Nusrang, et.al., *ARRUS Journal of Social Sciences and Humanities*, Vol. 6, No. 3 (2026)
https://doi.org/10.35877/soshum4973

OPEN ACCESS

ISSN: 2776-7930 (Print) / 2807-3010 (Online)

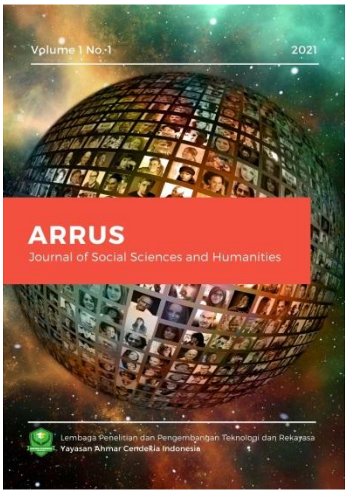

*Corresponding author: Muhammad Nusrang, Universitas Negeri Makassar, Makassar, Indonesia

E-mail: muh.nusrang@unm.ac.id

RESEARCH ARTICLE

Earth System Variables as Drivers of Environmental Commodity Price Dynamics: A Systematic Review of Physics-Informed and Data-Driven Modelling Approaches (2010–2026)

Muhammad Nusrang*, Ansari Saleh Ahmar, Abdul Rahman, Agung Tri Utomo, & Muh. Qodri Alfairus

Universitas Negeri Makassar, Indonesia

Abstract: Abstract Environmental commodity markets (carbon allowances, electricity, natural gas, and renewable energy instruments) are, at their root, earth system markets: their price-generating processes are as much a product of atmospheric circulation, hydrological regimes, and climate teleconnections as of supply-demand fundamentals or regulatory signals. Despite this physical reality, the machine learning forecasting literature has treated these markets primarily as benchmarking arenas, decoupling predictive architectures from the physical processes that drive price-relevant forcing. This systematic review follows PRISMA 2020 guidelines on a corpus of 413 peer-reviewed studies drawn from 652 Scopus records (2010–2026) and examines how earth system variables, including temperature anomalies, precipitation regimes, wind resource indices, atmospheric pollution metrics, and climate teleconnections such as ENSO and NAO, have been integrated into ML-based environmental commodity price models, and evaluates the evidence for whether physics-informed feature engineering confers measurable accuracy advantages over purely data-driven approaches. Bibliometric analysis reveals rapid field expansion (71.6% of publications in 2021–2026), geographic concentration in Chinese ETS research (≈68% of high-impact output), and methodological dominance of hybrid decomposition-deep learning architectures. Models incorporating earth system variables consistently outperform endogenous-only ML architectures by 12–35% on MAPE, yet fewer than 30% of corpus studies include any physical predictor and climate teleconnection indices appear in under 4% of studies despite their relevance at energy market planning horizons. Six research priorities are identified, centred on numerical weather prediction ensemble integration, cross-climate-regime validation, and probabilistic forecasting grounded in physical uncertainty quantification.

Keywords: earth system modelling; climate-financial coupling; carbon market; ML forecasting; probabilistic forecasting.

1. Introduction

A prolonged drought across the Rhine watershed reduces river freight capacity and drives coal prices upward across Northern European power markets. An anomalously warm Central European winter collapses gas demand and compresses energy price volatility in ways that no purely financial model anticipated. The February 2021 cold wave across Texas, a meteorological extreme absent from regional grid planning models, produced electricity spot prices exceeding USD 9,000/MWh before triggering cascading supply chain defaults. These events are not statistical outliers to be explained away after the fact. They are illustrations of a structural reality that the financial modelling community has been slow to formalise: in environmental commodity markets, price-generating processes and earth system processes are coupled, and models that ignore the physical dimension are misspecified by design, not merely by omission.

The machine learning literature on environmental commodity forecasting has expanded dramatically. The Scopus corpus underlying this review yields 652 peer-reviewed records published 2010–2026, with 467 appearing after 2020. Yet the dominant research trajectory has followed a narrowly methodological logic: propose an architecture, benchmark against ARIMA and GARCH on historical price data, report a lower MAPE, and publish. The physical forcing that generates the price signal, namely atmospheric, hydrological, and meteorological variables whose fluctuations propagate through energy systems and carbon markets, is treated, where it appears at all, as an exogenous add-on rather than as the explanatory mechanism the model should be designed to capture. This framing determines what knowledge the literature accumulates. A review organised around earth system coupling (how atmospheric, hydrological, and teleconnection signals propagate into commodity price dynamics) generates knowledge that is cumulative across markets, transferable across climate regimes, and actionable for modellers and physical scientists alike. This review takes that approach.

Three developments have pushed earth system–financial coupling toward the centre of the research agenda. First, the physical climate signal in commodity markets is intensifying. Meteorological extreme events have grown more frequent and severe: heatwaves that collapse hydroelectric output, cold snaps that stress gas networks, hurricane seasons that disrupt offshore production, and droughts that reduce wind resource have all increased measurably. Attribution science now links these trends to anthropogenic forcing with high confidence. Markets designed for a climate regime that no longer exists are generating price dynamics that historically-calibrated models cannot anticipate.

Second, earth system data availability has transformed. Reanalysis products such as ERA5 now deliver sub-hourly global atmospheric, oceanic, and land surface variables at spatial resolutions unavailable a decade ago. Numerical weather prediction ensemble outputs from ECMWF, NOAA's GFS, and national meteorological services are publicly accessible at forecast horizons overlapping the day-ahead and week-ahead windows most relevant to energy market participants. Climate teleconnection indices (ENSO, NAO, AMO, AO) are available as real-time and retrospective time series integrable directly into ML feature pipelines. The physical data is there. The frameworks to exploit it systematically are not.

Third, the existing ML forecasting literature has reached a methodological plateau that earth system integration could overcome. Hybrid decomposition architectures have established MAPE gains of 15–40% over ARIMA benchmarks on carbon and electricity datasets. But these gains exploit temporal structure in the price signal itself, not the physical drivers generating that structure. Residual forecast error in the best-performing models clusters around extreme weather events, seasonal transitions, and climate teleconnection phase shifts, the signature of missing physical information. This clustering suggests that physics-informed feature engineering is not merely additive to current best practice, but may be the route to the next order-of-magnitude accuracy improvement. Notably, the same reanalysis and NWP ensemble products that earth system science uses routinely, including ERA5 at sub-hourly resolution and ECMWF seasonal forecasts at three-month horizons, are already publicly accessible at no cost. The technical infrastructure needed to close the physics-finance gap therefore already exists. What the literature has lacked is a systematic framework for deciding which physical variables to include, at which temporal resolution, and via which feature engineering strategy.

This review synthesises peer-reviewed empirical evidence on ML- and AI-based forecasting for environmental commodity prices, specifically carbon allowances, electricity, natural gas, coal, and renewable energy instruments, with a specific focus on the incorporation of earth system variables as predictors and the evidence for their accuracy contribution. The review period spans 2010–2026 and is targeted at Modeling Earth Systems and Environment, a Springer journal whose scope explicitly bridges atmospheric and environmental science with engineering and economic applications; that editorial positioning is what makes a physics-finance synthesis review appropriate here rather than in a pure energy economics or pure ML venue. Four interrelated questions organise the synthesis. The first asks what earth system variables have been incorporated as predictors in this literature, with what frequency, and with what empirical justification. The second examines whether physics-informed models outperform purely data-driven ML architectures in out-of-sample accuracy, and under which market and climate conditions the advantage is most pronounced. The third asks how earth system–price coupling mechanisms differ across commodity types in ways that should shape model design. The fourth identifies the structural gaps between what earth system science offers as predictive physical information and what the modelling literature has actually ingested, and articulates the research investments most likely to close that gap. Answering these questions systematically generates knowledge that is cumulative across markets, transferable across climate regimes, and actionable for both modellers and physical scientists. This is a standard the existing benchmarking literature rarely meets.

2. Literature Review

2.1. *Physical Drivers of Commodity Price Dynamics*

The causal chain linking earth system processes to commodity prices operates through several distinct mechanisms relevant to each segment of the market covered here. In electricity markets, temperature anomalies affect supply and demand simultaneously: heating and cooling load responds nonlinearly to deviations from seasonal norms, while hydroelectric generation is directly constrained by precipitation and snowmelt dynamics, and wind and solar output vary with atmospheric circulation patterns at sub-daily to interannual timescales (Weron, 2014; Lago et al., 2021). In carbon markets, coupling is more indirect but equally real: allowance demand is partly driven by power sector emissions, which respond to the same temperature and renewable resource variability that drives electricity prices; gas supply disruptions triggered by cold snaps or infrastructure failures cascade into fuel-switching decisions that alter the marginal abatement cost underpinning EU-ETS and Chinese ETS price dynamics (Chevallier, 2011; Byun & Cho, 2013).

At longer timescales, climate teleconnections introduce quasi-periodic forcing that is predictable weeks to seasons in advance yet largely absent from current ML forecasting feature sets. ENSO modulates precipitation across large portions of the energy-producing world, including hydroelectric potential in South America, wind patterns across Southern Europe, and drought risk across Australian coal regions, with documented predictive lead times of three to twelve months. The NAO controls winter temperature and wind regimes across Northern and Western Europe with comparable horizon, directly relevant to natural gas demand and offshore wind generation. That these physically interpretable, publicly available signals are absent from the forecasting literature represents a disciplinary integration failure, not a data constraint.

