



UNIVERSITÀ
DEGLI STUDI
DI BRESCIA

UNIVERSITÀ DEGLI STUDI DI BRESCIA

DIPARTIMENTO DI ECONOMIA E MANAGEMENT

DOTTORATO DI RICERCA IN MODELLI E METODI PER L'ECONOMIA E IL
MANAGEMENT - ANALYTICS FOR ECONOMICS AND MANAGEMENT

(AEM)

CYCLE XXXVII

SSD: SECS-S/06

METHODS AND MODELS FOR FINANCE, ENERGY AND THE
ENVIRONMENT

**Exploring the European Union Emissions Trading System:
Non-Parametric Models for Price Determinants, Forecasting, and
Volatility Causality**

Supervisor:

Chiar.ma Prof.ssa Maria Elena De Giuli

Ph.D. Candidate:

Cristiano Salvagnin

ID: 739385

Academic Year: 2024/2025

Abstract (IT)

Il Sistema Europeo di Scambio delle Emissioni (EU ETS) svolge un ruolo fondamentale nella politica climatica europea, regolando le emissioni di gas a effetto serra attraverso un meccanismo di cap-and-trade. Questa tesi presenta tre analisi volte ad approfondire la comprensione del mercato EU ETS. In primo luogo, una revisione bibliometrica di 367 articoli scientifici (2004-2024) esamina i modelli di citazione, le reti di co-autori e le tendenze delle parole chiave per identificare autori, istituzioni e temi fondamentali, oltre a mettere in luce lacune nella ricerca. Tra i risultati principali si evidenzia un tasso di crescita annuo del 12,99%, con picchi di produzione scientifica corrispondenti a cambiamenti normativi, e contributi rilevanti da Germania, Cina e Francia. La revisione mette in evidenza tematiche emergenti come il carbon pricing e la volatilità del mercato, con una crescente attenzione agli aspetti finanziari. Nonostante i progressi, si riscontrano ancora diverse lacune, in particolare nello sviluppo e nell'applicazione di metodi non parametrici e nello studio dell'impatto delle variabili macroeconomiche. La seconda e la terza analisi affrontano tali lacune utilizzando metodologie innovative, non parametriche e model-free. Il metodo dell'Information Imbalance (II) identifica le variabili più informative per la previsione dei prezzi dell'EU ETS, rivelando che, mentre la Fase 3 risulta influenzata da indici energetici e materie prime, la Fase 4 è dominata da variabili finanziarie, probabilmente a causa di shock economici come la pandemia di COVID-19 e la crisi energetica. I risultati metodologici dimostrano che il metodo II seleziona efficacemente la scala temporale settimanale come la più informativa per effettuare previsioni accurate. I Gaussian Processes (GP) vengono utilizzati per affrontare il problema della frequenza mista integrando le variabili macroeconomiche nell'analisi e migliorando ulteriormente l'accuratezza nelle previsioni sia attuali (nowcasting) sia future

(forecasting). Inoltre, questa tesi introduce il Differentiable Information Imbalance (DII), un approccio non lineare, basato sull'Information Imbalance, che rileva relazioni causali tra il mercato EU ETS e variabili finanziarie. Il DII viene confrontato con il modello Vector AutoRegression (VAR) utilizzando dati sintetici, dimostrando la sua capacità di catturare interazioni non lineari. Nell'ambito del DII, viene sviluppata una nuova metrica di forza causale, l'Imbalance Gain (IG), per misurare l'influenza causale delle variabili predittive, in modo analogo alla statistica F. L'analisi empirica dal 2013 al 2024 rivela effetti causali significativi dell'IBEX35 e dei futures sul carbone sui prezzi delle EUA, evidenziando i legami complessi tra mercati energetici, finanziari e del carbonio. Attraverso l'impiego di queste metodologie avanzate per colmare le lacune nella letteratura, questo lavoro mira a migliorare i sistemi di scambio delle emissioni e a sostenere politiche climatiche solide per contrastare il cambiamento climatico.

This page has been intentionally left blank.

Acknowledgments

First and foremost, I would like to express my deepest gratitude to my supervisor, Maria Elena De Giuli, for her guidance, support, and encouragement throughout my PhD journey and academic career.

I am profoundly grateful to my collaborators, particularly Aldo Glielmo and Antonietta Mira, whose insights and contributions have greatly enriched this research and profoundly impacted me with their incredible support and expertise. I also acknowledge Vittorio Del Tatto and all the other scholars who have been involved in this journey.

My heartfelt thanks go to my PhD colleagues, who have provided a stimulating and supportive environment for both academic and personal growth. Additionally, I am grateful to the administrative and technical staff, whose assistance and professionalism have greatly facilitated my research efforts.

Most importantly, I owe my deepest thanks to my family, Sandro, Elena, and Michelle, for their unwavering love, patience, and support, without which this journey would not have been possible.

Finally, to Riccardo, I promise.

Contents

Contents	v
List of Figures	ix
List of Tables	xxi
1 Introduction	1
1.1 The EU ETS	1
1.2 History and Development of the EU ETS	3
1.2.1 Early Foundation and Policy Context	3
1.2.2 Phase 1 (2005-2007): Pilot Phase	5
1.2.3 Phase 2 (2008–2012): Expansion and Alignment with Kyoto Protocol	6
1.2.4 Phase 3 (2013–2020): System Strengthening	8
1.2.5 Phase 4 (2021–2030): Expansion and Ambition	10
1.3 Market Organisation and Compliance	12
1.3.1 How the Market Works	12
1.3.2 Compliance Mechanism	14
1.3.3 Market Oversight	15
1.4 Sectoral Coverage	16
1.5 Future Directions	18

1.5.1	The EU ETS 2	18
1.5.2	Integration with International Carbon Markets	19
1.6	Thesis Motivation and Objectives	21
1.7	Thesis Organisation	23
2	Literature Review	26
2.1	Exploring the EU ETS Market: Bibliometric and Review Literature	26
2.2	EU ETS Pricing Dynamics and Forecasting Methods	29
2.3	Volatility Patterns and Causal Links in the EU ETS	35
3	Methods and Models	40
3.1	The Information Imbalance	40
3.1.1	Introduction to Information Content and Feature Selection	40
3.1.2	Theory of Copulas in Defining Information Imbalance	41
3.1.3	Connection to Information Theory	42
3.1.4	The Information Imbalance Metric	43
3.1.5	Beyond Linear Correlation	47
3.2	The Differentiable Information Imbalance	49
3.2.1	Introduction to Causal Discovery and Relationships: Moving Beyond Granger Causality	49
3.2.2	Linear Causal Discovery: VAR and Granger's F-statistic	50
3.2.3	Non-Linear Causal Discovery: the Differentiable Information Imbalance and Imbalance Gain	53
3.2.4	Two Case Studies on Synthetic Data	60
3.3	Gaussian Processes	64
3.3.1	Kernel Function	65

3.3.2	Prediction and Uncertainty Estimation	66
3.3.3	Applications in Forecasting, Imputation, and Aggregation	67
4	Empirical Results I: Bibliometric insights into the EU ETS market	68
4.1	Data Collection and Selection Criteria	68
4.2	Bibliometric Tools	70
4.3	Descriptive Analysis	72
4.4	Document Analysis	78
4.5	Conceptual Structure	81
4.6	Discussion	84
5	Empirical Results II: Price Determinants and Forecasting in the EU ETS	92
5.1	Data	92
5.1.1	Descriptive Statistics	95
5.2	Information Imbalance Analysis	99
5.3	Greedy Selection of Variables	104
5.4	Time-Scale Aggregation and Forecasting	107
5.4.1	Data Frequency Selection	107
5.4.2	Selection of Predictor Variables	110
5.4.3	Prediction Performances	113
5.5	Discussion	116
6	Empirical Results III: Volatility Causality and Interaction with Financial Variables in the EU ETS	120
6.1	Data	120

6.1.1	Computation and Descriptive Analysis of Financial Re-	
	turns	122
6.2	Empirical Analysis	125
6.2.1	Granger’s F-statistic vs. Imbalance Gain: A Compar-	
	ative Analysis	126
6.2.2	Estimating Causal Weights: VAR and Information Im-	
	balance Approaches	128
6.3	Discussion	130
7	Conclusions	133
8	Further Research	137
8.0.1	Causal Graph Reconstruction Algorithm	137
	Bibliography	141
A	Appendix	CLXIV
A.1	Gradient-base Optimisation	CLXIV
A.2	Details on the DII optimisation	CLXVIII
A.3	Mixed frequency forecasting with Gaussian Processes .	CLXX

List of Figures

1.1	EU ETS Price Dynamics (2005-2007) - Phase 1. The plot illustrates the price evolution of the EU ETS Phase 1 future market over time. A 30-day rolling average is applied to smooth short-term fluctuations, providing a clearer trend. The highlighted region marks the period associated with the Global Financial Crisis, indicating its potential impact on price movements.	5
1.2	EU ETS Price Dynamics (2008-2012) - Phase 2. The plot depicts the evolution of EU ETS Phase 2 prices over time. A 30-day rolling average is applied to smooth short-term fluctuations, revealing broader trends. The highlighted region marks the period of Economic Recovery following the Global Financial Crisis and the EU Banking Crisis, illustrating potential market responses to these events.	7
1.3	EU ETS Price Dynamics (2013-2020) - Phase 3. The plot illustrates the evolution of EU ETS Phase 3 prices over time. A 30-day rolling average is applied to smooth short-term fluctuations, highlighting overall trends. The shaded region marks the Post-2012 Transition period and the Brexit event, emphasising their potential impact on price dynamics.	9

1.4	EU ETS Price Dynamics (2021-2030) - Phase 4.	This plot illustrates the evolution of EU ETS Phase 4 prices over time. A 30-day rolling average is applied to smooth short-term fluctuations and highlight broader trends. The shaded region marks the period of the Energy Crisis, driven by the Ukraine-Russia war and escalating geopolitical tensions, reflecting their potential impact on carbon market dynamics.	10
1.5	Sectoral coverage and emissions.	This figure presents the distribution of greenhouse gas (GHG) emissions across economic sectors, categorised using the Common Reporting Format (CRF). It highlights the relative contributions of different sectors to overall emissions, providing insight into their role in the emissions trading system. The figure is sourced from the <i>International Carbon Action Partnership</i> website, <i>Copyright © 2024 International Carbon Action Partnership, accessed 10.07.2024.</i>	16
1.6	ICAP Eurasia ETS Map.	This figure displays the geographical distribution of active and emerging emissions trading systems (ETS) across the Eurasian region. It highlights the expansion of cap-and-trade mechanisms as a policy tool for addressing climate change. The map provides insight into regional participation and collaboration in carbon markets. This image is sourced from the <i>International Carbon Action Partnership</i> website, <i>Copyright © 2024 International Carbon Action Partnership, accessed 10.07.2024.</i>	25

3.1	The Information Imbalance.	This figure illustrates the concept of Information Imbalance for various intuitive relationships between independent variables (X) and a dependent (target) variable (Z). Panel A presents specific datasets, with each dataset corresponding to markers of the same colour in Panel B . The Information Imbalance metric effectively captures different types of relationships: the trivial dependencies in A1 and A2 , the non-linear (quadratic) relationship in A3 , and the multivariate dependency in A4 , where the linear correlation coefficient (ρ) is not applicable. This highlights the advantage of Information Imbalance in detecting complex dependencies beyond simple correlation measures.	45
3.2	The Differentiable Information Imbalance.	This figure demonstrates the application of the Differentiable Information Imbalance (DII) to three spiral datasets with varying levels of noise. The first row displays the datasets in three-dimensional space, while the second row presents two-dimensional projections to highlight the marginal relationships between the target variable z and the predictors x^1 and x^2 . The third row shows the evolution of the weights w_{x^1} and w_{x^2} assigned by the DII to x^1 and x^2 during the gradient descent process across different epochs. The left panels illustrate that the DII considers both x^1 and x^2 as equally important for predicting z , with weights satisfying $w_{x^1} \approx w_{x^2} \approx 1$. In contrast, the middle and right panels show that as the noise level increases between z and x^2 , the DII assigns a smaller weight to x^2 , with w_{x^2} approaching zero.	56

- 3.3 **False Negatives from Linear Causal Discovery.** This figure compares the performance of the VAR/multivariate Granger Causality (GC) model and the Differentiable Information Imbalance (DII) method applied to a stochastic process defined by Equations (3.22), with x^1 and x^2 tested as potential causal variables for the target variable z . The left panel shows the true causal relationships among the three variables: x^1 and x^2 evolve independently, while both x^1 and x^2 influence z . The central bar plot displays the multiplicative weights assigned to the variables by the unrestricted VAR(1) model (blue bars) and the optimized DII ($[\mathbf{x}_0, z_0] \rightarrow z_1$) (orange bars). The right panel shows the F-statistics for each variable, calculated within the multivariate GC framework, alongside the Imbalance Gain computed using the DII approach. 61
- 3.4 **False Positives from Linear Causal Discovery.** This figure compares the performance of the VAR/multivariate Granger Causality (GC) model and the Differentiable Information Imbalance (DII) method applied to the stochastic process described by Equations (3.23), with x^1 and x^2 tested as potential causal variables for the target variable z . The left panel illustrates the true causal relationships: x^2 acts as a common driver for both x^1 and z , highlighting the structure of the underlying process. 62

4.1	Data Processing Workflow. This diagram illustrates the sequential phases involved in selecting and processing research documents, from initial data collection to final inclusion, highlighting key decision points and criteria used throughout the process.	69
4.2	Country Scientific Production. This plot shows the volume of scientific publications related to the EU ETS across various countries. It provides a comparative view of publication output by country. For more detailed figures and data on each country’s contributions, refer to Table 4.1.	71
4.3	Annual Scientific Production (A) and Average Citations per Year (B). This figure shows the trends in scientific publications and average citations per year from 2004 to 2024. Panel (A) illustrates the annual increase in the number of publications, while Panel (B) presents the average number of citations per publication each year.	72
4.4	Core Sources by Bradford’s Law. This plot illustrates Bradford’s Law, which characterises the distribution of research articles across various sources. It highlights the concentration of publications in a small number of core sources, with the remainder distributed across a larger number of less prolific sources.	76
4.5	Most Relevant Authors. This plot highlights the most prolific authors in the field, based on their publication counts. It also showcases the impact of their work, as measured by the total number of citations, providing insight into both productivity and influence.	77

4.6	Most Relevant Authors. This plot highlights the most prolific authors in the field, based on their publication counts. It also showcases the impact of their work, as measured by the total number of citations, providing insight into both productivity and influence.	78
4.7	Most Globally Cited Documents. This plot presents an overview of the top 10 most cited documents worldwide, showing the total number of citations each paper has received across all sources.	79
4.8	Word Cloud and Term Frequency. This figure consists of two parts: on the left, a word cloud that visually emphasises the most prominent themes and terms within the dataset, with larger words indicating higher frequency. On the right, a bar chart displays the frequency of specific terms, providing a quantitative overview of the term distribution across the textual data.	80
4.9	Thematic Map Analysis. This figure presents a thematic map that illustrates the distribution and centrality of key terms across various research clusters. It categorises the research areas into four main types of themes: niche themes, motor themes, basic themes, and emerging or declining themes. This classification allows for a clearer understanding of the relative importance, development, and dynamics of different thematic areas within the literature, highlighting how they evolve and interact over time.	83

5.1	Correlation Analysis - Phase 3. This figure examines the correlations between the daily dataset and EUA. It highlights the variables that provide the greatest insights individually, as identified using the Information Imbalance (II) method, which are shown in orange in Fig. 5.3. Additionally, the dashed bars represent the top five most insightful variables, selected using a greedy algorithm based on the Information Imbalance approach. This variable selection process is further illustrated in Fig. 5.5, corresponding to Phase 3 of the analysis.	97
5.2	Correlation Analysis - Phase 4. This figure examines the correlations between the daily dataset and EUA. It highlights the variables that provide the greatest individual insights, as identified using the Information Imbalance (II) method, which are shown in orange in Fig. 5.4. Additionally, the dashed bars represent the top five most insightful variables, selected using a greedy algorithm based on the Information Imbalance approach. This variable selection process is further detailed in Fig. 5.6, corresponding to Phase 4 of the analysis.	98
5.3	Daily Information Imbalance Analysis - Phase 3. This figure evaluates the informational value of each variable within the dataset. The legend ranks the variables from most to least informative, with darker orange indicating higher informativeness and lighter yellow-orange representing less informative variables. This analysis is specifically presented for Phase 3.	99

