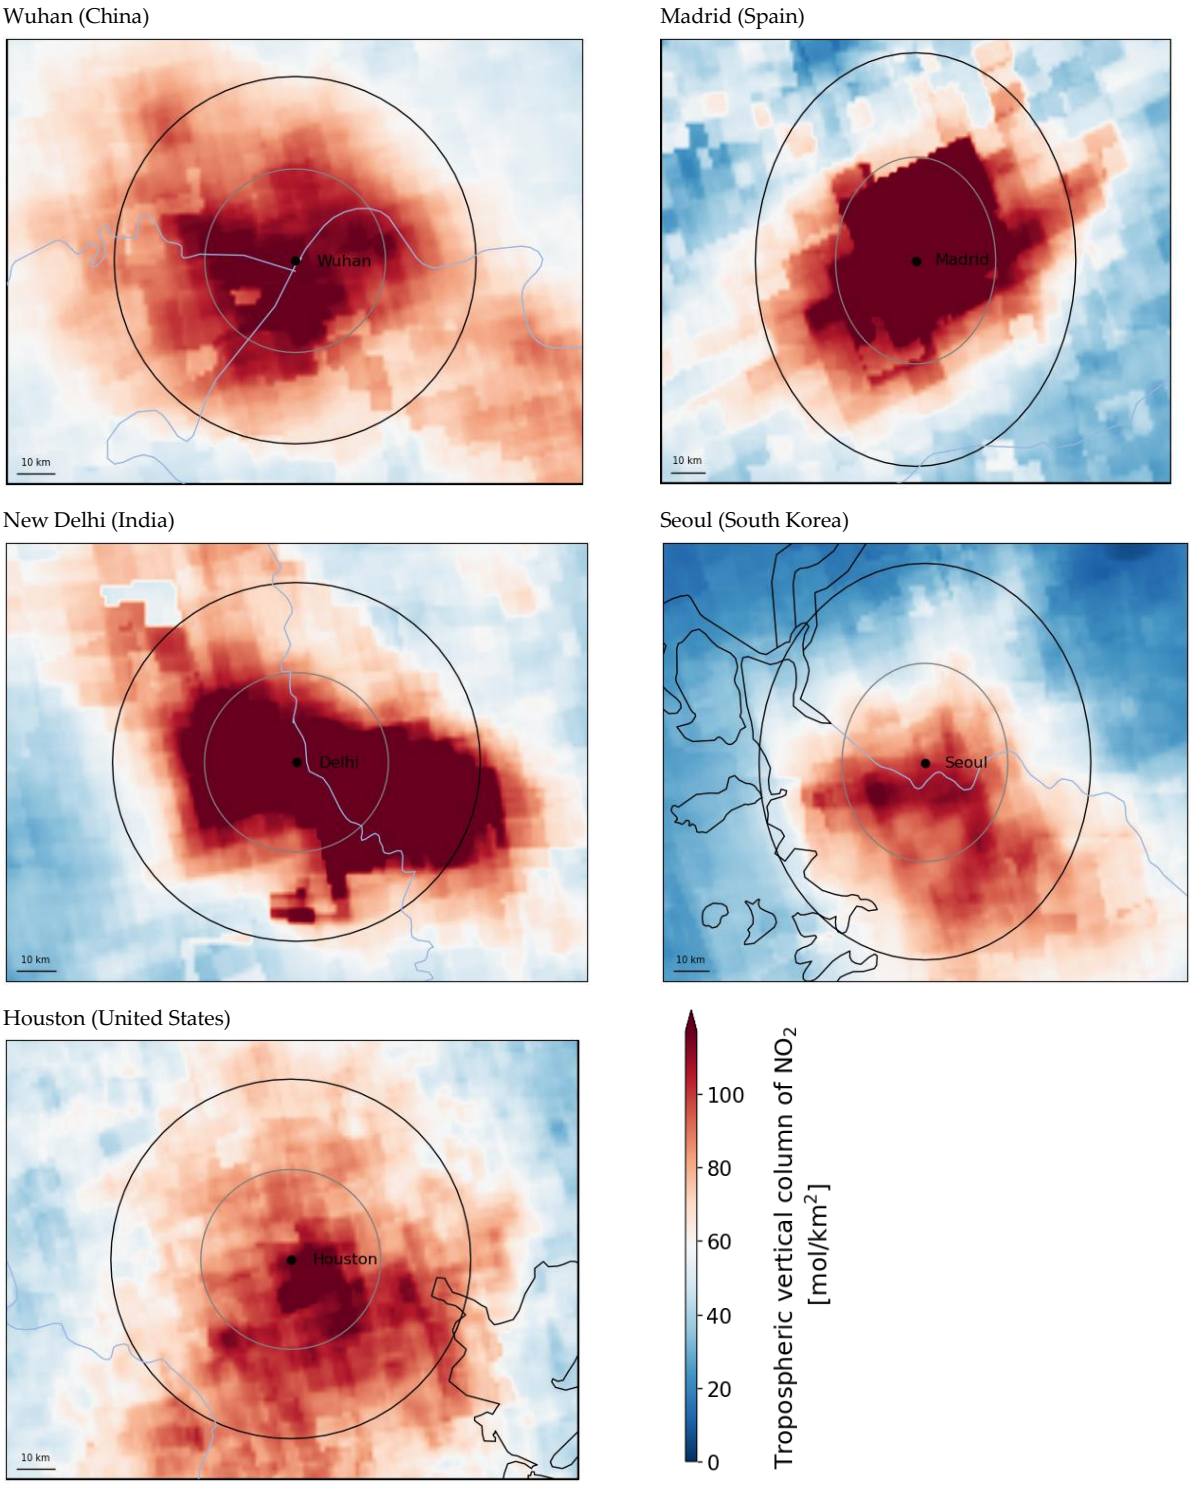

Sources: European Spatial Agency and authors.
Notes: Grey and black rings represent 20- and 40-km radius, respectively, around city centre. Rings can be egg-shaped in some cities due to distortions caused by higher latitudes. All images follow the same pre- and post-processing outlined in the article.

Figure A2. Out-of-sample RMSE of model with satellite-based NO_2 pollution – alternative specifications

(average across countries)

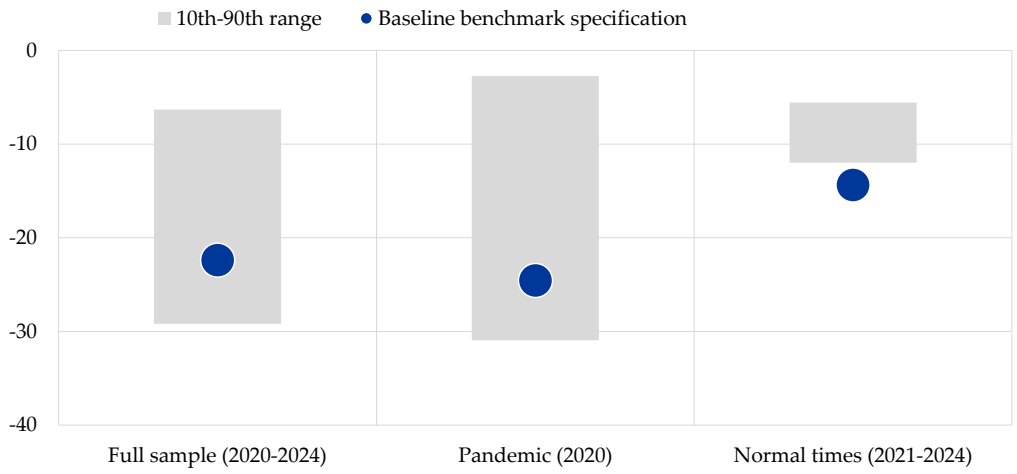

Sources: Haver Analytics, US Energy Information Administration, International Energy Agency, OECD, Bloomberg and authors.
Notes: Countries in the sample are Australia, China, France, India, Italy, Japan, South Korea, Spain the United Kingdom and the United States. The "baseline benchmark specification" model incorporates an AR term, a variable for industrial activity (industrial production), a variable for prices (PPI), a variable for the number of vehicles (new car registrations), a variable for weather (temperature) with satellite-based NO_2. Alternative specifications include other variables for industrial activity (PMI manufacturing), prices (CPI and oil prices), and weather conditions (heating degree days, cooling degree days, humidity and precipitation). The grey bar represents the 10^{th}-90^{th} range of the average relative performance of models with satellite-based NO_2 relative to models without satellite-based NO_2, across specifications incorporating all combinations of possible input variables.