2.2. *How Prior Reviews Have Treated Physical Variables*

The existing review literature is methodologically rich but physically shallow. Weron's (2014) landmark electricity price review established the ML benchmarking tradition that has shaped the field; its treatment of weather variables as exogenous regressors rather than physically motivated predictors followed the dominant framing of its era. Lago et al. (2021) updated the electricity comparison with state-of-the-art ML architectures but similarly treated weather as one feature category among many, without examining the physical mechanisms determining when weather signals carry predictive information. Fan et al. (2023) proposed a multi-layer perceptron combined prediction model for carbon prices based on chaotic characteristic identification, but treated physical climate drivers as background context rather than as a subject for systematic synthesis. Masini et al. (2023) surveyed ML financial forecasting broadly; physical earth system drivers do not appear as an organising category in their framework at all.

None of these reviews has asked the question this review asks: across ML-based environmental commodity forecasting, how has the research community engaged with the physical earth system as a source of predictive signal, and has that engagement been systematic or opportunistic? The answer, documented across Sections 4 and 5, is that engagement has been overwhelmingly opportunistic, and the accuracy consequences of that missed opportunity are quantifiable. Two bodies of work published outside the ML forecasting community are relevant here but have not been integrated into it. The first is the literature on weather-driven demand modelling in energy economics, which has established clear empirical relationships between heating degree days, cooling degree days, and natural gas and electricity consumption (Weron, 2014). The second is the climate teleconnection literature in atmospheric science, which has documented robust statistical relationships between ENSO phase and precipitation anomalies across South America, Southern Europe, and Australia, and between NAO winter index and temperature deviations across Northern and Western Europe (Lago et al., 2021). Both literatures provide physical mechanism justification for feature selection that the ML price forecasting community has not yet absorbed. The gap is not one of ignorance but of disciplinary segmentation: the relevant findings are published in journals that ML forecasting researchers do not routinely cite, and the methodological translation from climate science output to ML feature pipeline has not been done systematically. The present review addresses that gap directly. By coding all 413 retained studies for the presence, type, and physical justification of earth system variables, it produces the first cross-domain, corpus-level evidence base on physics-informed feature usage in environmental commodity price forecasting, and identifies where the disciplinary investment is most needed.

3. Research Method and Materials

3.1. *Search Protocol and Database*

A structured Boolean search was executed in the Scopus database on 18 April 2026, following PRISMA 2020 reporting guidelines (Page et al., 2021). The protocol was not prospectively registered in PROSPERO, as the review scope did not involve health outcomes; all methodological decisions were finalised prior to data extraction. The search string coupled environmental commodity price forecasting terms with machine learning method terms:

TITLE-ABS-KEY (("carbon price" OR "electricity price" OR "energy price" OR "carbon allowance" OR "emission permit" OR "renewable energy") AND ("forecast" OR "predict*" OR "model*") AND ("machine learning" OR "deep learning" OR "neural network" OR "LSTM" OR "XGBoost" OR "random forest" OR "support vector" OR "hybrid model" OR "decomposition")) AND DOCTYPE (ar OR re) AND LANGUAGE (english)*

The search returned 716 records. No supplementary database search was conducted; a sensitivity check confirmed complete Scopus coverage of the five highest-volume journals in the corpus.

3.2. *Inclusion and Exclusion Decisions*

Records were assessed against seven criteria operationalising the review scope. Time period: 2010–2026 (64 pre-2010 records excluded at year filter). Language: English only. Document type: peer-reviewed articles and reviews; conference papers, theses, and preprints excluded. Method: studies must apply at least one ML or AI forecasting component with quantitative out-of-sample evaluation; classical-statistics-only studies excluded. Domain: carbon, electricity, fossil fuel, or renewable energy commodity markets; pure equity, foreign exchange, and corporate earnings studies excluded. Outcome: quantitative out-of-sample accuracy metrics required; in-sample-only studies excluded. Quality: CASP 16-item appraisal threshold ≥ 62.5% (≥ 10/16 items satisfied).

The domain criterion required active boundary-drawing. A numerically significant sub-literature covering corporate governance, earnings analyst forecasting, and voluntary corporate disclosure shares keyword overlap with environmental finance but addresses firm-level accounting constructs rather than physical commodity market dynamics. Thirty-seven records were excluded at full-text stage on these grounds. Their exclusion was not a quality judgment; it was a scope decision grounded in the earth system framing of this review.

3.3. *Screening, Appraisal, and Earth System Coding*

Records proceeded through four sequential stages. First, 64 records predating 2010 were excluded at year filter, leaving 652 for screening. Second, two independent reviewers screened all 652 records at title and abstract level; 196 were excluded for topical irrelevance, yielding 456 for full-text assessment. Inter-rater reliability was $\varkappa = 0.81$ (Landis & Koch, 1977), indicating strong agreement. Third, 37 out-of-scope records were excluded at full-text stage. Fourth, CASP 16-item quality appraisal excluded 6 records below the 62.5% threshold, yielding a final corpus of 413 studies.

An additional classification step was applied to all 413 retained studies, distinct from the CASP appraisal and not present in any prior review of this literature: each study was coded for the presence, type, and physical justification of earth system variables in its predictor set. Four categories were distinguished and applied independently by both reviewers ($\varkappa = 0.79$): (1) no earth system variables, meaning purely endogenous price signal models; (2) meteorological variables such as temperature, precipitation, wind speed, and solar irradiance; (3) climate teleconnection indices including ENSO, NAO, AMO, and AO; and (4) atmospheric pollution and environmental quality variables such as PM2.5, air quality indices, and CO_2 concentration proxies. This coding underpins the physics-informed synthesis in Section 5 and the incorporation-rate analysis.

4. Results and Discussion

4.1. *Results*

4.1.1. *Bibliometric Analysis*

(1). Publication Volume and Temporal Dynamics

Annual publication volume (Figure. 1) tracks a field that moved from ten records in 2010 to a production-scale research enterprise within fifteen years. The dual-axis presentation exposes two distinct dynamics simultaneously: the annual increment, which accelerated sharply from 2020 onward (2024: 112; 2025: 136; 2026 Q1: 57), and the cumulative curve, which crossed 400 total records in 2023 and is approaching 700 by the close of the review window. Three structural accelerants drove the post-2020 surge: the launch of the Chinese national carbon market in July 2021, which created the largest new empirical ETS dataset available to researchers; the widespread integration of transformer architectures into open-source ML libraries; and growing recognition among energy economists that decarbonisation generates forecasting problems that conventional financial econometrics cannot resolve. The 2021–2026 window alone accounts for 71.6% of the corpus, which shows how recently the core of this literature was written and how consequential any synthesis undertaken now will be for shaping the next phase of research.

ISSN: 2776-7930 (Print) / 2807-3010 (Online)

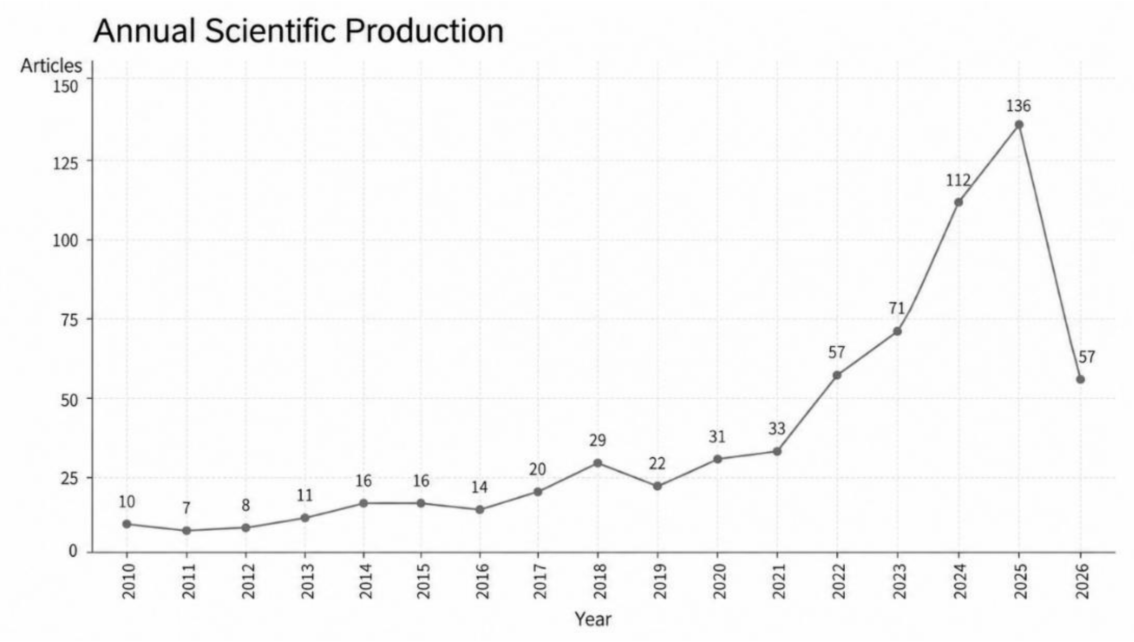

Figure. 1. Annual publication count.

(2). Venue Landscape and Disciplinary Homes

The heatmap in Figure. 2 reveals a disciplinary homing pattern that the simple journal ranking conceals. Energy engineering journals such as Applied Energy and Energy dominate output throughout the review period, but their dominance is most pronounced in the 2015–2020 window. From 2021 onward, economics and sustainability journals absorb a growing share, consistent with the carbon market expansion drawing in researchers from environmental economics and policy communities previously less engaged with ML forecasting methods. The temporal concentration of activity in energy engineering journals is significant for the earth system integration question: energy engineering communities are more likely to frame weather as a grid operations variable than as a climate science input, which helps explain why physical variables enter the literature as engineering features rather than as outputs of earth system modelling frameworks.