5.4	Information Imbalance by Variable in Phase 4. This figure quantifies the relative informational value of individual variables in predicting system dynamics during Phase 4. Variables are ranked from highest to lowest information content (left to right in the legend), with colour intensity indicating their relative importance—darker orange signifies higher information contribution, while lighter yellow-orange denotes lower contribution.	101
5.5	Information Imbalance Analysis of Phase 3 and Phase 4 EUA Price Determinants. This figure illustrates the Information Imbalance, denoted as $\Delta(X_t \rightarrow EUA_t)$, as it moves from the set of predictors to the European Union Allowance (EUA) price. The analysis tracks the changes in imbalance as an increasing number of variables are added to the set X_t . The x-axis labels correspond to each variable added to the set, providing a clear view of their cumulative effect on the EUA price.	105
5.6	Information Imbalance Greedy Optimization of Phase 3 and Phase 4 EUA Price Determinants. This plot illustrates the greedy optimisation process used to select the variables that provide the most significant insights within the Information Imbalance framework. The approach aims to identify the most informative predictors for the European Union Allowance (EUA) price by iteratively adding the variables that maximise the imbalance.	106

5.7	Frequency identification through the Information Imbalance.	
	The figure illustrates the Information Imbalance, denoted by $\Delta(X_t \rightarrow EUA_t)$, highlighting how the transition from the predictor variable set X_t to the European Union Allowance (EUA) pricing structure changes as the number of variables in X_t increases and across different data frequencies.	108
5.8	Variable selection through the Information Imbalance.	
	The plot shows variable selection based on the Information Imbalance, highlighting how the predictor variables are ranked and selected in relation to the European Union Allowance (EUA) pricing structure.	109
5.9	Information Imbalance plane for nowcasting EUA price.	
	The Information Imbalance simultaneously selection of each variable in relation to the target EUA. The analysis is based on weekly data. Variables with the highest informational value are placed at the top and highlighted in deep orange, as indicated in the legend.	110
5.10	Information Imbalance plane for forecasting EUA price.	
	The Information Imbalance 1-day ahead selection of each variable in relation to the target EUA. The analysis is based on weekly data. Variables with the highest informational value are placed at the top and highlighted in deep orange, as indicated in the legend.	111

5.11	Performance of a GP model built on: the 3 most informative variables; all variables; 3 randomly selected predictors (average over 10 replications). In the contexts of nowcasting (on the left) and forecasting (on the right), a scatter plot is presented to illustrate the relationship between the predicted and actual EUA (European Union Allowance) time series. The figure includes both the residuals and the observed time series for a comprehensive comparison.	119
6.1	EUA futures financial returns. European Union Allowance (EUA) futures returns are analysed alongside significant events such as the EU Banking Crisis (2012/11 - 2014/09), Brexit (2016/01 - 2020/12), and the Energy Crisis (2021/03 - 2024/04). Both the returns and their 30-day moving averages are plotted, with key EU ETS phases marked. Periods of increased market uncertainty, reflected by volatility spikes, align with major European crises, which are shaded in light blue for context. . .	123
6.2	The Imbalance Gain as a non-linear alternative to the GC F-statistic. The outcomes derived from the multivariate Granger causality (GC) and DII methods, with EUA as the target variable, are presented. The left panel shows the sorted F-statistics from the GC analysis, calculated using VAR models of order 1. The central panel illustrates the sorted Imbalance Gain (IG) measures. Bold highlights indicate the top four variables among the first ten identified by both methods as having a causal influence on EUA. The right panel provides a scatter plot that correlates the F-statistics with the Imbalance Gain values.	127

6.3 **Beyond VAR causal weight estimation.** The first panel ranks the estimated weights from a VAR(1) model in descending order. The second panel shows the causal weights, arranged in descending order according to the DII measure. Bold highlights emphasize the top 5 weights among the initial 10 that both methods identify as causal for EUA. The third panel features a scatter plot that correlates the VAR weights with those from the DII measure. 128

A1 **Imputation and aggregation using GPs.** In the left panel, a GP is used to impute the GDP time series variable from quarterly to weekly frequency. In the right panel, a GP is used to aggregate the target time series of EUA price from daily to weekly frequency. CLXX

List of Tables

4.1	Publication Frequency and Citations by Country. This table summarizes the frequency of academic publications and the total number of citations for each country. It provides an overview of the publication output and citation impact across different regions.	74
4.2	Most Relevant Venues per Article. This table presents the most prolific academic venues, ranked by the number of articles published. It highlights the key sources contributing to the body of research in this field.	75
5.1	Dataset Description. This table lists 34 distinct time series used in this research, along with their respective sources, organised into seven categories. All time series are collected in euros (EUR).	93
5.2	Descriptive Statistics. This table presents the descriptive statistics for a collection of 34 distinct time series. For each time series, the analysis includes the mean, standard deviation, minimum, and maximum values. Additionally, the values for three specific percentiles are provided to offer a deeper understanding of the data distribution and variability.	96

5.3	Prediction performance (MSE). Mean Squared Error (MSE) for nowcasting ($\delta t = 0$) and forecasting ($\delta t = 1$) GP models built using all predictors (<i>All</i>), a set of 3 random predictors (<i>Rand.</i>) and the set of 3 variables selected via the iterative greedy optimisation of the Information Imbalance (<i>Inf. Imb.</i>). The GP model built on the 3 variables selected via Information Imbalance performs best.	115
6.1	Variable descriptions and data sources. The variables used in the analysis with their full names, abbreviations, codes, categories, and data sources. All data is collected in euros (EUR).	121
6.2	Descriptive statistics of financial returns. Some key descriptive statistics for 35 financial time series, including the mean, standard deviation (STD), minimum (Min), maximum (Max), skewness, kurtosis, and percentile values (25%, 50%, and 75%).	124

Abstract

The European Union Emissions Trading System (EU ETS) plays a critical role in Europe's climate policy, regulating greenhouse gas emissions through a cap-and-trade mechanism. This thesis presents three analyses to deepen the understanding of the EU ETS market. First, a bibliometric review of 367 research papers (2004-2024) examines citation patterns, co-authorship networks, and keyword trends to identify pivotal authors, institutions, and themes, as well as uncovering research gaps. Notable findings include an annual growth rate of 12.99%, with research peaks corresponding to policy changes, and key contributions from Germany, China, and France. The review highlights emerging topics such as carbon pricing and market volatility, with a growing emphasis on financial implications. Despite the progress, several research gaps can be identified in the development and application of non-parametric methods and the impact of macroeconomic variables. The second and third analyses address these gaps using innovative, non-parametric, and model-free methodologies. The Information Imbalance (II) method identifies the most predictive variables for EU ETS price forecasting, revealing that while Phase 3 was influenced by energy indices and commodities, Phase 4 is dominated by financial variables and may be due to economic shocks like the COVID-19 pandemic and the energy crisis. The methodological results demonstrate that the Information Imbalance method effectively selects the most informative weekly timescale for accurate predictions and addresses the mixed-frequency problem by incorporating macroeconomic variables into the analysis. Gaussian Processes (GP) further enhance prediction accuracy in nowcasting and forecasting scenarios. Additionally, this thesis introduces the Differentiable Information Imbalance (DII), a non-linear, model-free approach, based on the Informa-

tion Imbalance, that detects causal relationships within the EU ETS market. DII is compared with the Vector AutoRegression (VAR) model using synthetic data, demonstrating its ability to capture non-linear interactions. A new causal strength metric, Imbalance Gain (IG), is developed within the DII framework to measure the causal influence of predictor variables, similar to the F-statistic. The empirical analysis from 2013 to 2024 reveals significant causal effects of IBEX35 and coal futures on EUA prices, highlighting the intricate linkages between energy, financial, and carbon markets. By leveraging these advanced methodologies to fill research gaps, this work aims to enhance emissions trading systems and support robust climate policies to combat climate change.

Chapter 1

Introduction

1.1 The EU ETS

The European Union Emissions Trading System (EU ETS), established in 2005, is the world's first large-scale carbon market and a cornerstone of the EU's climate policy. It operates under the cap-and-trade principle, where a total emission limit is set, and companies receive or purchase allowances to emit carbon dioxide or other greenhouse gases. Companies that emit less than their allocation can sell surplus credits, while those exceeding their limits must buy additional allowances or face penalties. The system covers sectors such as power generation, heavy industry, and aviation, with plans for further expansion.

The EU ETS aligns economic incentives with environmental goals, encouraging businesses to reduce emissions, invest in low-carbon technologies, and drive innovation in energy efficiency and renewables. It is a key component of the EU's target to reduce GHG emissions by 55% by 2030, in line with the European Green Deal and the Paris Agreement. The system is designed to gradually tighten the emission cap to ensure cost-effective emissions

reductions over time.

The EU ETS also serves as a benchmark for other regions considering similar market-based climate policies, influencing global discussions on addressing climate change.

In both empirical papers included in this thesis, we use the front-month futures contract (MO1 COMDTY) from Bloomberg as the primary data source for EUA prices. MO1 COMDTY is a continuous rolling contract that always references the nearest-expiration EUA futures contract (e.g., CO2Z23 for December 2023, then CO2H24 for March 2024 upon rollover). Since it tracks the most liquid contract at any given time, it serves as a reliable benchmark for price analysis.

Front-month contracts are the closest to expiration and typically exhibit the highest trading volume, minimizing bid-ask spreads and ensuring price efficiency. Although specific contracts (e.g., CO2Z23) refer to a fixed expiry month, they share identical price dynamics with MO1 COMDTY while active as the front-month contract. This is because MO1 COMDTY simply mirrors the nearest expiry until the roll date, after which it shifts to the next contract without disrupting the price continuity used in analysis [21].

In the context of the EU ETS market, the spot price represents the immediate cost of purchasing EU Allowances (EUAs) for near-term delivery, reflecting current supply and demand conditions. In contrast, the futures price (as captured by MO1 COMDTY) reflects the agreed-upon price today for future delivery, incorporating expectations on regulatory changes, economic trends, and market risk premia. The data used in this thesis aligns with standard market practices, sourced from Bloomberg to ensure consistency with observed EU ETS dynamics.

This chapter introduces the EU ETS, explaining its structure, objectives,

and significance in international climate policy.

1.2 History and Development of the EU ETS

1.2.1 Early Foundation and Policy Context

The idea of emissions trading emerged in the 1960s and 1970s, when environmental concerns began to grow alongside industrialisation and increased pollution [68]. Early economic approaches to environmental policy focused on using market-based solutions to address environmental problems, instead of relying only on government regulations.

The concept of *cap-and-trade* systems, where a pollution cap is established and emission permits are allocated or traded, was influenced by ideas of economic efficiency and the belief that market forces could incentivise companies to reduce pollution at the lowest possible cost [72]. The theoretical foundation for emissions trading was laid in the 1960s by economists such as Ronald Coase in [45], who argued in his seminal work *The Problem of Social Cost*, that market-based solutions could be used to allocate resources, or pollution rights, efficiently.

In the 1970s and 1980s, the concept was further developed as scholars and policy makers explored how to address the increasing levels of air pollution in industrialised countries. The first major application of emissions trading was in the US during the 1990s with the implementation of the Acid Rain Programme under the Clean Air Act Amendments [62, 76]. This programme allowed for the trade of sulphur dioxide (SO₂) emissions, which contributed to acid rain.

The success of this programme demonstrated the potential of emissions trading as an effective tool to reduce pollution and provided a model for

future carbon trading systems [147]. The concept of emissions trading, particularly in relation to CO₂, gained prominence on the global stage in the early 1990s. Following the 1992 Earth Summit in Rio de Janeiro, where the United Nations Framework Convention on Climate Change (UNFCCC) was established, the international community began to seriously address the issue of climate change [160].

The Kyoto Protocol, adopted in 1997, set binding targets for industrialised countries to reduce their GHG emissions [46]. One of the key mechanisms outlined in the Kyoto Protocol to help countries meet their emissions targets was the use of market-based instruments as emissions trading [94].

The idea of a carbon market, where countries or companies could trade carbon allowances to meet their emission targets, gained significant traction. The European Union, already active in international climate negotiations, sought to implement an emissions trading system that would allow member states to meet their Kyoto goals in a cost-effective manner.

The EU's policy on climate change was formalised through the adoption of the Burden Sharing Agreement (BSA) in 1997, which established individual emission reduction targets for EU member states for the period 2008-2012, in accordance with the Kyoto Protocol. The EU's commitment to achieving these targets led to the creation of the EU ETS as a key policy tool. The EU ETS was designed to ensure that emissions reduction across the EU could be achieved at the lowest possible cost by providing economic incentives for companies to reduce emissions.

The legislative framework for the EU ETS was first proposed in 2001 by the European Commission, and in 2003, the European Parliament and the Council of the EU adopted Directive 2003/87/EC, establishing the EU ETS [136]. The first trading period began in 2005, and the system initially




Figure 1.1: **EU ETS Price Dynamics (2005-2007) - Phase 1.** The plot illustrates the price evolution of the EU ETS Phase 1 future market over time. A 30-day rolling average is applied to smooth short-term fluctuations, providing a clearer trend. The highlighted region marks the period associated with the Global Financial Crisis, indicating its potential impact on price movements.

covered the power generation and heavy industrial sectors [47]. This period was a crucial moment in the development of carbon markets, marking the EU’s leadership in establishing a market-based approach to climate policy.

1.2.2 Phase 1 (2005-2007): Pilot Phase

The first phase of the EU ETS, which spanned from 2005 to 2007, was designed primarily as a pilot phase to establish the trading system and allow market participants to adjust to the mechanics of emissions trading.

During this period, the EU ETS covered approximately 12,000 power and industrial installations across the EU, representing about 45% of the total EU greenhouse gas emissions. The emission cap set for this phase was based on the anticipated emissions reductions that the participating sectors could

achieve [75].

However, the first phase of the EU ETS was fraught with issues that hindered its effectiveness in driving emissions reductions. The most notable problem was the overallocation of allowances. Excessive allowances were allocated to companies than they actually needed, in part due to the lack of historical emissions data and uncertainties about future emissions.

This overallocation led to a situation in which the carbon price was driven down by an oversupply of credits, undermining the market's ability to create a financial incentive for companies to reduce emissions [103]. The carbon price fluctuated modestly, with a downward trend throughout the phase, culminating in a significant drop toward the end of 2007, as the market adjusted to the excessive surplus of allowances; see figure 1.1.

The lessons learnt in Phase 1 were crucial to refining the EU ETS. It became clear that a more stringent cap on emissions was necessary to create a viable carbon market. The excess allowances and low carbon price prompted the European Commission to reconsider its approach, paving the way for reforms in the subsequent phases.

1.2.3 Phase 2 (2008–2012): Expansion and Alignment with Kyoto Protocol

Phase 2 of the EU ETS, which coincided with the commitments of the EU under the Kyoto Protocol, ran from 2008 to 2012. This phase marked a significant change, with stricter emission limits and an expansion of the system's coverage. The total emission cap was reduced by approximately 6.5% compared to 2005 levels, in accordance with the objectives of the EU under the Kyoto Protocol [50]. This tightening was designed to provide a stronger price signal and drive emission reductions across the sectors covered by the




Figure 1.2: **EU ETS Price Dynamics (2008-2012) - Phase 2.** The plot depicts the evolution of EU ETS Phase 2 prices over time. A 30-day rolling average is applied to smooth short-term fluctuations, revealing broader trends. The highlighted region marks the period of Economic Recovery following the Global Financial Crisis and the EU Banking Crisis, illustrating potential market responses to these events.

system.

In this phase, the EU ETS was extended to cover additional sectors, including domestic aviation and certain industrial sectors such as cement and chemicals. This expansion increased the overall potential for the reduction of emissions from the system. Furthermore, the allocation method was harmonised between member states, and a higher proportion of the allowances was auctioned rather than free allocation. This was seen as a step towards improving the efficiency and fairness of the market [77].

However, Phase 2 still faced challenges, particularly due to the global financial crisis that began in 2008. The economic downturn led to a significant reduction in industrial output, which in turn caused emissions to fall. As a result, the demand for allowances decreased, leading to a glut of allowances in

the market and causing the carbon price to drop significantly [134]. By 2009, the carbon price had fallen to around 10€per tonne, despite the tightened cap. This volatility highlighted the vulnerability of the EU ETS to external economic factors, which was a key lesson for future reforms [114].

Despite these challenges, Phase 2 established important groundwork for future development. The system mechanisms were tested on a larger scale and the phase demonstrated the need for market stabilisation measures, which would be introduced in later phases.

The carbon price in Phase 2 fluctuated significantly due to the global financial crisis. Prices started at around 25€per tonne in early 2008, then dropped sharply to 10€per tonne by 2009 as industrial activity slowed. Prices averaged between 10€and 15€per tonne throughout the phase, with a significant drop toward the end, see figure 1.2.