Figure A3. Out-of-sample RMSE of model with satellite-based NO_2 pollution – *daily* nowcasts relative to naïve AR model

(average across countries)

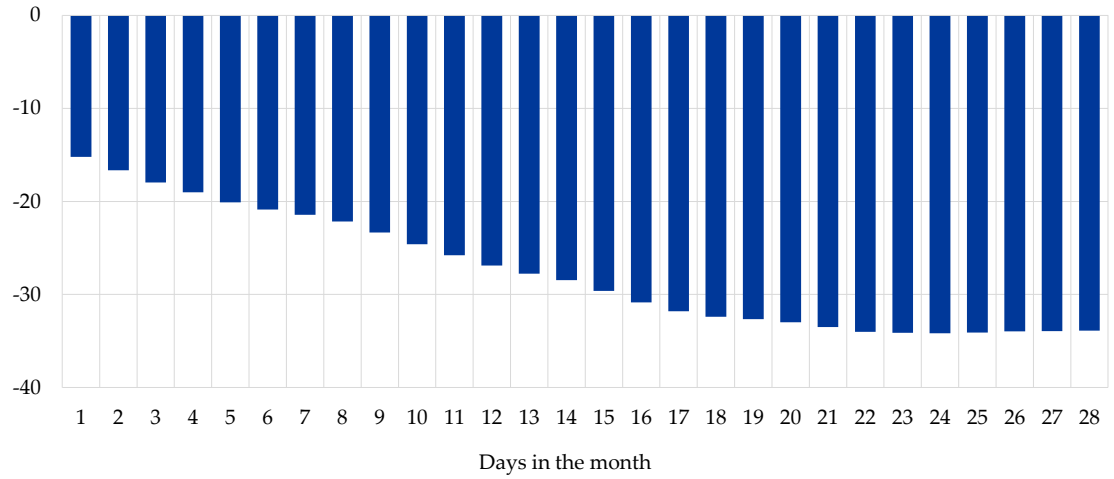

Sources: Haver Analytics, US Energy Information Administration, International Energy Agency, OECD, Bloomberg and authors.
Note: Countries in the sample are Australia, China, France, India, Italy, Japan, South Korea, Spain the United Kingdom and the United States.

Acknowledgements

We are very grateful to S. Laurent, R. Correa, E. Dumistrescu, E. Flachaire, E. Gallic, J. Marcucci, V. Mignon, N. Woloszko, as well as participants to the FRB–BdI 2020 and 2021 conferences on non-traditional data & statistical learning, the OECD-BdF conference on new approaches to macroeconomic forecasting, the BdF seminar, the BCEAO-BdF seminar on forecasting during crisis episodes, and DNB conference for useful comments.
The views expressed in this paper are those of the authors, and not those of the Banque de France, the European Central Bank, the Laboratoire d'Économie d'Orléans, Paris I University, QuantCube, or Equancy.

Jean-Charles Bricongne
Banque de France, Paris, France; Paris 1 Panthéon-Sorbonne Université, Paris, France; Laboratoire d'Economie d'Orléans (LEO), Orléans, France; email: jean-charles.bricongne@banque-france.fr

Joao Macalos
QuantCube Technology, Paris, France; email: j.macalos@quant-cube.com

Baptiste Meunier
European Central Bank, Frankfurt am Main, Germany; email: baptiste.meunier@ecb.europa.eu

Julia Milis
European Central Bank, Frankfurt am Main, Germany; email: julia.milis@ecb.europa.eu

Thomas Pical
Equancy, Paris, France; email: tpical@equancy.com

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.ecb.europa.eu

PDF ISBN 978-92-899-7726-5 ISSN 1725-2806 doi:10.2866/8144275 QB-01-26-086-EN-N