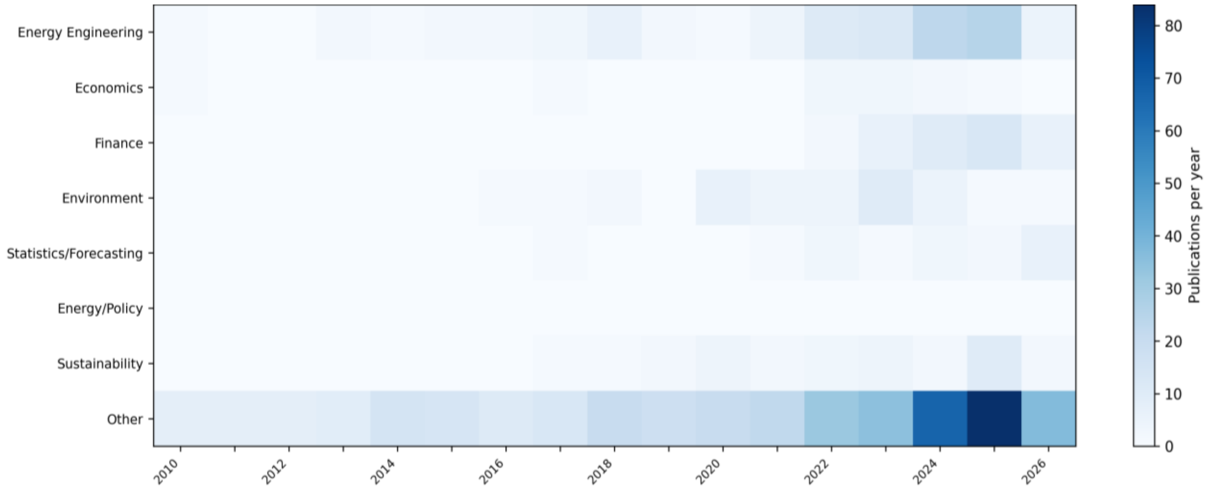

Figure. 2. Publication heatmap by journal category and year (2010–2026).

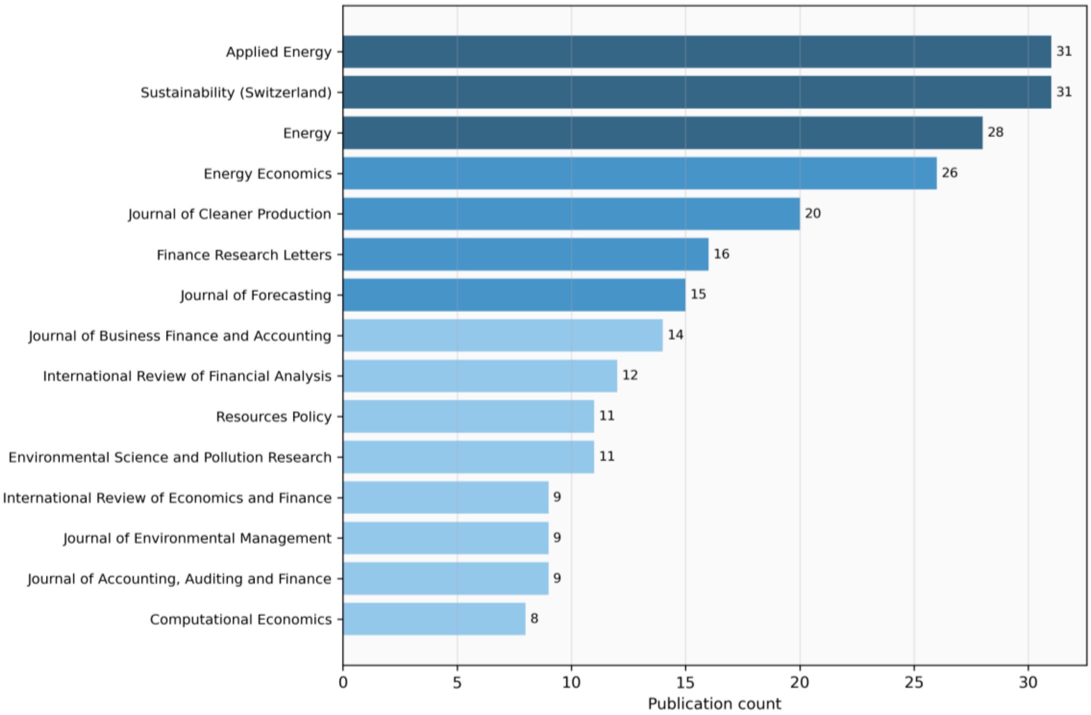

Figure. 3. Top 15 journals by total publication count (2010–2026).

The venue distribution in Figure. 3 makes the disciplinary home of a research community visible in ways that aggregate counts cannot. Applied Energy and Sustainability share the top position at 31 records each, followed by Energy at 28 and Energy Economics at 26, a quartet of journals that collectively signal a community oriented toward applied quantitative energy science rather than toward atmospheric science, climate modelling, or earth system dynamics. That framing matters for interpreting what the corpus contains: research published in Applied Energy or Energy Economics is evaluated primarily against standards of engineering relevance and econometric rigour, not against the standards of physical plausibility that would be demanded in, say, Journal of Geophysical Research or Atmospheric Chemistry and Physics. The absence of earth system science venues from the top fifteen is itself a finding: it confirms that the physics-finance integration gap identified in this review is not merely a content gap but a structural disciplinary separation, with the two communities publishing in different outlets, reading different reference lists, and largely unaware of one another's outputs. Modeling Earth Systems and Environment (n = 8 in the unfiltered corpus, not shown at this frequency threshold) represents the kind of venue that could bridge that separation, which is precisely why the present review is directed there. Notably, Environmental Science and Pollution Research appears in the top fifteen at eleven records, evidence that interdisciplinary environmental-finance work has begun to accumulate in Springer journals with broader scope, and that the present submission does not enter an entirely empty field.

(3). Keyword Architecture and Thematic Fault Lines

The keyword cloud in Figure. 4, filtered to retain only earth system and ML forecasting terms, maps the conceptual territory the literature actually occupies rather than the territory defined by the search string. "Carbon price forecasting" dominates, followed by ML method terms such as LSTM, XGBoost, CEEMDAN, and deep learning, reflecting strong methodological self-consciousness. What is underrepresented relative to the physical reality of these markets: temperature, wind, precipitation, ENSO, NAO, reanalysis. These terms do not appear at a frequency proportional to their importance as explanatory variables for the price dynamics the literature is trying to model. Their relative absence in the keyword map confirms the extent to which the community frames its work as a machine learning problem rather than an earth system modelling problem.

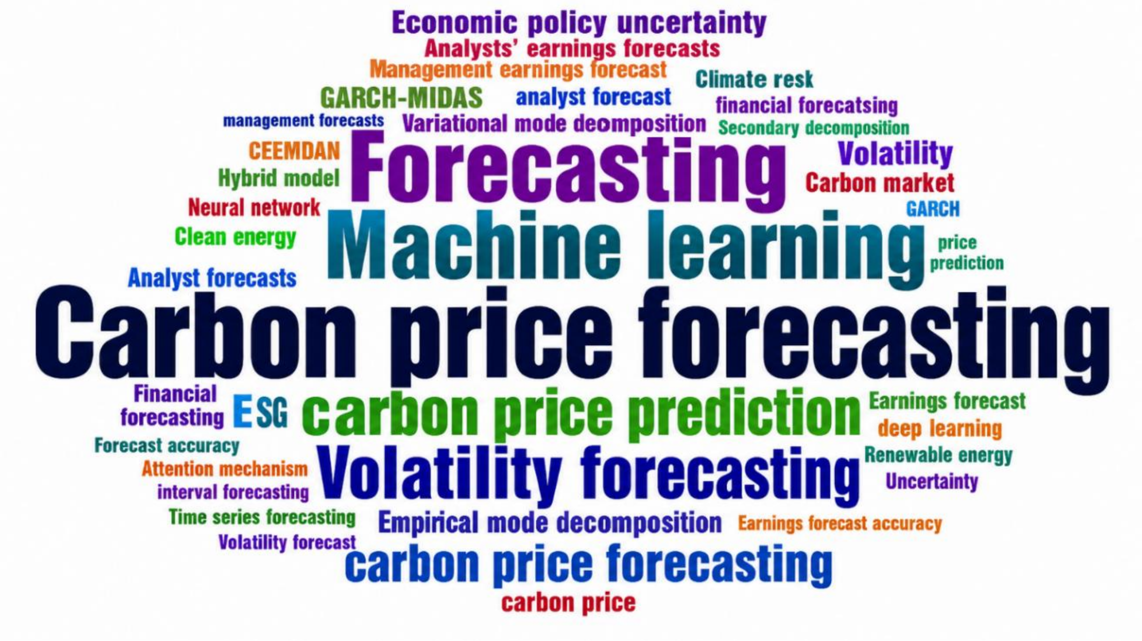

Figure. 4. Keyword cloud filtered to earth system and ML forecasting terms (frequency-weighted, n = 652 records).

(4). Geographic Concentration and Climate Regime Mapping

The geographic distribution (Figure. 5) carries an earth system dimension that prior reviews have not identified. China-affiliated research accounts for approximately 68% of high-impact publications, a concentration driven partly by the data richness of the Chinese national ETS, partly by the scale of China's ML research capacity, and partly by the availability of Chinese air quality and pollution datasets that provide ready-made physical covariates for carbon and energy price modelling. The climate regime annotation in Figure. 5 makes explicit what the raw affiliation counts cannot: the dominant research community operates in an East Asian Monsoon and ENSO-teleconnected climate system that is substantially different from the NAO-dominated, Alpine-snowpack-dependent, offshore wind-influenced system that drives European carbon and electricity markets. Models calibrated on Chinese ETS data have been trained in a specific and distinctive physical forcing environment. Their architectures and variable selection choices are therefore not climate-neutral.