1.2.4 Phase 3 (2013–2020): System Strengthening

Phase 3, which ran from 2013 to 2020, represented a major reform of the EU ETS. This phase aimed to address the issues of over-allocation and price volatility that plagued the earlier phases. Several key reforms were introduced during this period, starting with the shift to a centralised cap for the entire EU rather than allowing each member state to set its own cap. This created a more uniform approach across the EU and increased the overall efficiency of the system.

Alongside these changes, the Market Stability Reserve (MSR) was introduced in 2019 to absorb surplus allowances when demand was low and release them when demand was high [48]. This measure aimed to address the issue of surplus allowances that had caused price volatility in earlier phases and stabilise the carbon market.




Figure 1.3: **EU ETS Price Dynamics (2013-2020) - Phase 3.** The plot illustrates the evolution of EU ETS Phase 3 prices over time. A 30-day rolling average is applied to smooth short-term fluctuations, highlighting overall trends. The shaded region marks the Post-2012 Transition period and the Brexit event, emphasising their potential impact on price dynamics.

In terms of price dynamics, Phase 3 saw a gradual recovery from the low prices of earlier phases, but the market remained unstable for much of the period. Carbon prices began to increase steadily after the MSR was introduced in 2019, as the reserve began to absorb excess allowances and reduce oversupply on the market. At the end of Phase 3, carbon prices had reached around 25€ per tonne, marking a significant increase compared to earlier phases; see figure 1.3.

The introduction of the MSR was a turning point, as it helped stabilise the carbon market and provided a clearer price signal for companies, encouraging them to invest in cleaner technologies. However, the phase still faced criticism for not doing enough to drive emissions reductions at the pace needed to meet long-term climate goals.

In addition, the New Entrant Reserve (NER 300) was introduced to sup-




Figure 1.4: **EU ETS Price Dynamics (2021-2030) - Phase 4.** This plot illustrates the evolution of EU ETS Phase 4 prices over time. A 30-day rolling average is applied to smooth short-term fluctuations and highlight broader trends. The shaded region marks the period of the Energy Crisis, driven by the Ukraine-Russia war and escalating geopolitical tensions, reflecting their potential impact on carbon market dynamics.

port the deployment of innovative low-carbon technologies, such as carbon capture and storage (CCS) and renewable energy projects [49].

1.2.5 Phase 4 (2021–2030): Expansion and Ambition

Phase 4, which began in 2021 and will run until 2030, represents the most ambitious phase of the EU ETS that has been rolled out to date. This phase aligns with the European Green Deal and the Fit for 55 package, which set the target of reducing EU emissions by 55% by 2030 compared to 1990 levels. As part of this effort, the EU ETS emission cap will decrease by 2.2% annually, a faster rate of reduction compared to previous phases [54, 146]. Furthermore, the share of emission allowances sold through auctions is increased, reducing the allocation of free allowances, which aims to eliminate concerns regarding

overallocation and inefficiencies that were present in earlier phases [52].

Phase 4 also includes several new developments and improvements. The Carbon Border Adjustment Mechanism (CBAM) was introduced to address carbon leakage by imposing a carbon price on imports from countries with less stringent climate policies, ensuring a level playing field for domestic and foreign producers [57]. The issue of carbon leakage first emerged as a policy concern with the launch of the EU Emissions Trading System (EU ETS) in the early 2000s, which initially relied on free allowances to prevent industries from relocating to regions without carbon costs. However, free allowances were found to weaken decarbonisation incentives, leading to the development of CBAM as a more effective solution [161, 11].

Adopted as part of the EU's Fit for 55 package, CBAM targets high-emission sectors such as steel and cement, requiring importers to report embedded emissions during its transition phase (2023-2026) and purchase CBAM certificates from 2026 onwards [57]. By replacing free allowances, CBAM aligns carbon costs across borders, reducing leakage risks while maintaining incentives for decarbonisation. Future expansions of CBAM to include additional sectors, such as chemicals, could further enhance its effectiveness in addressing carbon leakage and supporting global decarbonisation efforts [28].

Another major aspect of Phase 4 is the expansion of the EU ETS to additional sectors. While the system has already covered energy and industry, it now includes maritime transport [53].

In terms of price dynamics, the introduction of more stringent emission reductions, coupled with the strengthened MSR, has driven carbon prices higher since 2021. The carbon price has exceeded 50€/per tonne in early 2023, reflecting the market's growing expectations of tighter supply and stronger

demand for allowances in the future. The EU's increased ambition for emissions reductions and the expansion of the carbon market are expected to continue driving prices upward in the coming years, creating a stronger price signal for companies to invest in low-carbon technologies and further accelerate the EU's decarbonisation efforts.

Carbon prices in Phase 4 have been marked by a significant rise in price, particularly after the market's recovery from the COVID-19 pandemic. By 2023, prices reached levels above 50€/per tonne, reflecting both the tightening of the emissions cap and increasing market confidence in the long-term sustainability of the EU ETS, see figure 1.4.

1.3 Market Organisation and Compliance

1.3.1 How the Market Works

The EU ETS operates as a cap-and-trade system, where a cap is established on the total amount of greenhouse gas emissions allowed in the covered sectors. This cap gradually decreases over time, leading to a reduction in emissions across the EU.

The market functions by allocating a certain number of European Union Allowances (EUAs) to companies, which correspond to the right to emit a specific amount of carbon dioxide or equivalent greenhouse gases. These allowances are tradable, meaning companies that reduce their emissions below their allocated cap can sell their surplus allowances to those who exceed their emission limit.

The cap-and-trade mechanism operates by imposing a limit on total emissions within the market. Companies are incentivised to reduce their emissions to avoid purchasing additional allowances, creating a market where unused

allowances can be traded for financial gain. This structure encourages emissions reductions while allowing flexibility in how companies achieve their targets. If companies reduce emissions beyond their allocated amount, they can sell their excess allowances, making emission reduction financially attractive. The trading element of the system creates a flexible and market-driven mechanism for emissions reductions, where the price of allowances is determined by supply and demand. If the supply of allowances is high relative to demand, prices will fall and vice versa.

In the EU ETS, the primary market involves the initial distribution of EU allowances (EUA) through auctions conducted by platforms such as the European Energy Exchange (EEX) or the Intercontinental Exchange (ICE). These auctions are the first point of sale for EUAs, where compliance entities and financial institutions purchase allowances directly. In contrast, the secondary market encompasses trading on exchanges (e.g., ICE, EEX) and over-the-counter (OTC) markets, where bilateral trades occur between parties. OTC trading offers flexibility for customised transactions but is less transparent than exchange-based trading, with trades often reported to registries to ensure market integrity.

Indeed, allowances in the EU ETS are either allocated for free or sold through auctions. The method of allocation varies between different sectors and phases of the programme. Initially, many allowances were assigned for free based on historical emission levels, but over time the system has increasingly relied on auctions as a method of allocation [56]. Auctioning helps avoid overallocation and ensures that allowances are priced in a way that reflects their scarcity, providing a clear price signal to market participants. For industries exposed to significant competition from non-EU countries, free allocation continues to be an option, but the overall share of free allowances

is decreasing.

The trading of allowances is conducted on various exchanges, such as the European Energy Exchange (EEX), where companies can buy and sell allowances. The price of allowances is determined by market forces, with supply and demand balancing out the cost of carbon emissions. Over time, market participants, including companies, traders and financial institutions, have developed sophisticated trading strategies, creating a dynamic market for carbon allowances.

1.3.2 Compliance Mechanism

To ensure the integrity of the EU ETS, the system includes a rigorous compliance mechanism that involves the monitoring, reporting, and verification (MRV) of emissions. The purpose of the compliance mechanism is to ensure that all participating entities meet their emission reduction obligations [60].

Under the EU ETS, companies are required to monitor their emissions and report them annually. Emissions are tracked transparently and the data are verified by independent third-party auditors. Companies must report their total emissions for the previous year and provide data on their operations to ensure accuracy. This reporting is crucial for the European Commission to determine whether companies comply with their allocated allowances and achieve their emission reduction targets.

The MRV process is highly regulated and standardised across the EU to ensure uniformity and consistency. Verified emission data must be submitted by a fixed deadline, and failure to comply with these deadlines can result in penalties. The reporting system also ensures that any discrepancies or inaccuracies are identified and corrected before they cause significant compliance problems.

Non-compliance with the EU ETS regulations is met with severe financial penalties to ensure that companies are incentivised to meet their obligations. If a company does not surrender enough allowances to match its emissions, it faces a penalty of 100€ per ton of excess emissions, in addition to the obligation to cover the shortfall with allowances in the following year. This heavy penalty serves as a strong deterrent against non-compliance, and penalties are structured to be higher than the cost of purchasing allowances, ensuring that companies have a strong incentive to comply with the rules [60].

1.3.3 Market Oversight

Effective market oversight is essential for the integrity of the EU ETS and to prevent potential manipulation or fraud. Several regulatory bodies are involved in the supervision of the EU ETS to ensure its proper functioning and maintain confidence in the system.

The main regulatory body overseeing the EU ETS is the European Commission, which sets the overall policy and ensures that the system is aligned with the EU's climate objectives. The Commission is responsible for overseeing the setting of caps, the allocation of allowances, and the overall functioning of the system, including any changes in policy or reform measures.

In addition to the European Commission, individual member states are also responsible for regulating and overseeing the activities of companies within their jurisdictions. National authorities monitor companies' emissions, ensure compliance with reporting and verification requirements, and enforce penalties for non-compliance. In many cases, national regulatory bodies cooperate with each other to ensure consistency and effective enforcement across borders.

The European Securities and Markets Authority (ESMA) also plays an

Overall GHG emissions by CRF category (MtCO₂e)




Figure 1.5: **Sectoral coverage and emissions.** This figure presents the distribution of greenhouse gas (GHG) emissions across economic sectors, categorised using the Common Reporting Format (CRF). It highlights the relative contributions of different sectors to overall emissions, providing insight into their role in the emissions trading system. The figure is sourced from the *International Carbon Action Partnership* website, *Copyright © 2024 International Carbon Action Partnership, accessed 10.07.2024.*

important role in the supervision of the financial aspects of the EU ETS [81]. ESMA ensures that the trading of carbon allowances is conducted in a fair and transparent manner, and monitors the market for potential signs of manipulation, fraud, or other irregularities [61]. This oversight is crucial to maintaining market integrity and ensuring that companies have confidence in the trading system.

To prevent market manipulation, several safeguards have been put in place. These include surveillance mechanisms to monitor trading activity and identify suspicious patterns that might indicate price manipulation or other illicit behaviour.

1.4 Sectoral Coverage

The EU ETS covers a wide range of sectors, but is primarily focused on industries that produce significant greenhouse gas emissions. The power sector and the heavy industry are the largest sectors covered, accounting for the

majority of emissions within the system. These sectors are highly energy-intensive and are seen as key to achieving substantial emissions reductions across the EU. The inclusion of these sectors ensures that the EU can effectively target and reduce the largest sources of emissions.

In addition to power generation and industry, the aviation sector was included in the EU ETS starting in 2012. The inclusion of aviation in the EU ETS marked a significant expansion of the system, as it addressed the growing emissions from air travel [59]. Airlines operating in the EU must surrender allowances based on their emissions, with the number of allowances allocated to them determined by the total number of flights they operate within the EEA. This inclusion was initially controversial, particularly with non-EU countries and international aviation organisations, but it reflects the EU's ambition to address emissions from all major sectors.

Other sectors included in the EU ETS are those involved in industrial installations that emit large amounts of greenhouse gases, such as the production of cement, steel, chemicals, and paper cited in the previous section. These sectors have been integral to the EU ETS since its inception, as their emissions are highly concentrated and difficult to reduce without significant changes in technology or operational processes.

Certain exemptions exist in the EU ETS, particularly for smaller emitters or entities that are subject to stringent international competition. These industries may struggle to compete with counterparts in countries with less stringent environmental policies. To address this concern, the EU has explored mechanisms such as carbon adjustment on the border, which would impose a carbon price on imports from countries without equivalent emission regulations, ensuring a level playing field [77].

1.5 Future Directions

1.5.1 The EU ETS 2

As the European Union strengthens its climate policies and works towards its ambitious emission reduction targets, the upcoming EU ETS 2 marks a key development. This stand-alone market will include new sectors such as road transport and buildings, which were previously excluded [58]. Their inclusion is now considered important for achieving the climate objectives of the EU, particularly the goal of achieving net zero emissions by 2050. The integration of these sectors into the EU ETS 2 is part of a broader effort to enhance the effectiveness of carbon pricing and ensure that emissions from all significant sources are addressed to meet the EU's ambitious climate commitments.

The transport sector, particularly road transport, is a major source of emissions, and its inclusion in the EU ETS will help address one of the largest remaining sources of greenhouse gas emissions in Europe.

Road transport emissions are mainly from passenger vehicles, freight trucks, and buses, making this a complex sector to regulate. It requires measuring emissions in a large number of vehicles, many of which are privately owned, and integrating them into a trading system. Similarly, the building sector is a substantial contributor to emissions, primarily from heating and cooling systems. The inclusion of buildings in the EU ETS presents a set of challenges, particularly in terms of measuring emissions from millions of individual buildings. Policymakers are considering a combination of approaches, such as requiring large buildings or landlords to purchase carbon allowances based on their energy use. However, there are also concerns about the potential impact on consumers, especially low-income households, so complementary policies, such as targeted subsidies or energy efficiency

programmes, will be necessary to mitigate any social or economic impacts.

The implementation of these new sectoral expansions is not without its challenges. There are concerns about how to measure and monitor emissions effectively in these sectors and how to ensure that the inclusion of new sectors does not create excessive administrative burdens.

In sectors such as road transport and buildings, the primary challenge lies in the indirect economic effects of carbon pricing, particularly the cost-pass-through mechanism. Unlike traditional industries, where carbon costs are more directly reflected in product prices, the complexity of supply chains and delayed price adjustments in these sectors can obscure the link between carbon pricing and higher costs for consumers and businesses. This lack of transparency creates uncertainty in market responses and undermines the effectiveness of climate policies [31].

1.5.2 Integration with International Carbon Markets

Another key direction for the future of the EU ETS is its integration with other international carbon markets. As climate change is a global problem, the EU recognises that carbon pricing cannot be isolated to Europe alone. One of the most significant opportunities for expansion is the potential linking with other ETS systems, such as the Swiss ETS.

Switzerland has been in discussions with the EU for years about linking their systems, and in 2020, the Swiss ETS became officially linked with the EU ETS, enabling the exchange of allowances between the two systems [80]. This was a milestone for the EU ETS, as it laid the groundwork for future integration with other national or regional carbon markets.

Looking beyond Europe, the EU ETS may potentially link with other carbon markets globally, especially in regions increasingly adopting emissions

trading to meet climate targets. For example, countries such as China, Japan, and South Korea, which have developed or are in the process of developing carbon markets, could potentially link to the EU ETS in the future.

Moreover, the EU ETS will play a significant role in international climate agreements, particularly in the context of the Paris Agreement and its Article 6, which provides mechanisms for countries to cooperate in reducing emissions [142]. Article 6 allows for the creation of international carbon markets through which countries can trade emission reductions to meet their national climate targets. The EU has been a strong advocate for the development of such markets and has worked to ensure that emissions trading is incorporated into the global climate architecture. In this context, the EU ETS could potentially serve as a model for other countries, offering valuable experience and lessons in the design and operation of carbon markets.

Despite the potential benefits of international integration, there are also significant challenges, including concerns about the uniformity of standards, the quality of carbon credits, and the fairness in how carbon pricing impacts different economies. Countries with lower carbon prices might seek to avoid stricter regulations by purchasing credits from regions with lower costs, potentially undermining the overall effectiveness of the system, cherry-picking effect.

The EU will need to address these concerns through careful coordination, transparent governance, and the establishment of clear rules for linking carbon markets to ensure that the expansion of the EU ETS into global markets is both effective and equitable.

1.6 Thesis Motivation and Objectives

The primary goal of this thesis is to perform a comprehensive bibliometric and empirical analysis of the literature of the EU ETS market, spanning 2004 to 2024, to assess the extent of scientific research on the EU ETS market and improve the comprehension of the dynamics of carbon pricing, as well as the factors driving its fluctuations.

By analysing articles indexed in the Scopus database, this study aims to identify key trends in the academic literature, track the evolution of research focus, and uncover significant areas of development within EU ETS research. This includes evaluating the volume of research, identifying influential authors, institutions, and countries, and performing citation analyses to highlight the most impactful studies. In addition, the thesis explores the evolution of research themes, examining emerging trends and areas of consensus or debate. The creation of a thematic map will further reveal niche, motor, and emerging topics, providing information on future directions of research.

On the empirical side, the study seeks to provide a deeper understanding of the factors driving carbon price fluctuations, with particular emphasis on identifying the exogenous variables that influence the price of European Union Allowances (EUA).

In addition, the study aims primarily to explore the dynamics of carbon pricing, with a focus on identifying the factors that drive its fluctuations. This exploration aims to improve the management of market risk for participants, enabling them to make more informed decisions.

Furthermore, the research aims to deepen the understanding of strategic tools available to European policymakers to regulate carbon prices, thereby enhancing the effectiveness of the EUA in the broader effort to combat cli-

mate change.

Moreover, the study introduces a novel non-parametric approach, the Information Imbalance (II) measure [89], to identify key external variables influencing EUA prices. The II is used to compare the informative variables of a dataset of different financial variables, spanning from January 2014 to April 2023, across two different phases, examining whether there is a difference in price determinants between Phase 3, characterised by relatively stable permit prices, and Phase 4, which is marked by price volatility due to COVID-19 and energy crises.

Methodologically, the study proposes combining II with Gaussian Process regression (GP) to integrate mixed-frequency data at the most informative time scale, creating models for forecasting or nowcasting. Additionally, it demonstrates how II can help select a small set of highly informative variables for these predictive models.

Indeed, another objective of this thesis is to enhance the comprehension of how energy and commodity markets are connected by introducing an innovative technique aimed at identifying non-linear causal relationships within financial time series data. This technique leverages the II metric and uses its differentiable variant called Differential Information Imbalance (DII) [168, 5]. The DII offers a model-independent novel alternative, providing a contrast to traditional methods like Vector Autoregression (VAR) models and Granger causality tests. To evaluate the efficacy of DII, its application is compared with that of a conventional VAR model using a dataset of financial returns spanning from January 2013 to April 2024, encompassing a diverse range of assets and indices pertinent to energy markets. A specific aim of this study is to investigate the causal mechanisms driving the EUA market, examining how external variables influence the financial return dynamics of EUA.

The influence of each variable is assessed on the basis of the improvement it brings to EUA prediction accuracy, in contrast to scenarios where the variable is absent. This is assessed via the F-statistic in multivariate Granger causality along with the newly introduced Imbalance Gain (IG) within the DII framework.

This thesis primarily provides an empirical contribution by applying innovative methodologies to the EU ETS market, a context where these methods have not been previously explored. Although the methodological framework is not newly developed, its application in this specific setting offers novel insights and extends the existing literature.

1.7 Thesis Organisation

The structure of this thesis is organised into several chapters, each focussing on a specific analysis carried out following the European Union Emissions Trading System (EU ETS) and its dynamics.

Chapter 1, **Introduction**, begins by offering a comprehensive overview of the EU ETS, its objectives, and its evolution over time. This chapter also establishes the motivation for the thesis and outlines its main research objectives. Following this, the structure of the thesis is detailed to guide the reader through the subsequent chapters.

In Chapter 2, **Literature Review**, the existing body of research surrounding the EU ETS is examined. This chapter addresses key topics such as the state-of-the-art of literature reviews, the pricing mechanisms and forecasting models related to the EU ETS, as well as the volatility and causality aspects that have been explored in the literature.

Chapter 3, **Methods and Models**, outlines the methodological ap-

proaches used in the analysis of the EU ETS market. This chapter introduces various analytical tools and models used throughout the research, such as linear correlation, copulas, the Information Imbalance (II), the Differentiable Information Imbalance (DII) and Gaussian Process regression (GP).

In Chapter 4, **Empirical Results I: Bibliometric Insights on the EU ETS Market**, the results of the bibliometric analysis of the EU ETS are presented. This chapter details the data collection process, the criteria for selecting relevant literature, and provides both descriptive analysis and the conceptual structure derived from the data in order to analyse trend, networks and highlight fundamental contributions and research gaps.

Chapter 5, **Empirical Results II: Price Determinants and Forecasting in the EU ETS**, focusses on the factors that determine prices within the EU ETS and the forecasting models used to predict future market trends, using a novel non-linear and non-parametric metric called Information Imbalance (II). Includes both the empirical and methodological results achieved in this study.

Chapter 6, **Empirical Results III: Volatility Causality and Interaction with Financial Variables in the EU ETS**, focusses on the proposal of a new non-parametric, model-free methodology for detecting non-linear causal relationships in financial returns data, the Differentiable Information Imbalance (DII) and the related causal metric, the Imbalance Gain (IG). A comparison with the Vector AutoRegression (VAR) model is provided, along with methodological insights based on the analysis of two sets of synthetic data. Furthermore, the empirical findings on the EU ETS financial returns causal relationships are presented.

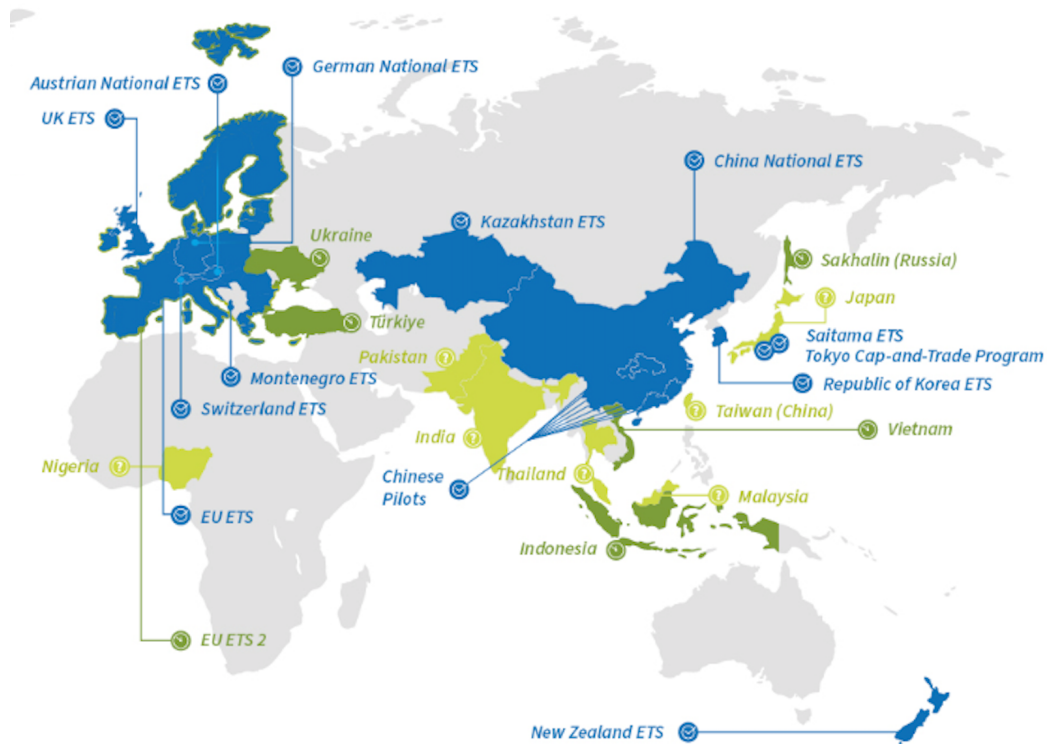


Figure 1.6: **ICAP Eurasia ETS Map.** This figure displays the geographical distribution of active and emerging emissions trading systems (ETS) across the Eurasian region. It highlights the expansion of cap-and-trade mechanisms as a policy tool for addressing climate change. The map provides insight into regional participation and collaboration in carbon markets. This image is sourced from the *International Carbon Action Partnership* website, Copyright © 2024 International Carbon Action Partnership, accessed 10.07.2024.

Chapter 2

Literature Review

2.1 Exploring the EU ETS Market: Bibliometric and Review Literature

The following section highlights key scientific contributions that have shaped the research landscape of the EU ETS, focusing on market performance, pricing mechanisms, allowance allocation, and policy implications. These studies provide valuable information on the evolution of the scheme and areas for further investigation.

A comprehensive review of the literature on economics and econometrics on the EU ETS is provided by [39], with a focus on the second phase of the market. This study summarises key findings on the CO₂ price mechanism, highlighting the influence of policy measures, market fundamentals, and macroeconomic factors. It also highlights gaps in understanding how carbon prices adjust during global recessions, calling for enhanced forecasting techniques and policy-support tools for stakeholders.

Expanding on this, [101] evaluates the performance of the EU ETS us-

ing five criteria: effectiveness, efficiency, coherence, EU added value, and relevance. Covering 250 publications, this review addresses significant issues such as allowances over-allocation and investment leakage. The authors identify 14 research gaps within three thematic areas: methodological, data, and coverage. They stress the importance of standardised baselines, improved data collection on investment leakage, and a wider exploration of emerging trends.

The allocation of allowances and its impact on market behaviour is a key focus of [174], which examines the overallocation in Phase 1 and its adverse effects on carbon prices. The study investigates the complexity of the price, highlighting how energy costs and weather conditions influence the prices of EU ETS allowances. Methodological gaps are identified, particularly in understanding price mechanisms and equilibrium. Future research is suggested to explore how rising carbon prices could affect energy companies' emissions and operations.

In [107], a bibliometric analysis of carbon pricing within the EU ETS reveals rapid growth in the field, with China and the USA as leading contributors. The study identifies Energy Policy as the most productive journal and highlights four main research areas: carbon price dynamics, design of price mechanisms, policy analysis, and the effects of carbon pricing. The analysis categorises the factors that influence policy, emissions levels, energy prices, weather and climate, and financial markets. Key challenges include low carbon prices, oversupply of allowances, and structural issues in the energy sector. The study emphasises the need for more quantitative analyses and improved market mechanisms to strengthen carbon trading systems.

Similarly, [165] explores carbon emission reduction mechanisms across disciplines like environmental science and economics. Influential researchers

such as Pizer W. and Nordhaus W. are identified, and major journals include Energy Policy and the Journal of Environmental Economics and Management. Research topics include uncertainty-based decision making, comparisons of carbon taxes and trading systems, and practical design of carbon markets. The authors note a lack of extensive international collaboration in this domain.

A detailed scientometric review by [106] examines the origins, mechanisms, and future trajectories of the European carbon market. The study underscores the importance of the ETS framework in addressing climate change cost-effectively. A key challenge is the allocation of carbon allowances. Although free allocation based on historical emissions was appropriate during the transitional phases (Phases 1 and 2), future systems are likely to adopt methods such as benchmarking and auctioning to improve efficiency.

The influence of EU ETS on the adoption of low carbon technologies is evaluated by [158], which reviews 22 empirical studies split into econometric and qualitative analyses. Econometric studies often find a positive correlation between EU ETS and low-carbon patenting or R&D investments among regulated firms, although some studies report no significant impact on emissions intensity. Qualitative studies suggest that while the EU ETS has spurred some innovation, energy costs and other factors often drive technological change. The findings highlight the modest but important role of the EU ETS in advancing low-carbon technologies and the regulatory framework's impact on firms' investment decisions.

Challenges related to overallocation of allowances are explored in depth by [120], who examine its reduced impact on emissions during economic downturns. Effective monitoring and favourable carbon pricing have led to some emission reductions, but further research and data are essential to refine

methodologies and inform the design of future trading systems. These findings underscore the need for an ongoing investigation into the relationship between free allocation mechanisms and economic performance.

By consolidating these insights, the scientific community has significantly advanced our understanding of the EU ETS, while also identifying critical gaps to guide future research and policy innovation.

2.2 EU ETS Pricing Dynamics and Forecasting Methods

A key challenge in understanding EU ETS price dynamics lies in identifying the most informative factors driving market fluctuations. To address this, we introduce the Information Imbalance measure, a non-parametric approach that quantifies the influence of different variables on carbon price movements. This method provides a fresh perspective on the evolving role of macroeconomic, financial, and energy-related factors, offering new insights into market behaviour and forecasting accuracy.

To construct this literature review, we began by analysing the foundational work on Information Imbalance by [89], as well as the study by [17] on carbon market dynamics. From there, we expanded our review to include papers that cite and are cited by these key studies, tracing the development of methods used to interpret EU ETS price behaviour.

The literature on the EU ETS offers a diverse array of methods and models that serve as interpretive tools, aiding in the navigation of the challenges and opportunities associated with the energy transition.

In their analysis, [18] examined the balance of supply and demand in the carbon market, identifying policy and regulation as key factors influenc-

ing supply dynamics and noting that carbon prices are determined by the interaction of allowance supply and demand.

Subsequent studies have found that fuel prices, particularly those of oil, natural gas, and coal, are significant determinants of permit prices (e.g., [128], [3], [114], [102], [38], [65], [30]).

These studies identify economic activity as a key driver of European Union Allowance (EUA) prices, often measured using stock market indices as proxies.

In [159], natural gas is identified as a key factor influencing EUA prices. The authors also observed a dependency between the crude oil sector and the European emissions market. In addition, rising electricity prices are related to higher allowance prices.

In [1] the determinants of the EUA permit price are studied through a time series analysis. Companies, by producing goods and reducing emissions, influence the price of permits. Empirical data from 2005 to 2010 show a strong relationship between the price of permits and fundamentals such as electricity, gas and coal prices, confirming the importance of these variables in determining the price of EUAs.

For example, [97] used a quantile regression approach to examine how fluctuations in oil, natural gas, coal and electricity prices affect the distribution of EUA price. The main conclusions are that an increase in crude oil prices causes a significant drop in EUA price.

Gas prices negatively affect the EUA at low levels, but positively when those are high. Electricity positively impacts the high end of the distribution, while coal prices have a negative influence on the price of EUA.

In [98], a Bayesian Structural Vector Autoregressive Model (BSVAR) was used to examine the short-term dynamics of CO₂ emission prices in response

to changes in oil, coal, natural gas and electricity prices. The results show that a positive shock in crude oil prices initially raises the CO₂ allowance prices, but then leads to a negative impact. An unexpected increase in natural gas prices reduces CO₂ prices, while a positive shock in coal prices has minimal and statistically insignificant effects. However, when electricity prices are excluded from the BSVAR framework, coal prices show a significant positive effect on CO₂ allowances. A positive shock in electricity prices negatively impacts CO₂ allowance prices. The study also highlights the persistent effects of energy price shocks on CO₂ prices, with the most pronounced impacts occurring six months after the shock, particularly for shocks to natural gas and crude oil prices. However, other authors emphasise the time-varying relationship between carbon prices and these variables.

[166] and [108] use a VAR model to highlight the significant impact of Brent oil yields on price returns and volatility of carbon and energy prices. [33] and [141] explore the correlation between permit and bond prices. Furthermore, [154] used variance decomposition to examine the directional connections in the Carbon-Energy-Finance system, showing that the carbon market is correlated with both the stock and non-energy commodity markets. In [38], carbon prices are found in Europe to have a weak negative association with both oil and natural gas prices during Phases 2 and 3, as assessed through the dynamic relationship between oil, gas, and carbon prices.

Moreover, in [42], the dependency structure between EUA yields and primary energy price yields (coal, gas, oil, and electricity), modelled through a Vine copula, shows that EUA prices are only correlated with energy prices, and the link with oil and gas prices is negative.

Furthermore, the authors use the Granger causality approach not only to examine the relationship between the stock market and the EUA spot prices

but also to improve the forecast predictions for future EUA values [111], highlighting the causality between common factors.

[88] found that CO₂ emission allowance prices exhibit significant persistence, with integration orders close to or slightly below 1, indicating that price shocks tend to have enduring effects. It is also observed that the causal relationship between the stock markets and the EUA spot prices provides valuable information for decision-makers.

After accounting for structural breaks in the data, the persistence decreases, suggesting that recent shocks have become less enduring over time. However, it is observed that the causal relationship between stock markets and EUA spot prices provides valuable insight for decision-makers.

This suggests that in earlier periods, stronger and more immediate policy interventions were required to restore CO₂ levels back to normal following economic or market shocks. In contrast, in more recent periods, the effects of such shocks have become less persistent, meaning that the system has become more resilient, and CO₂ levels return to equilibrium more quickly. The reduced persistence of these shocks can be attributed to several factors, including the financial crisis, which may have led to structural changes in markets, as well as the implementation of more active and effective climate and energy policies that help mitigate the long-term impacts of such shocks.

In their study, [108] and [166] examine the asymmetric volatility spillover effect between the EUA carbon market and the prices of WTI oil, Brent oil, and EU natural gas. They focus on how information connections and spillover effects operate between the carbon and energy markets, analysing the interaction between returns and volatility within the carbon-energy system. Crude oil, clean energy, and coal are identified as key factors that influence both return and volatility patterns. Specifically, the electricity market is

highlighted as the primary recipient of information influenced by the carbon market.

Furthermore, in Phase 2, there was a strong coal spillover effect on the carbon market, while in Phase 3, natural gas became increasingly important [90]. Furthermore, overflow effects were found to be significantly more pronounced in the volatility system than in the returns system [173].