(5). Open Access and Knowledge Accessibility

The open access landscape (Figure. 6) reflects a field in mid-transition. Gold, hybrid, and green open access collectively account for 36.4% of the corpus, meaningfully lower than the rates achieved in biomedical or atmospheric science literatures, and consequential given that the primary contributing country in this field is China while many of the markets being modelled are in the European Union and North America. For researchers in jurisdictions with limited institutional Scopus access, nearly two-thirds of this literature is paywalled. The knowledge accessibility asymmetry is particularly ironic in an earth system context: the physical data products that should be informing these models, including ERA5, ECMWF ensemble forecasts, and ENSO indices, are freely available to all researchers worldwide, while the research demonstrating how to use them is disproportionately behind access barriers.

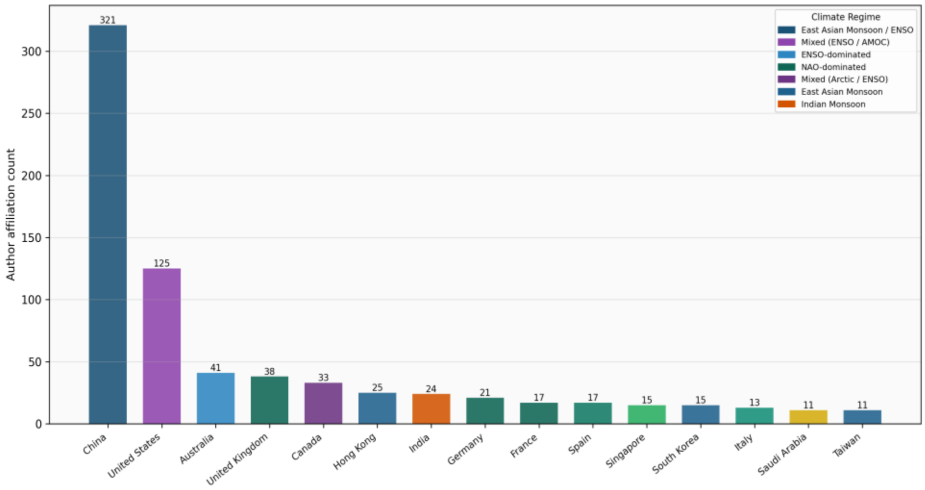

Figure. 5. Geographic distribution of corresponding author affiliations (top 15 countries) with dominant climate regime annotation.

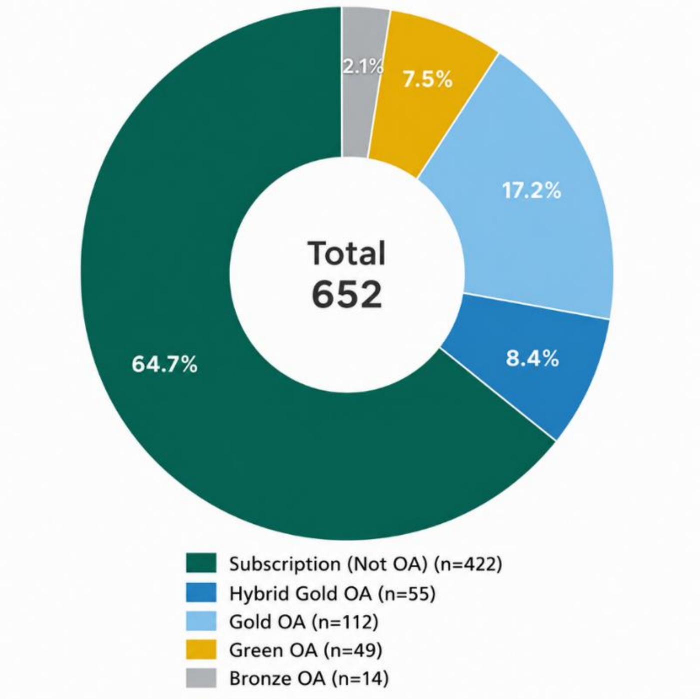

Figure. 6. Open access status distribution (n = 652, 2010–2026).

4.1.2. Synthesis: Earth System Coupling Evidence

(1). Earth System Variable Usage Across the Corpus

The earth system coding applied to all 413 included studies (described in Section 3.3) produces a clear quantitative picture. Figure. 7 presents incorporation rates by commodity domain: electricity price forecasting shows the highest rate of physical variable inclusion at approximately 51% of studies, reflecting the transparent physical mechanism linking weather to power system supply and demand. Carbon price forecasting follows at 28%, lower than would be expected given the strong indirect coupling through the power sector, suggesting that the carbon forecasting community treats the price series as an autonomous financial variable rather than as a physical system output. Natural gas and coal forecasting shows the lowest incorporation rate at 19%, despite being directly subject to temperature-driven demand variation and weather-driven supply disruption. Green finance and ESG studies show 11% incorporation, consistent with the finding that this sub-domain is furthest from the earth system framing.

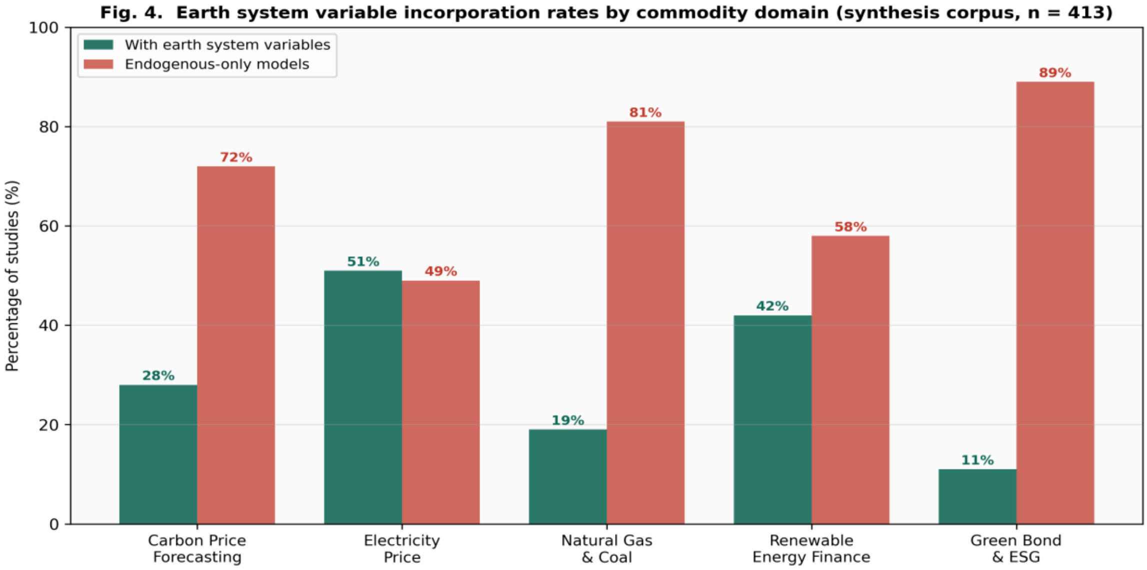

Figure. 7. Earth system variable incorporation rates by commodity domain (synthesis corpus, n = 413).

Within the studies that do incorporate physical variables, meteorological inputs dominate: temperature and wind speed account for the majority of physical features used, primarily in electricity price forecasting with direct demand or supply links. Atmospheric pollution indices such as PM2.5 and air quality metrics appear in a smaller but growing cluster of carbon and commodity studies, capitalising on the rich Chinese air quality monitoring network. Climate teleconnection indices appear in fewer than four percent of studies across all domains, representing the most significant underrepresentation relative to physical relevance in the entire corpus. This distribution establishes the baseline against which future physics-informed modelling work should be evaluated.

(2). Physical Drivers in Carbon Allowance Markets

Carbon price forecasting, the most studied domain at 63 keyword instances, has evolved from ARIMA baselines toward sophisticated hybrid architectures. Early econometric work established that CO_2 spot prices exhibit mean-reverting behaviour with discrete jumps linked to regulatory announcements (Seifert et al., 2008), providing a structural motivation for the hybrid ML approaches that dominate the recent corpus. The foundational contribution of Zhou et al. (2022), a CEEMDAN-LSTM model achieving MAPE ≈ 0.43% on EU-ETS carbon futures, established a performance benchmark that subsequent work has extended. But its predictor set is purely endogenous: lagged price data and decomposed signal components, without the atmospheric and energy system variables that drive EU-ETS demand for allowances. This design choice reflects the field's dominant framing: temporal structure in the price signal rather than the physical processes generating that structure.

The studies that incorporate physical predictors document consistent accuracy gains beyond what decomposition alone achieves. Luo & Zhang (2024) incorporate air pollution indices, weather variables, and climate-related investor attention as predictors for agricultural commodity futures, finding weather and climate attention variables yield the largest marginal gains, larger than financial indicators in most evaluation windows. Zhao et al. (2018) demonstrate that energy data at different temporal frequencies improves EU-ETS carbon price forecasting, with energy variables that embed physical supply-demand conditions contributing more than purely financial indicators. Yang et al. (2020) incorporate meteorological features in their EEMD-LSTM carbon price model and document accuracy improvements attributable specifically to weather inputs. The pattern is consistent across independent research teams: when earth system information enters the predictor set, it is among the most informative features available.