From a forecasting perspective, [167] use a Bayesian network to identify the most informative variables to predict permit prices. They highlight that natural gas and crude oil have a direct influence on carbon prices, while the S&P 500 and the Global Clean Energy Index affect the carbon market indirectly. Moreover, [175] discovered that macroeconomic volatility and energy disruptions, particularly during the COVID-19 pandemic and the energy crisis, significantly impacted market dynamics. These disruptions made EU ETS prices more sensitive to energy-related factors, highlighting the growing influence of energy markets on carbon pricing during times of crisis. This shift underscores the importance of adopting a broader approach to market analysis to better predict price fluctuations driven by changes in energy supply and demand. Using combination MIDAS regression models, the study forecasts weekly carbon prices based on daily data from the prices of coal, crude oil, gas, and EUROSTOXX 50. Among these, coal was identified as the most influential long-term predictor.

[2] studied the predictive power of the crude oil, natural gas and coal prices to forecast the European carbon price. The authors found that changes in carbon prices are strongly correlated with natural gas prices, while they are only weakly correlated with changes in coal prices on a weekly basis. The combined models, which outperform individual predictors, suggest a general decline in carbon prices over the next month. Additionally, incorporating

weather and air quality data could enhance prediction accuracy in future research.

The issue of mixed-frequency data arises frequently in econometrics and time series analysis, referring to situations where data is collected at different time intervals. For example, in economic and financial research, data might be collected on varying time scales such as daily, weekly, monthly, or quarterly. This presents challenges because most traditional time series models assume that data are recorded at a uniform frequency, making it difficult to integrate data with different time intervals effectively. As a result, techniques must be adapted to handle the complexities of combining such disparate data sources without compromising the accuracy or reliability of the model [143].

MIDAS models [87] are widely recognised in the scientific literature as the most relevant parametric models to combine high-frequency and low-frequency data. These models are used primarily to analyse and forecast macroeconomic indicators and to study the effects of monetary policy [86]. MIDAS models establish a relationship between variables and estimate parameters, typically assuming a specific functional form, such as a weighted sum or a regression with lagged terms, to address mixed-frequency data. The model parameters are estimated using statistical methods like maximum likelihood estimation or Bayesian techniques.

To the best of our knowledge, Wavelet decomposition is the only effective methodology for identifying inter-temporal information. However, its application in the study of parametric financial or energy markets is limited. Wavelet analysis is a mathematical technique that allows for the examination of signals and time series in both the time and frequency domains. The key parameters in Wavelet analysis are scale and translation, which correspond to the width and position of the Wavelet, respectively. The resulting Wavelet

coefficients reveal the contribution of different frequency components at various scales, helping to identify significant features of the signal [152, 163].

In their work, [153] investigate the relationships between the market performance of the renewable energy sector and the fossil fuel energy sector in Europe. The study applies a multiresolution analysis using Wavelet analysis tools, which decompose time series into their time-scale components linked to specific frequency ranges. This approach proves valuable for examining the comovements of fossil fuel and renewable energy index prices across different time horizons.

2.3 Volatility Patterns and Causal Links in the EU ETS

Understanding the causal drivers of EU ETS price fluctuations is crucial for risk management and decision-making. To this end, we introduce the Differentiable Information Imbalance (DII), a non-parametric tool from the physics literature, to identify variables causally linked to EUA financial returns. Comparing DII with multivariate Granger causality, we reveal both shared and distinct causal influences, highlighting the limitations of linear methods.

We began this review by examining papers that apply DII and Granger causality, which guided our decision to extend this methodology to the EU ETS market. From there, we collected foundational studies on causal inference in financial markets, including works by [168], the classical Granger causality framework [91], and [5] research on causal-information-theoretic methods. This approach allowed us to trace the development of key methodologies and adapt them to the analysis of carbon prices.

Examining the volatility of the EU ETS market and its links to other financial assets offers valuable insights for informed decision-making and effective risk management. [43] explore the volatility of the 2008 European Union Allowance (EUA) futures contract, revealing that while daily volatility appears to follow a normal distribution, deviations occur when returns are standardised using daily volatility measures. Furthermore, [41] find that options do not have a systematic impact on market stability and may even contribute to stabilising it. [36] analyse EUA volatility using three measures: the daily EGARCH model, implied volatility of options, and realised volatility from intraday data, identifying significant shifts in volatility, especially in the EGARCH and implied volatility models, linked to compliance events and post-Kyoto uncertainties.

A related study by [162] reveals that applying the EGARCH(1,1) model to high-frequency data uncovers significant long memory, power effects, and asymmetry, all of which are important for managing risks in futures contracts. Similarly, compliance events and post-Kyoto uncertainties contribute to the volatility instability in carbon prices [63].

Despite the EU ETS market being a closed system, shocks can still indirectly affect the market by influencing the behaviour and financial stability of participants, such as companies and investors. [131] demonstrate that the energy crisis and geopolitical shocks, such as the Russia-Ukraine conflict, significantly increased extreme risks in the EU ETS market. These events heightened volatility by impacting industrial activity, emissions, and allowance demand, with notable spillover effects from stock and commodity markets.

Moreover, industrial firms tend to increase trading during periods of low volatility, while financial intermediaries exhibit more adaptable trading be-

haviours [82]. Understanding these dynamics is essential for evaluating market efficiency, as noted in [176] and [13]. The inclusion of energy-related indices highlights the growing attention to the effects of the energy transition and sustainability on financial markets [139].

[35] examine how economic shocks lead to increased volatility in carbon futures markets, driven by uncertainty about economic activity and emissions production. [116] show that policy uncertainty leads to mispricing and increased volatility in carbon markets, with the ongoing Energy Crisis (2021–2024) exacerbating volatility due to rising energy prices and geopolitical tensions. These findings align with [29], who argue that energy prices and carbon markets are closely interconnected, with higher energy prices typically leading to increased volatility in carbon prices due to rising emissions costs. [10] highlight that EUA prices are primarily impacted by volatility in sectors such as finance, energy production, energy-intensive industries, and consumption.

In addition, [9] investigate the correlation between carbon spot and futures prices within European markets, identifying significant interactions and volatility spillovers. [74] provide empirical evidence of substantial volatility co-movement and spillover from the energy markets to the EUA market. In terms of the connection between the carbon and energy markets, [109] analyse returns and volatility spillovers, illustrating dynamic patterns with network diagrams and rolling windows, where electricity markets are found to be the primary recipients of information. The study uncovers a greater interconnectedness in volatility compared to returns, showing significant time-varying characteristics.

Concerning the non-linear bidirectional relationship within the EUA volatility structure, [123] utilise non-linear Granger causality methods to analyse

both mean and volatility spillovers. They find that volatility spillovers increase over time and are especially sensitive to financial crises and extreme events, underscoring their importance for market participants.

To identify key factors that affect the volatility of EUA in evolving dynamics, [122] employ the LASSO method in conjunction with a Markov switching vector autoregression model. This study reveals different impacts of energy and economic factors on the carbon market during the COVID-19 pandemic, with the energy markets having a negative influence and the stock markets a positive one. Notably, during low volatility phases, oil markets have a significant impact on carbon prices, while during high volatility periods, coal markets take precedence.

Furthermore, [127] examine the correlation between EUA prices and factors such as energy prices and economic risks, revealing a non-linear relationship that evolves over time. The high volatility regime is particularly aligned with periods of economic decline or institutional changes that affect confidence in emission targets [9]. [110] explore the influence of annual verified emission announcements in the EU ETS on carbon market returns and volatility, highlighting their impact on traders' expectations and strategies, which in turn affect carbon prices and volatility. However, the study also underscores potential market shocks due to early data leaks and stresses the importance of robust risk prevention measures.

[71] investigate the effect of economic policy uncertainty on carbon spot return volatility in the EU ETS using a GARCH-MIDAS-RV model. Similarly, [156] analyse the volatility of carbon prices in both the EU ETS and the Clean Development Mechanisms (CDM) markets, identifying influences from market mechanisms, external factors and price trends.

Focussing on heavily traded futures contracts in Phase 2 of the EU ETS,

[29] identify a new pricing regime driven by economic growth, energy prices, and weather conditions. Unlike earlier studies, they emphasise the higher stability of futures contracts against structural breaks, offering insight into the long-term relationship between determinants and EU carbon allowances. In [79], the dynamics of EUA Phase 2 and Phase 3 futures prices and volatility are analysed. Traditional modelling methods face challenges in handling excessive skewness and kurtosis in return series, mainly due to outliers. To address this, the study introduces a mixture of Gaussian distributions to account for endogenously determined jumps in the return process.

Nevertheless, a delicate balance must be maintained between providing information and ensuring market stability, which underscores the need for careful information release by regulatory authorities [104, 83]. [67] examine emission permits as financial assets and highlight 11 brief price surges, mainly in 2009 and 2011, where permit prices spiked rapidly but did not sustain the increase. Speculative and hedging activities in the European carbon futures market are associated with significant volatility, as noted by [126].

This research finds high speculation during Phase II listing, particularly in the first quarter, likely related to EU ETS deadlines. These dynamics contribute to increased volatility and have implications for EU ETS reforms and contract design, suggesting the need for adjustments to manage such volatility. Similarly, [73] focus on forecasting EUA returns and volatility in all phases of the EU Emission Trading Scheme.

Finally, [135] observe that despite major crises, such as the global financial and European sovereign debt crises, the returns of the carbon market follow a random walk pattern, indicating that past returns do not predict future returns.

Chapter 3

Methods and Models

3.1 The Information Imbalance

3.1.1 Introduction to Information Content and Feature Selection

Quantifying the information content of datasets plays an important role in tasks such as feature selection and dimensionality reduction. These methods aim to identify the most relevant features or to represent data in reduced dimensions while preserving original structures.

In this context, accurate assessment of the relative informativeness of feature subsets or distance metrics enables informed decisions regarding data representation. This section delves into the detailed theoretical foundation of the Information Imbalance (II), a measure designed to compare distance metrics using mathematical and statistical tools and concepts such as copulas.

3.1.2 Theory of Copulas in Defining Information Imbalance

Consider a data set $\{x_i\}_{i=1}^N$ sampled from a density $p(x)$ and two spaces equipped with distance measures $D_X(x, x')$ and $D_Z(x, x')$. The joint probability of sampling a point x' at distances d_X and d_Z from a reference point x is given by

$$p(d_X, d_Z | x) = \int dx' p(x') \delta(d_X - D_X(x, x')) \delta(d_Z - D_Z(x, x')). \quad (3.1)$$

Although $p(d_X, d_Z | x)$ represents the relationship between the two spaces, practical difficulties arise because X and Z may use different units or scales. These challenges can be addressed using Sklar's Theorem, [151], which decomposes the joint distribution into a product of marginal distributions and a copula function.

$$p(d_X, d_Z | x) = p(c_X(d_X), c_Z(d_Z) | x) p_X(d_X | x) p_Z(d_Z | x), \quad (3.2)$$

where c_X and c_Z are the copula variables defined as

$$c_X(d_X) = \int_0^{d_X} dw p_X(w | x), \quad c_Z(d_Z) = \int_0^{d_Z} dw p_Z(w | x). \quad (3.3)$$

These variables, which vary uniformly in $[0, 1]$, transform the distance measures into dimensionless quantities. The copula distribution $p(c_X, c_Z | x)$, therefore, represents the correlation structure between X and Z , independent of the marginal scales [133].

Distance ranks r_X and r_Z are statistically equivalent to copula variables and provide a means of analysing the information structure in spaces.

3.1.3 Connection to Information Theory

Mutual information $I(X, Z)$ between two spaces X and Z measures the reduction in uncertainty about one space given the knowledge of the other [64].

$$I(X, Z) = H(X) - H(X | Z), \quad (3.4)$$

where $H(X)$ is the entropy of X , and $H(X | Z)$ is the conditional entropy [149]. However, mutual information is symmetric and does not account for directional informativeness between spaces.

To address this, mutual information can be expressed in terms of copula variables

$$I(X, Z) = -H(c_X, c_Z) = -H(c_Z | c_X). \quad (3.5)$$

To generalise mutual information and introduce asymmetry, the h -restricted information is defined.

$$I_h(X \rightarrow Z) = - \int_0^h H(c_Z | c_X = \tilde{c}_X) d\tilde{c}_X. \quad (3.6)$$

For $h = 1$, this reduces to the standard mutual information. For $h < 1$, $I_h(X \rightarrow Z)$ becomes asymmetric, capturing directional information content.

3.1.4 The Information Imbalance Metric

The Information Imbalance (II) $\Delta(X \rightarrow Z)$, first introduced by [89], quantifies the asymmetry in information content between two spaces. Unlike traditional measures, the II is more general and model-free, capturing both linear and non-linear dependencies. As a non-parametric measure, it does not rely on any specific functional form or distribution of the data, making it especially versatile for uncovering complex and unknown relationships between variables.

The concept is deeply rooted in copula theory, which provides a framework for modelling and measuring dependence between random variables. By using copulas, the II quantifies the degree of dependence between variables in a manner that is independent of their marginal distributions, thus effectively capturing both linear and non-linear relationships without distributional assumptions.

For small h , the derivative of I_h is related to the entropy of $p(c_Z | c_X = \epsilon)$

$$\left. \frac{dI_h(X \rightarrow Z)}{dh} \right|_{h=0} = -H(c_Z | c_X = \epsilon). \quad (3.7)$$

By approximating $p(c_X | c_Z = \epsilon)$ with an exponential distribution, its entropy can be expressed as:

$$H(c_Z | c_X = \epsilon) \approx 1 + \log \langle c_Z | c_X = \epsilon \rangle, \quad (3.8)$$

which leads to:

$$\Delta(X \rightarrow Z) \approx \langle c_Z | c_X = \epsilon \rangle. \quad (3.9)$$

Alternatively, the entropy of the conditional copula distribution can quan-

tify the dispersion of $p(c_X | c_Z = \epsilon)$. The entropy is related to the conditional mean:

$$H(c_X | c_Z = \epsilon) = 1 + \log \langle c_X | c_Z = \epsilon \rangle. \quad (3.10)$$

Since the mean and entropy are monotonically related, either metric provides consistent results regarding the relationships of equivalence, partial equivalence, inclusion, or independence between spaces.

Given a variable X and two points, or statistical units, i and j , the rank r_{ij}^X of point j with respect to point i is defined by sorting the pairwise distances between i and all other points in ascending order. The rank r_{ij}^X represents the position of the distance between the points i and j in the ordered sequence. For example, if point j is the closest to point i , then $r_{ij}^X = 1$. Similarly, the rank r_{ij}^Z measures the relative position of point j with respect to point i based on the distances computed from variable Z . Generally, $r_{ij}^X \neq r_{ij}^Z$, reflecting the differences in how X and Z are spatially organised in the dataset.

The II from X to Z , denoted $\Delta(X \rightarrow Z)$, is then defined over a dataset with N data points, where the joint values for both X and Z are recorded. The Information Imbalance is formulated as

$$\Delta(X \rightarrow Z) = \frac{2}{N} \mathbb{E}[r_Z | r_X = 1], \quad (3.11)$$

where the expectation is taken only over the nearest neighbour points according to variable X . This measure quantifies the extent to which the proximity of points in X -space can be used to predict the proximity of points in Z -space. In the limit as $N \rightarrow \infty$, the II is constrained to the interval $(0, 1)$, where a value of $\Delta(X \rightarrow Z) \approx 0$ indicates that X provides almost all the

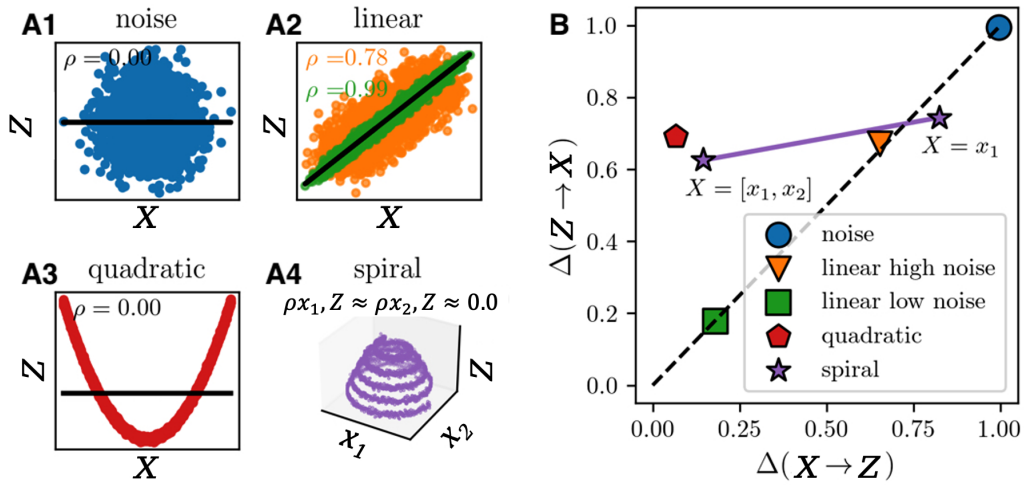


Figure 3.1: **The Information Imbalance.** This figure illustrates the concept of Information Imbalance for various intuitive relationships between independent variables (X) and a dependent (target) variable (Z). Panel **A** presents specific datasets, with each dataset corresponding to markers of the same colour in Panel **B**. The Information Imbalance metric effectively captures different types of relationships: the trivial dependencies in **A1** and **A2**, the non-linear (quadratic) relationship in **A3**, and the multivariate dependency in **A4**, where the linear correlation coefficient (ρ) is not applicable. This highlights the advantage of Information Imbalance in detecting complex dependencies beyond simple correlation measures.

predictive information for Z , and $\Delta(X \rightarrow Z) \approx 1$ suggests that X offers no useful information for predicting Z .

This limiting behaviour can be understood from Equation 3.11. When the variables X and Z are identical, the expected rank evaluates to 1, and the Information Imbalance approaches $2/N$, which is close to zero for large N . On the other hand, for completely unrelated variables, the expected rank is approximately $N/2$, and the Information Imbalance reaches a value of 1, reflecting the lack of predictability.

The Information Imbalance plane is a two-dimensional graph where $\Delta(X \rightarrow Z)$ is plotted against $\Delta(Z \rightarrow X)$. Each point in the plane corresponds to a

specific pair of variables X and Z and represents the directional predictive relationships between them. This is illustrated in figure 3.1, which shows the results of analysing four types of synthetic datasets through the II plane.

In Panel A1, X and Z are purely related by Gaussian noise, resulting in a Pearson correlation $\rho = 0$. Both the II values $\Delta(X \rightarrow Z)$ and $\Delta(Z \rightarrow X)$ are equal to 1, placing the point at the top-right of the II plane. This suggests that there is no predictive relationship between X and Z .

Panel A2 shows a scenario in which a correlation exists between the two variables. As the strength of the linear relationship increases, both the Pearson correlation ρ and the II change. The point moves diagonally from the top-right toward the bottom-left of the plane, reflecting a gradual decrease in II as the variables become more predictively related. In this case, lower noise levels in the linear relationship result in a point closer to the origin of the plane. Notably, the II captures these linear relationships effectively, but its real power lies in identifying more complex dependencies.

In Panel A3, a non-linear (quadratic) relationship exists between X and Z . Despite having a zero Pearson correlation ρ , which fails to capture the non-linearity, the II detects a meaningful relationship. Specifically, it shows that Z can be predicted by X better than the reverse, as reflected in the low value of $\Delta(X \rightarrow Z)$ and the high value of $\Delta(Z \rightarrow X)$.

Panel A4 illustrates a more complex multivariate relationship, where x_1 and x_2 together are much more effective in predicting Z than each variable individually. The linear correlation ρ would fail to capture this multivariate interaction, but the II shows a significant reduction in $\Delta(X \rightarrow Z)$ when both x_1 and x_2 are included in the model, as compared to when only x_1 is used. This reflects how the II can capture the value of multivariate predictors, revealing dependencies that are not detectable using traditional linear

correlation alone.

The II framework, rooted in the statistical theory of copulas, offers a robust and interpretable measure of the directional information content between distance measures. By focussing on local neighbourhoods and leveraging rank distributions, this approach is particularly well-suited for analysing non-linear data manifolds, enabling effective feature selection and dimensionality reduction.

3.1.5 Beyond Linear Correlation

The Information Imbalance (II) metric may initially appear similar to linear correlation, as both quantify relationships between two datasets or feature spaces. However, fundamental differences highlight the distinctive strengths of II in capturing complex relationships.

Linear correlation is a well-established measure of the strength and direction of a linear relationship between two variables. For two random variables X and Z , the Pearson correlation coefficient, [138], $\rho_{X,Z}$ is formally defined as:

$$\rho_{X,Z} = \frac{\text{Cov}(X, Z)}{\sigma_X \sigma_Z},$$

where $\text{Cov}(X, Z)$ is the covariance between X and Z , and σ_X and σ_Z are their standard deviations.

The covariance, given by $\text{Cov}(X, Z) = \mathbb{E}[(X - \mathbb{E}[X])(Z - \mathbb{E}[Z])]$, measures the extent to which the deviations of X and Z from their respective means are proportional. The Pearson correlation coefficient is bounded between -1 and $+1$, where values near ± 1 indicate strong linear relationships, and values near 0 indicate weak or no linear relationship.

Linear correlation focusses on linear dependencies and is symmetric, assigning equal significance to both variables regardless of directionality. This limits its use in non-linear or asymmetric scenarios.

In contrast, Π , accounts for non-linear, asymmetric relationships between distance measures in two spaces. By evaluating conditional rank distributions and their deviations, Π offers a richer framework for analysing information content and dependencies, surpassing the limitations of linear correlation.

Pearson correlation assumes a strictly linear relationship, failing to capture non-linear dependencies between variables. In contrast, Π focusses on preserving the local neighbourhood structure of one space X in another space Z , using the copula distribution $p(c_Z | c_X)$, where c_X and c_Z are cumulative distribution functions of distances in their respective spaces. While Pearson's correlation quantifies symmetric global linear relationships ($\rho_{X,Z} = \rho_{Z,X}$), Π captures asymmetric relationships by analysing the conditional distribution $p(c_Z | c_X = \epsilon)$ for small ϵ , emphasising local structures and information transfer at infinitesimal scales. Π ranges from 0 to 1, indicating equivalence, partial inclusion, or independence between spaces. For example, $\Delta(X \rightarrow Z) = 0$ implies that Z is completely contained in X , and $\Delta(X \rightarrow Z) = 1$ indicates independence.

Unlike correlation, which treats X and Z symmetrically, Π captures directional asymmetries, making it suitable for identifying dominant representations or dependencies in high-dimensional, non-linear, or manifold-based data representations.

3.2 The Differentiable Information Imbalance

3.2.1 Introduction to Causal Discovery and Relationships: Moving Beyond Granger Causality

Understanding causal relationships is essential for uncovering the mechanisms driving complex systems, especially in financial markets where dynamic interactions between variables play a critical role in decision making, monitoring, and forecasting.

Granger causality (GC), defined by [91], captures the idea of a directional influence between two variables, where past values of one variable (the cause) help predict future values of another (the effect). This approach does not imply true causality, which touches on a more philosophical aspect of understanding cause-and-effect relationships, but highlights predictive relationships, which are particularly useful in economics and financial analysis, where establishing true causality is often hindered by system complexity and confounding factors.

Traditionally, causality was studied through controlled experiments, allowing direct manipulation of variables to observe outcomes. However, in real-world scenarios such as financial markets, direct experimentation is impractical or unethical, making observational data crucial for inferring causal relationships [85]. This has led to the development of statistical methods to detect causality in time-ordered observational data [105, 137].

Conventional methods such as GC, mutual information, and other causal inference techniques have been widely used to identify causal relationships from observational data. However, these approaches can face challenges when dealing with complex, high-dimensional, or non-linear systems, which are common in financial markets [93]. Given the growing complexity of these

systems, it becomes important to explore methodologies that can more effectively capture non-linear dependencies and intricate interactions.

To address these challenges, a novel, non-parametric, and non-linear metric is proposed for detecting causal relationships in financial data, offering an alternative to traditional models like the VAR (Vector AutoRegression) model. Our approach introduces the Differentiable Information Imbalance (DII) as a potential method for modelling causal relationships in dynamic, high-dimensional systems, offering a parametric-free framework for causal discovery. Additionally, the Imbalance Gain (IG) metric is proposed as a relative measure of causal strength within the DII framework, which provides a non-linear, non-parametric approach similar to the F-statistic used in Granger causality.

3.2.2 Linear Causal Discovery: VAR and Granger’s F-statistic

Vector AutoRegressive (VAR) models are frequently employed alongside the Granger causality (GC) test in the analysis of financial datasets to assess causal relationships quantitatively. These models are particularly useful for representing the dynamic interconnections between dependent and independent variables over a temporal framework.

To illustrate, let us consider a VAR(p) model, with p denoting the quantity of time lags included, and the structure of the system of equations can be expressed as follows

$$\mathbf{v}_t = \sum_{i=1}^p \mathbf{A}_i \mathbf{v}_{t-i} + \mathbf{u}_t, \quad (3.12)$$

In this context, \mathbf{v}_t denotes a vector of variables at a given time t . The

matrix \mathbf{A}_i serves as the weight matrix corresponding to the i -th lag, and the term \mathbf{u}_t represents a white Gaussian random vector. The fundamental concept of a VAR model is that each variable present in the vector \mathbf{v}_t is expressed as a linear combination of both its own historical values and those of the other variables within the vector [96]. Typically, \mathbf{v}_t is assumed to incorporate not only a target variable, designated as z_t , but also possible predictor variables formulated as $\mathbf{x}_t^T = (x_t^1, \dots, x_t^D)$. Therefore, without loss of generality, \mathbf{v}_t can be expressed as $(z_t, x_t^1, \dots, x_t^D)^T$. To determine the causal weight w^α linked to the predictor x^α at a given lag i , one can calculate it as the absolute magnitude of the $(1, \alpha)$ -entry in the matrix \mathbf{A}_i . It is important to note that according to our notation, the first row of the matrix \mathbf{A}_i includes the weights that dictate the influence of all predictors at lag i on the target variable z .

The Akaike Information Criterion (AIC) is utilised to identify the optimal lag for the VAR model by minimising the balance between model fit and complexity. Hence, the lag corresponding to the lowest AIC score is considered the optimal choice.

Granger's F-statistic For each specified time lag, the GC test is used to determine to what extent the historical values of a particular predictor x_t^α (where α is an element of the set $\{1, \dots, D\}$) contribute to the prediction of the target variable z_t [92].

The procedure involves testing the null hypothesis that there is no causal link between the predictor and the target. This is accomplished by fitting and contrasting two distinct models: the first VAR model incorporates the lagged values of x_t^α as part of the potential predictors \mathbf{x}_t , and is known as the unrestricted model, while the second VAR excludes x_t^α , keeping all other

predictors constant, is known as the restricted model. Should the test result in rejecting the null hypothesis, it implies the existence of a causal connection between the predictor x_t^α and the target variable z_t .

The statistical test is based on the construction of an F statistic that assesses the joint significance of the weights associated with the predictor variable within the complete VAR model. Specifically, for each predictor variable x_t^α , the F-statistic tests whether the weights associated with the lagged values of x_t^α across all lags are jointly different from zero. A high F-statistic value indicates that the past values of the predictor have significant explanatory power for the target variable, suggesting the presence of a causal relationship [92].

The F-statistic is calculated as:

$$F = \frac{(RSS_r - RSS_u)/q}{RSS_u/(N - k)}. \quad (3.13)$$

The equation referenced earlier contains RSS_r , which stands for the restricted residual sum of squares. This value is derived from a model in which specific causality constraints are applied, meaning certain weights are assigned a value of zero.

Conversely, RSS_u indicates the unrestricted residual sum of squares, which is computed using the complete model where all weights are estimated without any imposed constraints. The symbol q signifies the total number of restrictions, equivalent to the number of weights assigned to zero in the restricted model. Typically, in standard implementations, the restricted VAR is created by nullifying all lagged values of one specific potential causal variable, hence q equals p .

To further clarify, let us define N as the total number of observations and k as the number of parameters estimated within the unrestricted model

framework. The expression in the numerator $(RSS_r - RSS_u)/q$ quantifies the reduction observed in the residual sum of squares attributed to applying the restriction; in other words, it reflects the extent to which the unrestricted model provides a better fit relative to the restricted model, with this improvement being adjusted by the number of constraints imposed.

Conversely, the denominator $RSS_u/(N - k)$ signifies the mean squared error for the unrestricted model, representing an estimation of the variance associated with the error terms.

Beyond merely deciding whether the null hypothesis should be accepted or rejected at a specified significance level, the F-statistic, as illustrated in Eq. (3.13), also serves as a valuable tool for comparing the causal impacts of various variables on the same target.

3.2.3 Non-Linear Causal Discovery: the Differentiable Information Imbalance and Imbalance Gain

Differentiable Information Imbalance (DII) generalises the Information Imbalance (II) measure, introduced in [89], to allow for a more flexible evaluation of the predictive relationship between a set of variables \mathbf{x} and a target variable z .

Although II quantifies predictivity by comparing the relative distances between the data points in \mathbf{x} and z , it assumes that all features are equally scaled. The DII extends this approach by introducing a weight vector \mathbf{w} , which scales each feature in \mathbf{x} based on its relevance. This modification enables the optimisation of feature contributions through gradient-based methods, making the DII framework more adaptable to datasets where feature importance varies. The following paragraph details the DII framework, which builds upon the foundational concepts of II.

From the Information Imbalance to a Differentiable version. The methodology employed in this work is a generalisation of the Information Imbalance measure, introduced in [89] and reviewed in the following.

Given a dataset of N points, the Information Imbalance allows to quantify how much a set of variables \mathbf{x} can predict a given target variable z .

In this context, \mathbf{x} is assumed to be multi-dimensional and z one-dimensional, as this setting is most relevant for the application, but it is important to stress that the method can be applied regardless of the dimensionality of the two sets.

The Information Imbalance measure is based on the idea that \mathbf{x} is predictive with respect to z when data points that are close in \mathbf{x} remain close in z .

In practice, Euclidean distance functions $d^{\mathbf{x}}$ and d^z are used, allowing the distances between two points i and j to be expressed as $d_{ij}^{\mathbf{x}} = |\mathbf{x}_i - \mathbf{x}_j|$ and $d_{ij}^z = |z_i - z_j|$, where \mathbf{x}_i and z_i represent the coordinates of point i in spaces \mathbf{x} and z , respectively. For each point i , the distances between i and all other points $j (\neq i)$ are sorted from smallest to largest. The distance rank $r_{ij}^{\mathbf{x}}$ (r_{ij}^z) is defined as the position of $d_{ij}^{\mathbf{x}}$ (d_{ij}^z) in the list of sorted distances.

For example, $r_{ij}^{\mathbf{x}} = 1$ if j is the closest point to i according to the distance $d^{\mathbf{x}}$, and $r_{ij'}^z = 5$ if j' is the fifth nearest neighbour of i according to d^z .

The Information Imbalance from \mathbf{x} to z is written as:

$$\Delta(\mathbf{x} \rightarrow z) = \frac{2}{N} \mathbb{E} [r^z \mid r^{\mathbf{x}} = 1] = \frac{2}{N^2} \sum_{\substack{i,j=1 \\ (i \neq j)}}^N r_{ij}^z \delta_{r_{ij}^{\mathbf{x}}, 1}, \quad (3.14)$$

where δ denotes the Kronecker delta function, which restricts the sum to pairs of points satisfying $r_{ij}^{\mathbf{x}} = 1$.

As Eq. (3.14) shows, $\Delta(\mathbf{x} \rightarrow z)$ is directly proportional to the average

distance rank in space z , conditioned over pairs of points that are nearest neighbours in \mathbf{x} . The prefactor $2/N$ allows an asymptotic normalisation of the Information Imbalance to 1, as a function of N , in the case of minimum predictivity, namely when \mathbf{x} carries no information about z . Indeed, in this case $\mathbb{E}[r^z | r^{\mathbf{x}} = 1] = N/2$, as the conditional distance ranks in space z will be uniformly distributed between 1 and $N - 1$.

In the opposite regime, namely when \mathbf{x} is maximally predictive of z , $\mathbb{E}[r^z | r^{\mathbf{x}} = 1] = 1$, as all nearest neighbour points in \mathbf{x} will also be nearest neighbours in z . In this second limit case, $\Delta(\mathbf{x} \rightarrow z) = 2/N$, which approaches 0 in the limit of large N . Therefore, $\Delta(\mathbf{x} \rightarrow z) \approx 0$ suggests that variables \mathbf{x} are highly informative in predicting variable z , while $\Delta(\mathbf{x} \rightarrow z) \approx 1$ indicates that variables \mathbf{x} provide little to no predictive information about variable z .

The DII framework. The capacity of predictor variables \mathbf{x} to mimic the neighbourhood structure determined by a target distance d^z is based on the choice of distance metric $d^{\mathbf{x}}$. Although the Euclidean distance is reasonable since the data manifold appears approximately flat on small scales, the scaling of variables is relatively arbitrary.

[168] addressed this issue by creating a differentiable adaptation of the Information Imbalance, known as the Differentiable Information Imbalance (DII).