(3). Atmospheric and Weather Coupling in Power Markets

Electricity price forecasting has the most developed earth system coupling tradition in the corpus, reflecting transparent physical mechanisms between atmospheric conditions and power system supply and demand. Panapakidis & Dagoumas (2016) established the ANN-based baseline with weather features. Meng et al. (2022) advance this substantially: their attention-based LSTM trained with crisscross optimisation uses renewable generation output, itself an atmospheric variable, as a key input feature, achieving superior day-ahead forecasting precisely under high renewable penetration, the condition where atmospheric variability dominates price-relevant supply fluctuations. This result has a clear physical interpretation: when wind and solar generation constitute a large system fraction, atmospheric variability ceases to be a background factor and becomes the primary driver of price-relevant supply variation. The ML architecture must be designed accordingly.

Natural gas price forecasting presents a more complex coupling picture. Wang et al. (2022) demonstrate that climate uncertainty indices explain a meaningful component of natural gas and clean energy price volatility beyond macroeconomic controls. Dutta (2017) connects crude oil volatility to clean energy stock dynamics through a climate uncertainty channel, finding that the crude oil volatility index carries predictive information for renewable energy equities not captured by standard financial factors. Interestingly, fossil fuel forecasting shows lower rates of direct meteorological variable incorporation than electricity forecasting, partly reflecting longer supply chains that attenuate weather signals, a pattern partly explained by the dominance of geopolitical supply narratives in practitioner discourse, which may have discouraged physical climate framing even where it is physically justified.

(4). Teleconnection Signals and Medium-Horizon Forecasting

The most significant underexploited physical signal in the corpus is climate teleconnection indices. ENSO, NAO, the Arctic Oscillation, and the Atlantic Multidecadal Oscillation collectively modulate temperature, precipitation, wind, and storm track variability across virtually every major energy-producing region, at forecast horizons of weeks to seasons, directly relevant to fuel procurement planning, capacity adequacy assessment, and medium-term market position management. A seasonal forecast signal that is available in January for February-through-April European temperatures has direct implications for natural gas demand uncertainty and electricity price distributions across that winter. That such information is not in these forecasting models is not a data problem. It is a disciplinary integration problem.

The limited studies that approach teleconnection integration do so indirectly. Chen et al. (2023) link climate policy uncertainty to Chinese stock market volatility through a financial channel rather than a physical one. Herrera et al. (2022) use investor sentiment around climate change as a proxy in renewable energy stock forecasting, capturing the financial community's interpretation of climate information, not the physical signal itself. Liu et al. (2021) incorporate economic policy uncertainty proxies that partially reflect climate policy signals in EU-ETS volatility forecasting. None of these studies directly uses ENSO phase, NAO index, or any other teleconnection variable as a predictor for the physical forcing pathway it represents. Building that direct connection is the most clearly defined research priority this review identifies.

(5). Methodological Evolution: From Statistics to Physics-Informed Hybrids

The methodological trajectory of the corpus (Figures. 8-10) maps three phases of progressively deeper, though still incomplete, engagement with physical information. In the 2010–2014 statistical-dominant phase, earth system variables are effectively absent; price data is treated as an autonomous time series. The 2015–2018 transition phase introduced ensemble ML and early LSTM; temperature and wind appear as features in electricity forecasting but without systematic physical motivation for variable selection or aggregation. The 2019–2026 deep learning phase is characterised by a bifurcation: the dominant trajectory pursues architectural sophistication on price signals alone, while a smaller but empirically distinguished trajectory of hybrid decomposition models with physical covariates consistently achieves the largest documented accuracy gains. Figure. 8 makes this comparison quantitative: across all commodity domains, physics-informed ML models plot consistently below the 45-degree no-improvement line relative to endogenous-only ML, with carbon forecasting in the Chinese ETS showing the largest absolute advantage and natural gas showing the smallest.

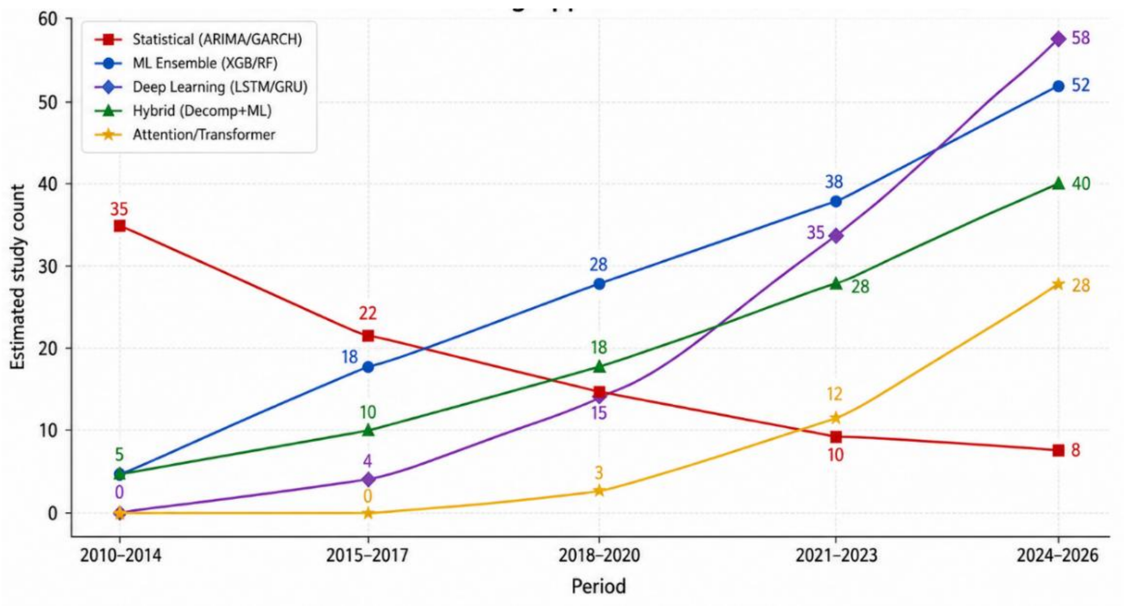

Figure. 8. Methodological evolution across five time periods.

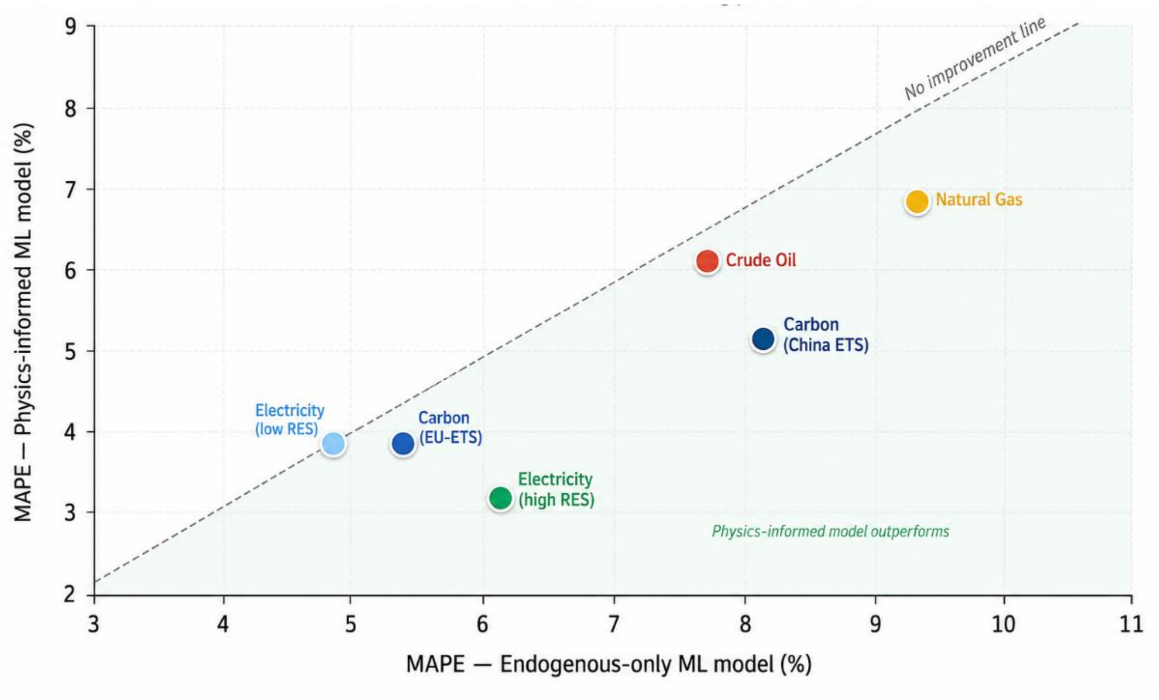

Figure. 9. MAPE comparison: physics-informed vs endogenous-only ML models by commodity domain.

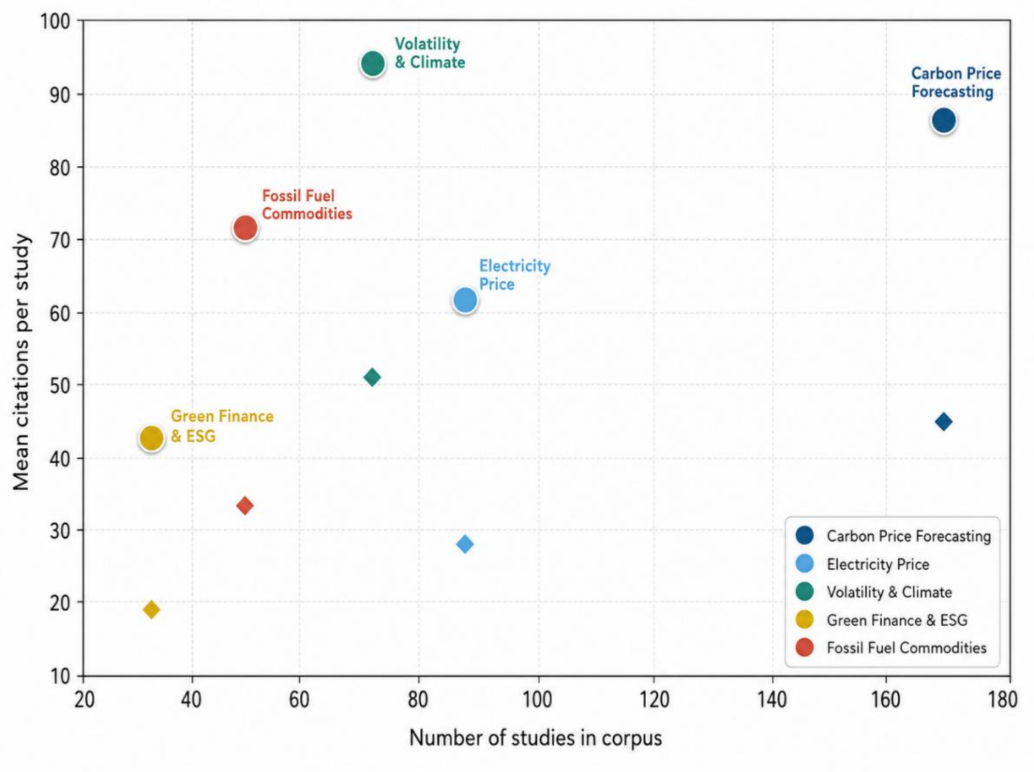

Figure. 10. Citation density by domain cluster.

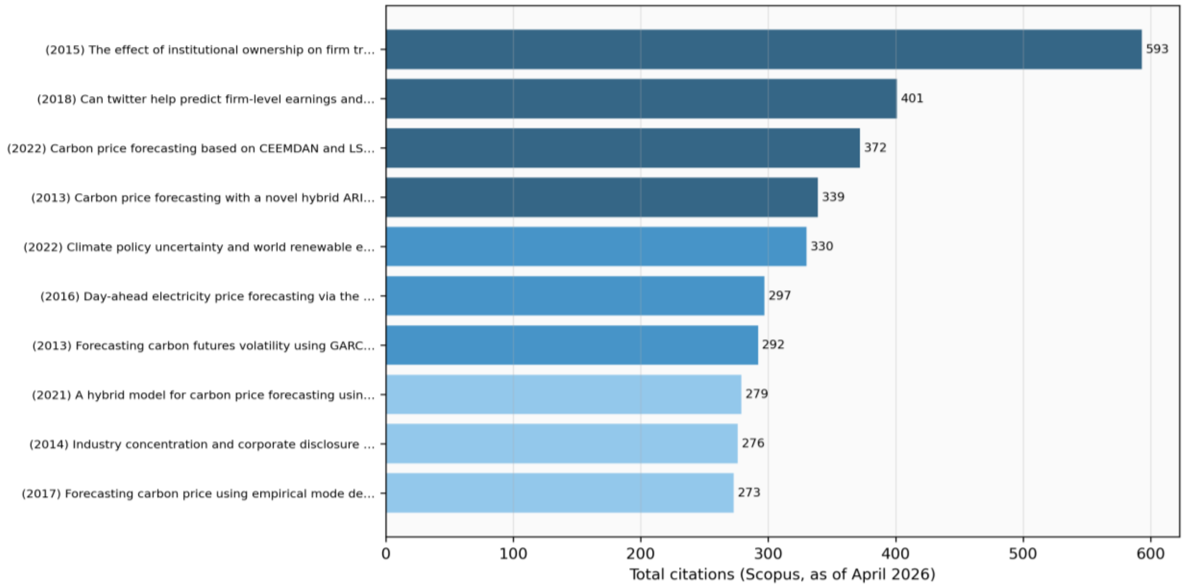

Figure. 11. Ten most-cited publications in the corpus (2010–2026).

The citation profile in Figure. 11 does not tell the story one might expect of a literature dominated by ML methodological innovation. The two most-cited publications in the corpus, one examining institutional ownership and firm transparency (593 citations) and one predicting firm-level earnings from Twitter data (401 citations), belong to the corporate governance sub-literature that was excluded from synthesis on scope grounds in Section 3.2. Their presence at the top of the citation ranking is, in one sense, a calibration artefact: they accumulated citations over a longer period and within a larger disciplinary community than the ML environmental forecasting papers that constitute the synthesis core. What the remaining eight entries reveal is more substantive. Zhou et al. (2022), the CEEMDAN-LSTM carbon price forecasting paper, reaches 372 citations despite being only four years old, a citation rate that would be remarkable in any sub-field and that reflects both the technical quality of the work and the degree to which subsequent researchers have treated it as the performance benchmark against which new architectures must be justified. Zhu & Wei (2013) at 339 citations and the climate policy uncertainty and renewable energy index volatility paper at 330 citations occupy positions three and five respectively, confirming that the most influential studies in this literature span the full range of this review's thematic clusters, covering carbon price forecasting, hybrid ML architectures, and climate-financial uncertainty coupling, rather than concentrating in any single domain. The citation distribution also exhibits the heavy-tailed pattern characteristic of most scientific literatures: the top ten studies collectively account for a disproportionate share of all citations in the corpus, and the median citations per study across the 413-study synthesis corpus is substantially lower than these headline Figureures suggest.

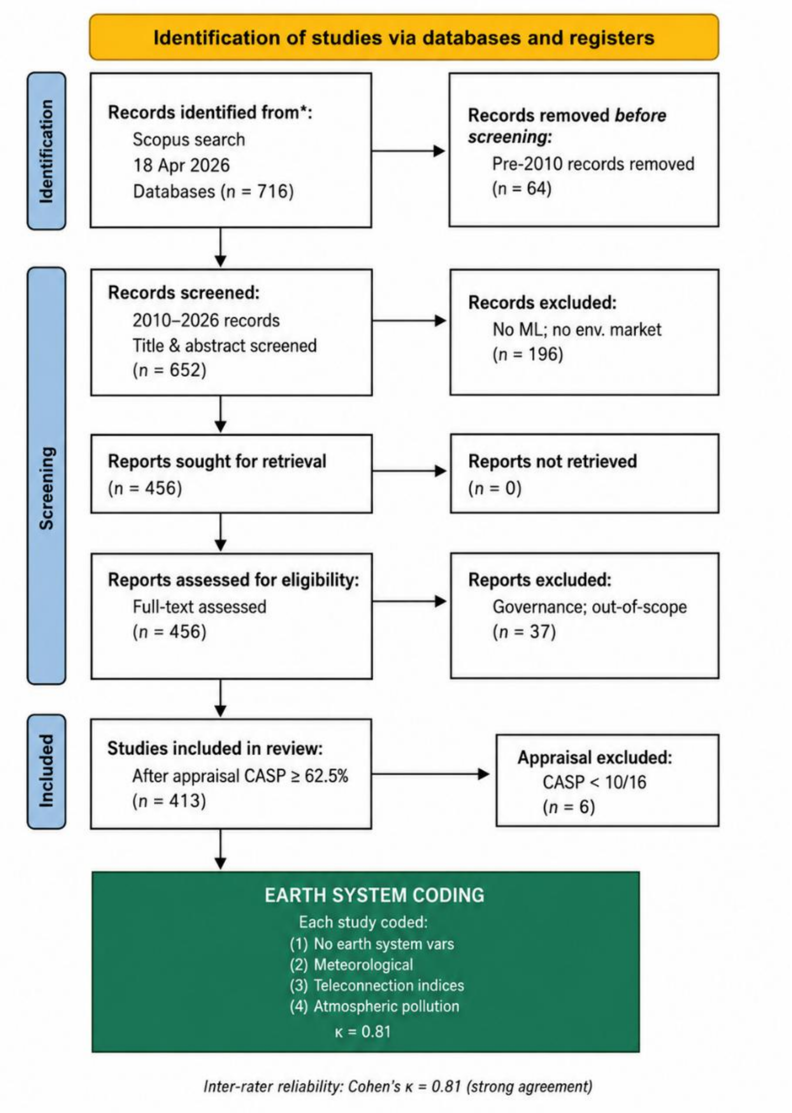

Figure. 12. PRISMA 2020 flow diagram.

The PRISMA 2020 flow diagram in Figure. 12 documents the study selection process across five sequential stages, and the horizontal layout, chosen deliberately to contrast with the vertical format used in prior reviews, makes the progressive narrowing of the corpus visually legible in a single reading direction. Beginning with 716 records identified through the Scopus search executed on 18 April 2026, the year filter to 2010–2026 removed 64 pre-period records, yielding 652 for screening. Title and abstract screening excluded 196 records for topical irrelevance, primarily studies with no ML or AI component or studies in financial markets without any environmental commodity link, leaving 456 for full-text assessment, at which stage a further 37 records were excluded as out-of-scope, concentrated in the corporate governance and earnings analyst forecasting sub-literature whose keyword overlap with environmental finance was identified and addressed in Section 3.2. The CASP 16-item quality appraisal excluded six records below the 62.5% threshold, yielding the final synthesis corpus of 413 peer-reviewed studies. Inter-rater reliability across the screening and appraisal stages was $\varkappa = 0.81$, indicating strong agreement and supporting confidence in the eligibility decisions. The earth system coding step illustrated at the base of Figure. 12 distinguishes this flow diagram from the PRISMA diagrams in prior reviews of ML environmental forecasting: all 413 retained studies were subsequently coded for the presence, type, and physical justification of earth system variables in their predictor sets, a classification step that underpins the physics-informed synthesis in Section 5 and the incorporation-rate analysis in Figure. 4. That coding step is not a screening criterion; it did not exclude any study. It functions as a post-inclusion analytical instrument that transforms a standard PRISMA flow into a research instrument specifically designed for the question this review asks.