To introduce it, the weighted space $\mathbf{x}(\mathbf{w}) = \mathbf{w} \odot \mathbf{x}$ must be defined first, where \odot represents the element-wise multiplication between a vector of arbitrary weights and the set of predictor variables. In the function $\mathbf{x}(\mathbf{w})$, every variable is adjusted by multiplying it by the elements of a weight vector. In conclusion, the distance is specified as $d_{ij}^{\mathbf{x}}(\mathbf{w}) = |\mathbf{w} \odot (\mathbf{x}_i - \mathbf{x}_j)|$, representing

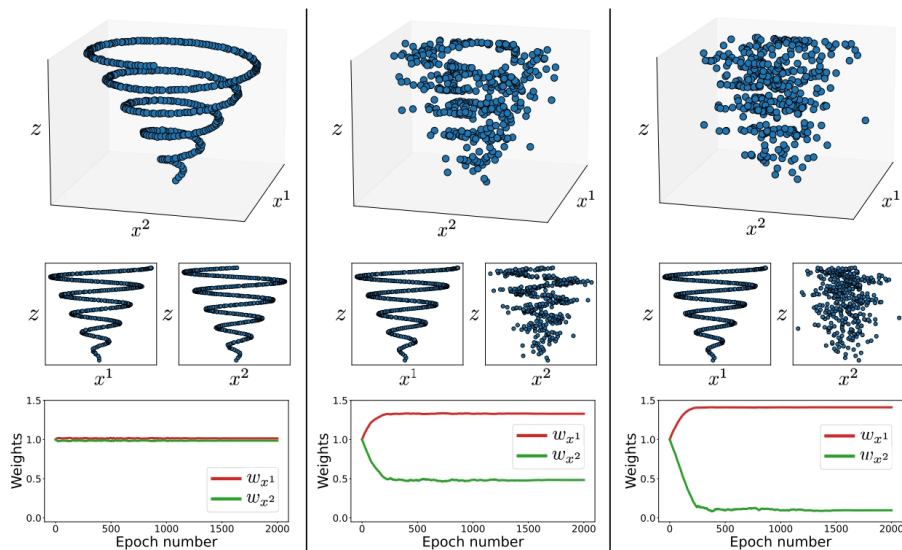


Figure 3.2: **The Differentiable Information Imbalance.** This figure demonstrates the application of the Differentiable Information Imbalance (DII) to three spiral datasets with varying levels of noise. The first row displays the datasets in three-dimensional space, while the second row presents two-dimensional projections to highlight the marginal relationships between the target variable z and the predictors x^1 and x^2 . The third row shows the evolution of the weights w_{x^1} and w_{x^2} assigned by the DII to x^1 and x^2 during the gradient descent process across different epochs. The left panels illustrate that the DII considers both x^1 and x^2 as equally important for predicting z , with weights satisfying $w_{x^1} \approx w_{x^2} \approx 1$. In contrast, the middle and right panels show that as the noise level increases between z and x^2 , the DII assigns a smaller weight to x^2 , with w_{x^2} approaching zero.

the Euclidean distance computed within the space weighted by \mathbf{w} .

The DII from $\mathbf{x}(\mathbf{w})$ to z is defined as

$$\text{DII}(\mathbf{x}(\mathbf{w}) \rightarrow z) = \frac{2}{N^2} \sum_{\substack{i,j=1 \\ (i \neq j)}}^N c_{ij}(\lambda, \mathbf{w}) r_{ij}^z, \quad (3.15)$$

where the weights c_{ij} are defined as:

$$c_{ij}(\lambda, \mathbf{w}) = \frac{e^{-d_{ij}^{\mathbf{x}}(\mathbf{w})^2/\lambda}}{\sum_{m(\neq i)} e^{-d_{im}^{\mathbf{x}}(\mathbf{w})^2/\lambda}}. \quad (3.16)$$

In this context, λ serves as a small, positive parameter that determines the scale of neighbourhoods within the space $d^{\mathbf{x}}(\mathbf{w})$. Specifically, as λ approaches zero, the softmax weights in Equation (3.16) converge to the Kronecker delta in Equation (3.15). This limits the selection to just the nearest neighbour, thereby restoring the usual Information Imbalance scenario. To facilitate optimisation, λ is typically maintained at a magnitude comparable to the average squared distance from the nearest neighbour. The advantage of the formulation in Eq. (3.15) is that derivatives of the DII with respect to weights \mathbf{w} can be explicitly computed. This allows minimising $\text{DII}(\mathbf{x}(\mathbf{w}) \rightarrow z)$ via gradient descent, namely following an update rule of the form:

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta \nabla_{\mathbf{w}} \text{DII}(\mathbf{x}(\mathbf{w}) \rightarrow z), \quad (3.17)$$

until convergence.

In this context, η represents a hyperparameter referred to as the learning rate, while $\nabla_{\mathbf{w}} \text{DII}$ signifies the gradient of DII concerning the weight vector. Consequently, employing gradient descent for the optimization of the DII inherently determines the weighted combination among the potential D variables that best forecasts the target variable z .

An illustrative example is presented in figure 3.2. Further details on the DII optimisation are reported in A.2.

The Imbalance Gain. In [157], the authors applied the standard Information Imbalance method to identify causal relationships between high-dimensional dynamical systems. This approach extends the predictability criterion of Granger causality (GC) into a non-linear framework.

Given two one-dimensional time series x_t and z_t , it can be said that x_0 at time $t = 0$ improves the prediction of z at a future time $t = \tau$ if a distance

measure built with both x_0 and z_0 is more informative about z_τ than a distance where only z_0 appears. This is satisfied when there exists a value of $\tau > 0$ such that the following inequality holds [157]:

$$\min_w \Delta ([wx_0, z_0] \rightarrow z_\tau) < \Delta (z_0 \rightarrow z_\tau) , \quad (3.18)$$

In this context, the two-dimensional space $[wx_0, z_0]$ is defined by the scaled variable wx_0 and the target variable z_0 . The weight w serves to enhance the expressiveness of the space that includes both x_0 and z_0 . By minimising over w , the aim is to identify the space that offers the most informative representation regarding z_τ . Distances are calculated between separate realisations of the dynamics. These realisations can be derived from a single stationary time series by sampling time frames, which are subsequently used as independent initial conditions.

In systems where long-memory effects are observed, analysing time lags $\tau > 1$ can improve the identification of dynamical interactions. This is because the influence of a causal relationship $x \rightarrow z$ often becomes more evident after a certain time interval after the interaction.

The aforementioned inequality can be reformulated in terms of the relative difference known as Imbalance Gain (IG):

$$\text{IG}(x \rightarrow z) = 1 - \frac{\Delta (z_0 \rightarrow z_\tau)}{\min_w \Delta ([wx_0, z_0] \rightarrow z_\tau)} , \quad (3.19)$$

which can be expressed as a percentage variation.

Eq. (3.18) is essentially equivalent to the condition $\text{IG}(x \rightarrow z) > 0$. Moreover, both of these equations can be extended to a multivariate context similarly to multivariate GC as shown in [157]. Nevertheless, this extension requires additional optimisation parameters, which results in a computation-

ally intensive minimisation on the right-hand side of Eq. (3.18).

Following the same notation presented in Sec. 3.2.2 for GC, the generalisation of Eq.(3.18) in a multivariate setting reads:

$$\min_{\mathbf{w}} \text{DII}(\mathbf{w} \odot \mathbf{v}_0 \rightarrow z_\tau) < \min_{\mathbf{w}'} \text{DII}(\mathbf{w}' \odot (\mathbf{v}_0 \setminus x_0^\alpha) \rightarrow z_\tau), \quad (3.20)$$

where $\mathbf{v}_0 = (z_0, x_0^1, \dots, x_0^D)^T$, \mathbf{w} and \mathbf{w}' are vectors of $D + 1$ and D optimisation parameters, respectively, and $\mathbf{v}_0 \setminus x_0^\alpha$ denotes the vector \mathbf{v}_0 without variable x_0^α .

Equivalently, the multivariate version of the IG can be written as:

$$\text{IG}(x^\alpha \rightarrow z) = 1 - \frac{\min_{\mathbf{w}'} \text{DII}(\mathbf{w}' \odot (\mathbf{v}_0 \setminus x_0^\alpha) \rightarrow z_\tau)}{\min_{\mathbf{w}} \text{DII}(\mathbf{w} \odot \mathbf{v}_0 \rightarrow z_\tau)}, \quad (3.21)$$

such that a causal effect of x^α on z is detected when $\text{IG}(x^\alpha \rightarrow z) > 0$ for some $\tau > 0$.

Importantly, while the optimal weights minimising the DII depend on the variances of the (unscaled) input variables, the minimum value of the DII, and consequently the IG measure, are independent of such original scales. Thus, deciding to standardise the input variables or not leads to distinct optimal weights. However, it does not influence the final estimate of the Imbalance Gain, as long as equivalent global minima are obtained. Consequently, in this context, the Imbalance Gain serves as a more dependable metric compared to the optimal weights for assessing the true effect of a predictor variable on the target variable.

Causal Sufficiency. Our analysis presupposes the condition known as causal sufficiency [144], implying that no variables, financial or otherwise,

acting as common influencers of European Union Allowance (EUA) and other predictors in the dataset are omitted. This assumption is important in both the DII method and the conventional Granger causality analysis. Although this condition may not hold perfectly in practice, it can be considered reasonably valid in the context of our study, given the comprehensive and diverse set of financial variables included in the analysis. These variables help mitigate the potential impact of omitted influences, making the assumption more plausible within the scope of our research.

3.2.4 Two Case Studies on Synthetic Data

In this section a comparison of the linear approach described in Sec. 3.2.2 with the novel non-linear methodology described in Sec. 3.2.3 on two synthetic datasets with a known data-generating process is proposed.

Both data sets were created by generating time series for x^1 , x^2 , and z , initialising the state of the system with three Gaussian random variables of unit variance. The first 5000 steps were discarded to avoid equilibration artifacts. Subsequently, the $N = 2800$ time steps were selected after equilibration, aligning the number of observations with those used in the empirical analysis of returns with the DII (see A.2). VAR models of order $p = 1$ and DII models with time lag $\tau = 1$ are considered.

False negatives from linear causal discovery. The first synthetic dataset is intended to illustrate how the linear causal discovery method can miss important causal predictors that are non-linearly related to a target variable. The dataset is constructed by simulating the following stochastic process, with one target variable z and two causal predictors x^1 and x^2 :

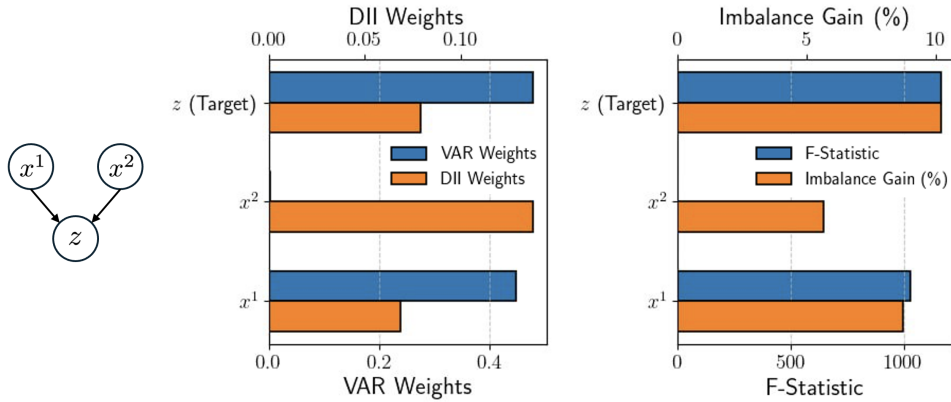


Figure 3.3: **False Negatives from Linear Causal Discovery.** This figure compares the performance of the VAR/multivariate Granger Causality (GC) model and the Differentiable Information Imbalance (DII) method applied to a stochastic process defined by Equations (3.22), with x^1 and x^2 tested as potential causal variables for the target variable z . The left panel shows the true causal relationships among the three variables: x^1 and x^2 evolve independently, while both x^1 and x^2 influence z . The central bar plot displays the multiplicative weights assigned to the variables by the unrestricted VAR(1) model (blue bars) and the optimized DII ($[\mathbf{x}_0, z_0] \rightarrow z_1$) (orange bars). The right panel shows the F-statistics for each variable, calculated within the multivariate GC framework, alongside the Imbalance Gain computed using the DII approach.

$$x_t^1 = u_t^1, \quad (3.22a)$$

$$x_t^2 = u_t^2, \quad (3.22b)$$

$$z_t = 0.5 z_{t-1} + x_{t-1}^1 + (x_{t-1}^2)^2 + u_t^z. \quad (3.22c)$$

Here, u_t^1 , u_t^2 and u_t^z are independent Gaussian white noise terms, sampled from $\mathcal{N}(0, 1)$.

According to Eq. (3.22c), variable z is driven by both x^1 and x^2 , since the state of z at time t is determined by both variables in the previous step.

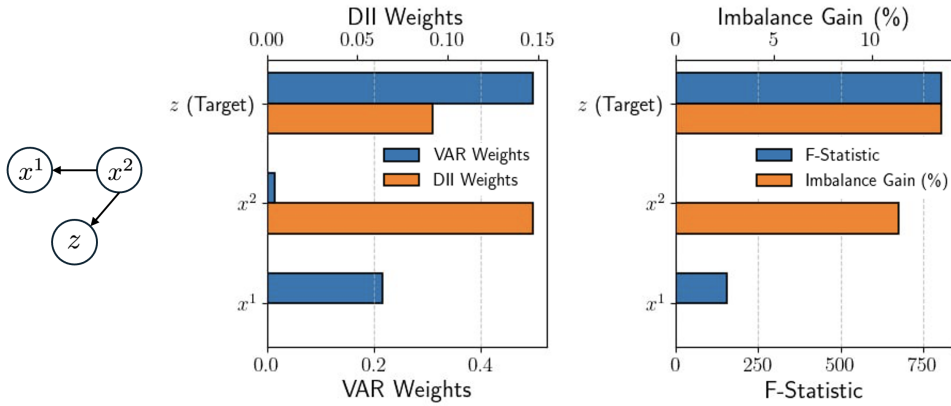


Figure 3.4: **False Positives from Linear Causal Discovery.** This figure compares the performance of the VAR/multivariate Granger Causality (GC) model and the Differentiable Information Imbalance (DII) method applied to the stochastic process described by Equations. (3.23), with x^1 and x^2 tested as potential causal variables for the target variable z . The left panel illustrates the true causal relationships: x^2 acts as a common driver for both x^1 and z , highlighting the structure of the underlying process.

Importantly, the causal relationship $x^1 \rightarrow z$ is linear, while the link $x^2 \rightarrow z$ involves a non-linear (quadratic) coupling.

The results of both methods applied to this simple stochastic process are shown in figure 3.3. As anticipated, the linear causal relationship $x^1 \rightarrow z$ is identified by both approaches, as evident from the variable weights (left panel) and the F-statistic and Imbalance Gain measures (right panels). However, the non-linear relationship $x^2 \rightarrow z$ is not detected by either the VAR weights or the F-statistic, while the DII method successfully identifies it. This example illustrates that standard multivariate GC methods are susceptible to false negatives when dealing with non-linear relationships.

False positives from linear causal discovery. The second synthetic dataset is intended to illustrate how the linearity assumption of VAR models

and multivariate GC can lead to incorrect detection of false causal relationships.

The dataset is constructed by simulating the stochastic process:

$$x_t^1 = 0.1 x_{t-1}^1 + (x_{t-1}^2)^2 + u_t^1, \quad (3.23a)$$

$$x_t^2 = 0.7 x_{t-1}^2 + u_t^2, \quad (3.23b)$$

$$z_t = 0.5 z_{t-1} + (x_{t-1}^2)^2 + u_t^z. \quad (3.23c)$$

In this second case study, u_t^1 , u_t^2 , and u_t^z are independent Gaussian white noise terms with standard deviations of 0.2, 0.5, and 1.0, respectively. In the system described by Equations. (3.23), x^2 acts as a common driver for both x^1 and z , influencing both variables through a non-linear quadratic relationship.

Figure 3.4 shows the results of applying multivariate GC (blue bars) and the DII method (orange bars) to the system, with z selected as the target variable. As in the previous example, it is observed that the linear methodology fails to detect the non-linear coupling $x^2 \rightarrow z$, whereas both the optimal weight from the DII optimisation and the Imbalance Gain associated with x^2 are significantly non-zero. Additionally, the linear method detects a spurious causal link $x^1 \rightarrow z$, as seen from the magnitudes of the VAR weight for x^1 and its corresponding F-statistic from the GC test. In contrast, both the DII weights and the Imbalance Gain for the $x^1 \rightarrow z$ direction show that the impact of x^1 on the dynamics of z is negligible.

This second example demonstrates that linear methodologies may suffer from false positives in the presence of common drivers and non-linear

relationships, while the DII method appears to be robust to this issue.

The two examples illustrated above motivate the introduction of non-linear generalisations, such as the DII, for analysing real-world time series in which the assumption of linear relationship may be, to different extents, violated.