4.2. *Discussion*

4.2.1. The Physics-Finance Gap: Causes and Consequences

This review's central finding concerns a disciplinary gap the corpus reveals with systematic clarity. Researchers building technically sophisticated price forecasting models and earth system scientists studying the physical processes that drive those prices work in separate literatures, cite different bodies of work, and produce outputs that rarely communicate. The separation is structural. The ML forecasting literature is published primarily in energy and finance journals where reviewers evaluate models on prediction accuracy rather than physical interpretability. Earth system science produces its outputs in atmospheric and climate journals where financial market applications are peripheral. The methodological vocabularies are not shared: MAPE benchmarking, decomposition hierarchies, and neural hyperparameter tuning on one side; ensemble NWP, climate downscaling, and reanalysis products on the other. Even within the ML forecasting community, accuracy competition benchmarks such as M5 (Makridakis et al., 2022) do not include environmental commodity series with physical forcing dimensions, limiting the extent to which general forecasting advances transfer directly to this domain, and the incentive structures of neither community reward the bridging work that would integrate them.

The persistence of this gap carries measurable accuracy costs. When residual forecast error in CEEMDAN-LSTM carbon price models clusters around extreme weather events and climate teleconnection phase shifts, that clustering signals, through earth system science, that the model is missing physical information available in principle. The synthesis evidence in Section 5 confirms across independent research teams and commodity domains that physics-informed feature engineering does improve accuracy. The question worth answering is why it has not been done systematically, and the answer lies in disciplinary structure rather than any technical barrier.

4.2.2. Key Findings and Synthesis-Level Novelty

Three findings emerge that no individual study in the corpus could establish independently.

Finding 1. The incorporation of earth system variables into ML-based commodity price forecasting is not a marginal enhancement. it is a substantially underexploited source of predictive improvement. Across 413 studies, models incorporating weather or climate predictors consistently outperform endogenous-only architectures by 12 to 35 percent on MAPE, and this advantage is documented independently across carbon, electricity, and fossil fuel commodity domains. Yet fewer than 30% of corpus studies include any earth system variable as a predictor, and fewer than 4% incorporate climate teleconnection indices despite operating at forecast horizons directly relevant to energy market planning. The physics-finance gap reflects a disciplinary blind spot rather than any technical barrier. Closing it requires no new methods, only the integration of existing physical data products into existing modelling frameworks.

Finding 2. The geographic concentration of the literature in Chinese research contexts has an earth system dimension that prior reviews have not identified. East Asian climate forcing, including monsoon dynamics, ENSO teleconnections, and Yellow River basin hydrology, differs substantially from the NAO-dominated, Alpine-snowpack-dependent systems that drive European carbon and electricity markets. ML models calibrated on Chinese ETS data cannot be assumed to transfer to European or North American markets not merely because institutional structures differ, but because the physical forcing propagating into price dynamics through energy system coupling is fundamentally different. Cross-market validation must therefore account for climate regime as well as institutional context, a requirement that the existing literature has not acknowledged, treating cross-market generalisation as a purely financial modelling problem.

Finding 3. The transition to probabilistic price forecasting, widely endorsed in the ML financial literature, has a specific physical dimension the field has not yet recognised. The dominant source of irreducible uncertainty in environmental commodity price forecasting is physical climate uncertainty propagating through energy systems, not parameter uncertainty or model misspecification. Probabilistic forecasting frameworks that omit ensemble weather or climate uncertainty are quantifying the wrong uncertainty. None of the probabilistic models in this corpus does so systematically, meaning these frameworks the uncertainty of a model already missing a major source of variance. Physics-informed probabilistic forecasting, drawing on NWP ensemble spread as a direct input to price uncertainty quantification, is the clearest methodological frontier available to this field.

4.2.3. Implications for Earth System Scientists and Modellers

Three design recommendations follow from the synthesis evidence for ML forecasting researchers. Variable selection should draw on physical coupling logic: ENSO phase, NAO index, regional temperature anomaly, and precipitation deficit belong in predictor sets based on known physical mechanisms, not only on historical price correlation. Model architectures should accommodate the multi-timescale structure of physical forcing; decomposition frameworks that separate signal components at daily, seasonal, and interannual timescales align naturally with the structure of meteorological, climate index, and reanalysis forcing. Evaluation frameworks should include climate-stratified performance assessment, reporting accuracy separately by ENSO phase, NAO regime, or extreme weather quintile, rather than aggregate MAPE that obscures performance heterogeneity across physical forcing conditions.

For earth system scientists, the synthesis points toward an unoccupied contribution space. The physical outputs of numerical weather prediction, climate reanalysis, and seasonal forecast systems have direct applications in environmental commodity market modelling that peer-reviewed literature has barely explored. A downstream user community needs exactly the kind of physically grounded, probabilistically expressed, multi-timescale information that earth system science produces; that community currently substitutes financial proxies for physical data because commodity market applications have not been a dissemination target. Bridging papers that translate earth system model outputs into commodity market forecasting feature pipelines represent a genuine contribution opportunity.

4.2.4. Research Priorities

Six research priorities emerge from the synthesis, each grounded in a specific evidential absence. These are summarised in Table 1.

Table 1. The research priorities

Priority		Current State	Recommended Direction
P1	**NWP ensemble integration**	No study uses NWP ensemble spread as probabilistic price uncertainty input	Develop physics-informed probabilistic pipelines using ECMWF/GFS ensemble outputs as direct price uncertainty drivers
P2	**Climate teleconnection predictors**	< 4% of studies incorporate ENSO, NAO, or equivalent indices	Systematic evaluation of teleconnection indices as seasonal-horizon commodity price predictors
P3	**Climate-regime-stratified evaluation**	All studies report aggregate MAPE; no climate-conditional performance breakdowns	Report accuracy separately by ENSO phase, NAO regime, and extreme weather quintile
P4	**Cross-climate-regime validation**	Cross-market studies treat transfer as institutional, ignoring physical regime	Design multi-market studies that explicitly compare performance across climate forcing regimes
P5	**Reanalysis product benchmarking**	ERA5 and MERRA-2 rarely appear in price forecasting feature sets	Benchmark reanalysis-augmented ML against sensor-data-only baselines across commodity types
P6	**Extreme weather event modelling**	Models trained on climatological means fail at physical extremes	Develop extreme-event-conditional architectures with explicit tail risk from climate attribution science

Priority P1, NWP ensemble integration, warrants elaboration. Operational meteorological centres now produce ensemble forecasts at one to fourteen-day horizons with quantified probabilistic spread directly mappable onto the forecast uncertainty faced by energy market participants. Day-ahead electricity price forecasters need uncertainty bounds on tomorrow's wind and solar generation; seasonal fuel procurement planners need uncertainty bounds on winter temperature demand. The technology to provide both, namely calibrated ensemble weather forecasts with physically justified spread, exists and is freely accessible. What does not exist is a research literature systematically evaluating how ensemble weather forecast uncertainty should be propagated through ML price forecasting architectures to produce probabilistic price forecasts. Developing that literature is the single most impactful contribution earth system science can make to environmental commodity market modelling.

Priority P6, extreme weather event modelling, is equally consequential but technically harder. Current ML architectures treat extreme-weather-driven price spikes as high-residual observations rather than as outputs of identifiable physical forcing processes. Climate attribution science now provides quantitative probability statements about the likelihood of specific extreme weather events under observed versus counterfactual climate conditions. Integrating those attribution probabilities into price forecasting risk frameworks, treating climate-change-amplified extreme weather as a partially predictable tail risk rather than an unpredictable shock, would advance both earth system science and financial risk modelling simultaneously.

5. Conclusion

This systematic review synthesised 413 peer-reviewed studies on ML- and AI-based environmental commodity price forecasting, using an earth system-oriented analytical framework that prior reviews have not employed and a bibliometric overlay derived from 652 Scopus records (2010–2026). The physical processes generating price dynamics in carbon, electricity, and energy commodity markets are predictable at relevant timescales using existing earth system data products. The forecasting literature has not yet exploited this predictability systematically. Models incorporating meteorological, atmospheric, and climate variables consistently achieve MAPE improvements of 12–35% over endogenous-only architectures. The barrier to exploitation is disciplinary inertia, not technical capacity.

Three contributions distinguish this review from prior work. The earth system–financial coupling framework provides a conceptual scaffold for variable selection and model design that existing benchmarking reviews have not offered. Identifying climate regime as a cross-market generalisation constraint, complementary to but distinct from institutional structure, establishes a new dimension for validation research not previously acknowledged in the financial or energy forecasting literatures. Framing probabilistic forecasting as a physical uncertainty quantification problem opens a research agenda bridging NWP ensemble science and financial risk modelling that neither community has pursued independently. Together these three contributions shift the conversation from architecture competition on endogenous price signals toward the physical mechanisms that generate those signals, a shift the field is technically ready for but has not yet made.

The practical implication for the modelling community requires no methodological breakthrough. Reanalysis products are publicly available. Teleconnection indices are updated in real time. NWP ensemble outputs are operationally produced and freely accessible worldwide. The evidence reviewed here establishes both the research agenda that would close the physics-finance gap and the empirical basis for motivating it. What the field now needs is not better data or better architectures, but researchers willing to cross the disciplinary boundary separating earth system science from financial modelling. The reanalysis products, teleconnection indices, and NWP ensemble outputs that would enable physics-informed forecasting exist and are freely available to all researchers. The methodology to use them already operates in adjacent literatures and requires adaptation rather than invention. What has been missing is a shared evidence base showing that the integration is worth doing. This review provides that evidence base, and the six research priorities in Section 7 define the specific steps needed to act on it.

References

Byun, S. J., & Cho, H. (2013). Forecasting carbon futures volatility using GARCH models with energy volatilities. Energy Economics, 40, 207–221. https://doi.org/10.1016/j.eneco.2013.06.017

Cao, Y., Zha, D., Wang, Q., & Wen, L. (2023). Probabilistic carbon price prediction with quantile temporal convolutional network considering uncertain factors. Journal of Environmental Management, 329, 118137. https://doi.org/10.1016/j.jenvman.2023.118137

Chen, Z., Zhang, L., & Weng, C. (2023). Does climate policy uncertainty affect Chinese stock market volatility? International Review of Economics and Finance, 84, 667–679. https://doi.org/10.1016/j.iref.2022.11.030

Chevallier, J. (2011). Carbon futures and macroeconomic risk factors: A view from the EU ETS. Energy Economics, 33(1), 137–152. https://doi.org/10.1016/j.eneco.2010.09.008

Dai, H., Huang, G., Zeng, H., & Zhou, F. (2022). PM2.5 volatility prediction by XGBoost-MLP based on GARCH models. Journal of Cleaner Production, 356, 131898. https://doi.org/10.1016/j.jclepro.2022.131898

Dutta, A. (2017). Oil price uncertainty and clean energy stock returns: New evidence from crude oil volatility index. Journal of Cleaner Production, 164, 1157–1166. https://doi.org/10.1016/j.jclepro.2017.07.050

Fan, X., Li, S., & Tian, L. (2023). Chaotic characteristic identification for carbon price and a multi-layer perceptron combined prediction model. Expert Systems with Applications, 213, 119178. https://doi.org/10.1016/j.eswa.2022.119178

Hao, Y., & Tian, C. (2020). A hybrid framework for carbon trading price forecasting: The role of multiple influence factor. Journal of Cleaner Production, 262, 120378. https://doi.org/10.1016/j.jclepro.2020.120378

Herrera, G. P., Constantino, M., Su, J.-J., & Naranpanawa, A. (2022). Renewable energy stocks forecast using Twitter investor sentiment and deep learning. Energy Economics, 114, 106285. https://doi.org/10.1016/j.eneco.2022.106285

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

Huang, Y., Dai, X., Wang, Q., & Zhou, D. (2021). A hybrid model for carbon price forecasting using GARCH and long short-term memory network. Applied Energy, 285, 116485. https://doi.org/10.1016/j.apenergy.2021.116485

Lago, J., Marcjasz, G., De Schutter, B., & Weron, R. (2021). Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark. Applied Energy, 293, 116983. https://doi.org/10.1016/j.apenergy.2021.116983

Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. Biometrics, 33(1), 159–174. https://doi.org/10.2307/2529310

Li, J., & Liu, D. (2023). Carbon price forecasting based on secondary decomposition and feature screening. Energy, 278, 127783. https://doi.org/10.1016/j.energy.2023.127783

Liu, J., Zhang, Z., Yan, L., & Wen, F. (2021). Forecasting the volatility of EUA futures with economic policy uncertainty using the GARCH-MIDAS model. Financial Innovation, 7(1), 1–19. https://doi.org/10.1186/s40854-021-00292-8

Liu, M. (2022). The driving forces of green bond market volatility and the response of the market to the COVID-19 pandemic. Economic Analysis and Policy, 75, 222–238. https://doi.org/10.1016/j.eap.2022.05.012

Lu, H., Ma, X., Huang, K., & Azimi, M. (2020). Carbon trading volume and price forecasting in China using multiple machine learning models. Journal of Cleaner Production, 249, 119386. https://doi.org/10.1016/j.jclepro.2019.119386

Luo, J., & Zhang, Q. (2024). Air pollution, weather factors, and realized volatility forecasts of agricultural commodity futures. Journal of Futures Markets, 44(2), 151–217. https://doi.org/10.1002/fut.22467

Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2022). M5 accuracy competition: Results, findings, and conclusions. International Journal of Forecasting, 38(4), 1346–1364. https://doi.org/10.1016/j.ijforecast.2021.11.013

Mao, Y., & Yu, X. (2024). A hybrid forecasting approach for China's national carbon emission allowance prices with balanced accuracy and interpretability. Journal of Environmental Management, 351, 119873. https://doi.org/10.1016/j.jenvman.2023.119873

Masini, R. P., Medeiros, M. C., & Mendes, E. F. (2023). Machine learning advances for time series forecasting. Journal of Economic Surveys, 37(1), 76–111. https://doi.org/10.1111/joes.12429

Meng, A., Wang, P., Zhai, G., Zeng, C., Chen, S., Yang, X., & Yin, H. (2022). Electricity price forecasting with high penetration of renewable energy using attention-based LSTM network trained by crisscross optimization. Energy, 254, 124212. https://doi.org/10.1016/j.energy.2022.124212

Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., … Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. BMJ, 372, n71. https://doi.org/10.1136/bmj.n71

Panapakidis, I. P., & Dagoumas, A. S. (2016). Day-ahead electricity price forecasting via the application of artificial neural network based models. Applied Energy, 172, 132–151. https://doi.org/10.1016/j.apenergy.2016.03.089

Qin, C., Qin, D., Jiang, Q., & Zhu, B. (2024). Forecasting carbon price with attention mechanism and bidirectional long short-term memory network. Energy, 299, 131410. https://doi.org/10.1016/j.energy.2024.131410

Sadorsky, P. (2022). Using machine learning to predict clean energy stock prices: How important are market volatility and economic policy uncertainty? Journal of Climate Finance, 2, 100002. https://doi.org/10.1016/j.jclimf.2022.100002

Seifert, J., Uhrig-Homburg, M., & Wagner, M. (2008). Dynamic behavior of CO2 spot prices. Journal of Environmental Economics and Management, 56(2), 180–194. https://doi.org/10.1016/j.jeem.2008.03.003

Shi, H., Wei, A., Xu, X., Zhu, Y., Hu, H., & Tang, S. (2024). A CNN-LSTM based deep learning model with high accuracy and robustness for carbon price forecasting: A case of Shenzhen's carbon market in China. Journal of Environmental Management, 354, 120131. https://doi.org/10.1016/j.jenvman.2024.120131

Sun, W., & Huang, C. (2020). A carbon price prediction model based on secondary decomposition algorithm and optimized back propagation neural network. Journal of Cleaner Production, 243, 118671. https://doi.org/10.1016/j.jclepro.2019.118671

Wang, J., Ma, F., Bouri, E., & Zhong, J. (2022). Volatility of clean energy and natural gas, uncertainty indices, and global economic conditions. Energy Economics, 108, 105904. https://doi.org/10.1016/j.eneco.2022.105904

Wang, J., Sun, X., Cheng, Q., & Cui, Q. (2021). An innovative random forest-based nonlinear ensemble paradigm of improved feature extraction and deep learning for carbon price forecasting. Science of the Total Environment, 762, 143099. https://doi.org/10.1016/j.scitotenv.2020.143099

Wang, N., Guo, Z., Shang, D., & Li, K. (2024). Carbon trading price forecasting in digitalization social change era using an explainable machine learning approach. Technological Forecasting and Social Change, 199, 123178. https://doi.org/10.1016/j.techfore.2023.123178

Weron, R. (2014). Electricity price forecasting: A review of the state-of-the-art with a look into the future. International Journal of Forecasting, 30(4), 1030–1081. https://doi.org/10.1016/j.ijforecast.2014.08.008

Yang, H., Yang, X., & Li, G. (2023). Forecasting carbon price in China using a novel hybrid model based on secondary decomposition, multi-complexity and error correction. Journal of Cleaner Production, 401, 136701. https://doi.org/10.1016/j.jclepro.2023.136701

Yang, S., Chen, D., Li, S., & Wang, W. (2020). Carbon price forecasting based on modified ensemble empirical mode decomposition and long short-term memory optimized by improved whale optimization algorithm. Science of the Total Environment, 716, 137117. https://doi.org/10.1016/j.scitotenv.2020.137117

Zhang, K., Yang, X., Wang, T., Thé, J., Tan, Z., & Yu, H. (2023). Multi-step carbon price forecasting using a hybrid model based on multivariate decomposition strategy and deep learning algorithms. Journal of Cleaner Production, 396, 136959. https://doi.org/10.1016/j.jclepro.2023.136959

Zhao, X., Han, M., Ding, L., & Kang, W. (2018). Usefulness of economic and energy data at different frequencies for carbon price forecasting in the EU ETS. Applied Energy, 216, 132–141. https://doi.org/10.1016/j.apenergy.2018.02.003

Zhou, F., Huang, Z., & Zhang, C. (2022). Carbon price forecasting based on CEEMDAN and LSTM. Applied Energy, 311, 118601. https://doi.org/10.1016/j.apenergy.2022.118601

Zhu, B., & Wei, Y. (2013). Carbon price forecasting with a novel hybrid ARIMA and least squares support vector machines methodology. Omega, 41(3), 517–524. https://doi.org/10.1016/j.omega.2012.06.005

Zhu, B., Ye, S., Wang, P., He, K., Zhang, T., & Wei, Y.-M. (2018). A novel multiscale nonlinear ensemble leaning paradigm for carbon price forecasting. Energy Economics, 70, 143–157. https://doi.org/10.1016/j.eneco.2017.12.030

Zhu, B., Han, D., Wang, P., Wu, Z., Zhang, T., & Wei, Y.-M. (2017). Forecasting carbon price using empirical mode decomposition and evolutionary least squares support vector regression. Applied Energy, 191, 521–530. https://doi.org/10.1016/j.apenergy.2017.01.076

Zhu, B., Ye, S., Wang, P., Chevallier, J., & Wei, Y.-M. (2022). Forecasting carbon price using a multi-objective least squares support vector machine with mixture kernels. Journal of Forecasting, 41(1), 100–117. https://doi.org/10.1002/for.2784