3.3 Gaussian Processes

A Gaussian Process (GP) is a powerful and versatile statistical tool used in various fields, including machine learning [169], statistics [150], and Bayesian optimisation [170]. GPs have gained significant popularity due to their flexibility, ability to model complex, non-linear relationships, and effectiveness in quantifying uncertainty. These properties make GPs an ideal choice for tasks such as function approximation, regression, and time-series forecasting, where understanding the underlying structure and uncertainty of the data is crucial.

A Gaussian Process can be viewed as an infinite-dimensional generalisation of a multivariate Gaussian distribution. In simpler terms, a GP defines a distribution over functions, where each function is drawn from a multivariate Gaussian distribution at any finite subset of points. More precisely, a GP is a collection of random variables where any finite subset of them follows a multivariate Gaussian distribution. This flexibility makes GPs particularly well-suited for modelling continuous functions in a non-parametric manner.

To define a Gaussian Process, a mean function must first be selected $\mu(\mathbf{x})$, which provides the expected value of the modelled function

$$\mu(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})]. \tag{3.24}$$

In practice, the mean function is often assumed to be zero, especially when the data have been standardised. This assumption simplifies the GP model without loss of generality. Throughout this work, this common practice is adopted for simplicity and consistency. The second and crucial component of a GP is the covariance function (or kernel), which characterises the relationships between different points in the function’s domain. The covariance function quantifies the similarity or correlation between function values at different input points, and it plays a central role in shaping the properties of the modeled function.

3.3.1 Kernel Function

The choice of the kernel function is crucial, as it determines the properties of the GP and how it models the dependencies between data points. A widely used kernel is the Radial Basis Function (RBF), which is commonly chosen for its smoothness properties. Other commonly used kernels include additive, multiplicative, and Matern kernels. For this work, the focus is on the Matern kernel, defined as:

$$k_{\text{Matern}}(\mathbf{x}, \mathbf{x}') = \frac{1}{\Gamma(\nu)2^{\nu-1}} \left(\sqrt{2\nu} \frac{\|\mathbf{x} - \mathbf{x}'\|}{l} \right)^\nu K_\nu \left(\sqrt{2\nu} \frac{\|\mathbf{x} - \mathbf{x}'\|}{l} \right), \quad (3.25)$$

where ν controls the smoothness of the function and l is the length scale parameter, which governs how rapidly the function varies with respect to changes in the input. The function $K_\nu(\cdot)$ is the modified Bessel function of the second kind, which provides a mathematically tractable solution to the covariance. The Matern kernel is particularly useful because it provides a flexible trade-off between smoothness and the scale of variations in the

function, making it a good fit for many real-world datasets.

Once the kernel is selected, a GP can be used as a prior distribution in Gaussian Process regression, a form of Bayesian regression. Given a dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$, where \mathbf{x}_i are the input features and y_i are the observed values, the GP model infers a posterior distribution over functions that fit the observed data. Importantly, the posterior distribution in GP regression is also a GP, and the posterior mean provides the best estimate of the underlying function at any point in the input space.

3.3.2 Prediction and Uncertainty Estimation

The posterior mean of a GP provides a predictive mean function that is used for forecasting. Given a test point \mathbf{x}^* , the predictive mean $\mu(\mathbf{x}^*)$ is calculated as:

$$\mu(\mathbf{x}^*) = \mathbf{k}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y}, \quad (3.26)$$

where \mathbf{K} is the kernel matrix with entries $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$, \mathbf{k}_* is the vector of kernel values between the test point \mathbf{x}^* and the training data, and σ_n^2 represents the noise variance in the data. The matrix \mathbf{I} is the identity matrix, and \mathbf{y} is the vector of observed output values.

In addition to the predictive mean, GPs also provide a measure of uncertainty through the predictive covariance. The variance at a test point \mathbf{x}^* gives an indication of how confident the model is in its predictions. This feature is particularly useful in applications like forecasting and decision-making, where not only the prediction but also the uncertainty of that prediction is essential.

3.3.3 Applications in Forecasting, Imputation, and Aggregation

Gaussian Processes (GPs) are widely used in time-series forecasting, where the goal is to predict future values based on historical data.

In addition to forecasting, GPs are also used in imputation tasks, where missing or incomplete data need to be estimated. Since GPs model dependencies between data points, they can predict missing values based on observed relationships in the data. This property is particularly useful in fields such as epidemiology, economics, and finance, where incomplete data are common.

Another important application of GPs is in aggregation, where multiple sources of information or data points need to be combined. GPs can effectively aggregate information by treating each source as a different realisation of a latent function and using the covariance function to model the relationships between these sources. This allows for the construction of more robust and reliable models when data from multiple sources or domains are available.

Chapter 4

Empirical Results I:

Bibliometric insights into the EU ETS market

4.1 Data Collection and Selection Criteria

The data set used in this study is derived from the Scopus database, chosen for its comprehensive coverage of peer-reviewed scholarly articles and its advanced citation tracking capabilities.

The documents collected span from 2004 to 2024, facilitating an in-depth examination over a period of nearly 20 years.

Predefined search terms related to the EU ETS are selected, including *European Union Emissions Trading System*, *EU ETS*, *emissions trading*, *carbon trading*, and *carbon pricing*. (AND, OR) Boolean operators helped refine our search results to ensure a broad yet targeted collection of relevant literature.

The initial search in Scopus returned 475 documents which were exported in BibTeX format. A custom Python script was developed to refine the

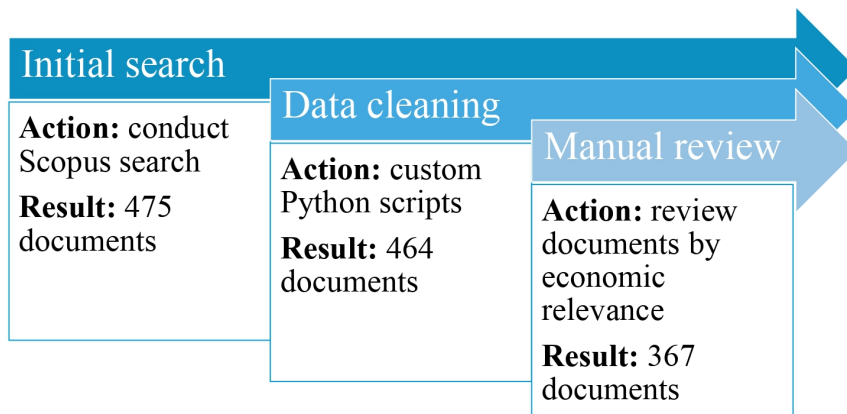


Figure 4.1: **Data Processing Workflow.** This diagram illustrates the sequential phases involved in selecting and processing research documents, from initial data collection to final inclusion, highlighting key decision points and criteria used throughout the process.

dataset by automatically identifying and removing duplicates, errors, and encoding issues. As a result, the dataset was reduced to 464 documents. Furthermore, a manual review ensured that only papers specifically addressing the EU ETS market from an economic or financial standpoint were included, resulting in a final selection of 367 papers, as shown in figure 4.1.

The relevance and quality of the included literature is guaranteed by selecting only articles from peer-reviewed journals classified as articles. To maintain a clear focus on the financial and economic dimensions of the EU ETS market, studies from specific fields, such as Chemistry, Materials Science, Medicine, Physics and Astronomy, Multidisciplinary areas, Arts and

Humanities, Biochemistry, Genetics and Molecular Biology, Pharmacy, Psychology, Neurosciences, Agriculture and Biological Sciences, Civil and Structural Engineering, Engineering, Sociology and Political Science, Earth and Planetary Sciences, Decision Sciences, Mathematics, and Computer Science, were excluded.

Furthermore, only articles written in English were included to ensure accessibility and clarity.

4.2 Bibliometric Tools

Bibliometric analysis was performed using the Bibliometrix R package [8]. The analysis started with descriptive statistics to examine the distribution of publications over time, revealing trends in research activity and the growth rate of publications in the EU ETS field.

Venue analysis was performed to identify the key journals that published economic and financial research on the EU ETS, focussing on the leading journals according to the volume of publications and the impact of citations. In addition, citation analysis was performed to explore the most frequently cited articles, the most influential authors, and the leading institutions within the field. This analysis provided a deeper understanding of the impact and reach of specific studies, highlighting the contributions of key researchers and institutions, and shedding light on the broader scope and significance of their work within the context of economic and financial EU ETS research.

Co-authorship and collaboration networks were mapped to investigate patterns of cooperation among researchers and institutions. This network analysis revealed prominent research networks, highlighting key collaborative connections, and identifying influential groups of scholars working together

in the field of economic and financial research on EU ETS.

In addition, a keyword analysis was performed to identify the most prevalent research themes and topics within the EU ETS literature. By analysing the co-occurrence of keywords, related terms were grouped and clustered, uncovering key thematic trends and highlighting emerging areas of research in the field.

Finally, the conceptual structure of the EU ETS research field was mapped through co-word analysis, which visualised the interrelationships between various research themes. This analysis helped identify and illustrate the main research clusters, providing a clearer understanding of how different topics are connected and the key areas of focus within the field.

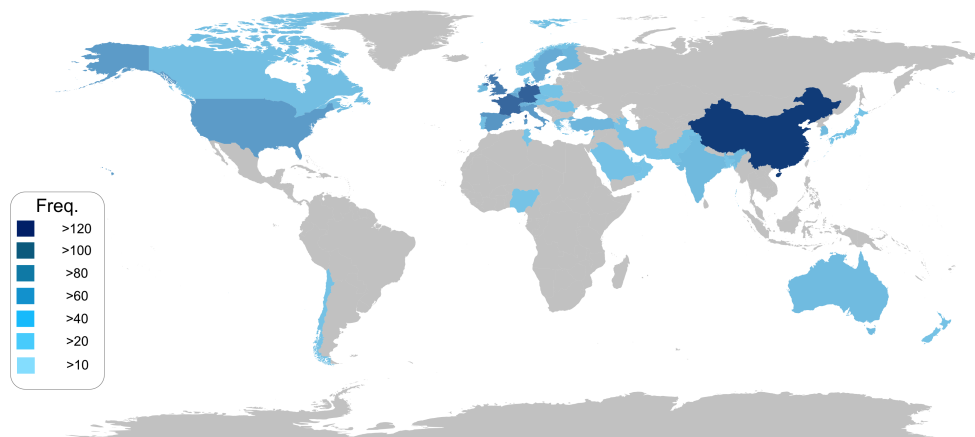


Figure 4.2: **Country Scientific Production.** This plot shows the volume of scientific publications related to the EU ETS across various countries. It provides a comparative view of publication output by country. For more detailed figures and data on each country's contributions, refer to Table 4.1.

4.3 Descriptive Analysis

Figure 4.3 presents an overview of publication trends, with a focus on the annual publication count. This analysis highlights notable fluctuations in research activity, illustrating significant increases or decreases over time.

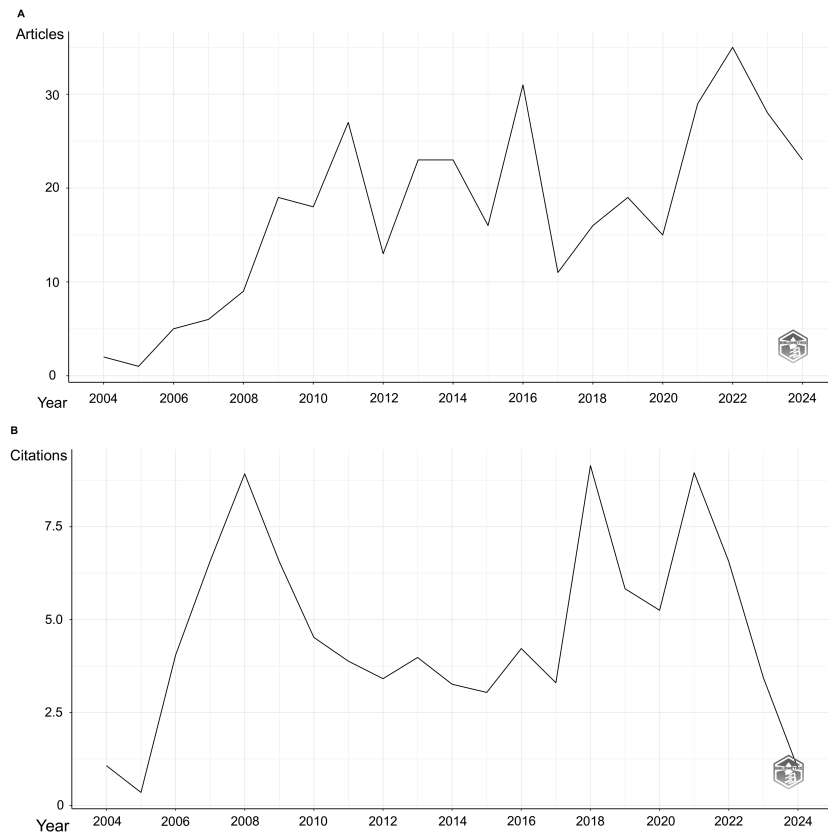


Figure 4.3: **Annual Scientific Production (A) and Average Citations per Year (B).** This figure shows the trends in scientific publications and average citations per year from 2004 to 2024. Panel (A) illustrates the annual increase in the number of publications, while Panel (B) presents the average number of citations per publication each year.

The exponential increase in publications on the EU ETS market began in 2005, coinciding with the onset of phase 1. Since then, the number of articles has fluctuated, yet it has consistently remained above 10 articles

annually. A surge in scientific production was observed in 2021, coinciding with the beginning of phase 4 and the subsequent announcement of EU ETS 2, leading to a peak of 35 articles in 2022.

The geographic distribution analysis, shown in figure 4.2, highlights the countries with the highest contributions to the literature, determined by the affiliations of the authors.

Table 4.1 presents data on the frequency of scientific publications by region, focussing on their contributions to the global scientific community.

China emerges as the most prolific country, with a total of 168 publications, demonstrating its dominant position in scientific output compared to other nations. Germany ranks second with 113 publications, highlighting its significant contribution to global scientific research. France follows closely in third place with 105 publications, while the UK holds fourth place with 81 publications. Spain takes the fifth position, with 49 publications, and Italy ranks sixth with 44 publications, reflecting important contributions, albeit slightly smaller than those of the leading countries.

Furthermore, the United States follows with 43 publications, followed by Belgium, Sweden, and Switzerland with 30, 26, and 20 publications, respectively. Finally, Finland and the Netherlands contribute 14 publications each.

Besides, table 4.1 ranks countries according to their citation metrics, taking into account both the total number of citations and the average citations per publication. The analysis reveals that Germany leads with the highest total citation count and an impressive average of 55.00 citations per paper. France and China follow, with France having a slightly higher total citation count, while China has a marginally lower average citation rate. Interestingly, the Netherlands stands out for having the highest average citation rate of 96.50, despite a lower total citation count of 386.

Country	Frequency	Total Citations	Average Article Citations
China	168	2003	32.80
Germany	113	2476	55.00
France	105	2150	53.80
UK	81	1290	44.50
Spain	49	932	44.40
Italy	44	445	22.20
USA	43	399	33.20
Belgium	30	290	29.00
Sweden	26	387	32.20
Switzerland	20	553	55.30
Finland	14	336	67.20
Netherlands	14	386	96.50

Table 4.1: **Publication Frequency and Citations by Country.** This table summarizes the frequency of academic publications and the total number of citations for each country. It provides an overview of the publication output and citation impact across different regions.

Table 4.2 presents a list of the top ten venues extracted from our dataset, along with the corresponding number of EU ETS articles published. The leading journal is Energy Policy, with a total of 64 articles, closely followed by Energy Economics, which has 61 articles. These two journals are prominent in terms of publication volume, underscoring their significant role in disseminating research on the EU ETS and energy markets in more general. However, there is a notable decline in the number of articles in the remaining venues, with the third-ranked journal, Environmental and Resource Economics, publishing 17 articles. Although these other journals still contribute valuable research, they seem to place less focus on EU ETS economics and finance topics, as reflected in their lower article counts.

Bradford’s Law is a bibliometric metric that illustrates the distribution of scientific literature across different venues. First introduced by [27], the

Venues	Articles
Energy Policy	64
Energy Economics	61
Environmental and Revenue Economics	17
Journal of Environmental Economics and Management	15
Environmental Science and Pollution Research	11
International Journal of Global Energy Issues	11
Energy Journal	10
Finance Research Letters	9
Journal of Cleaner Production	8
Revenue and Energy Economics	8

Table 4.2: **Most Relevant Venues per Article.** This table presents the most prolific academic venues, ranked by the number of articles published. It highlights the key sources contributing to the body of research in this field.

law proposes that articles in any given field are not evenly spread across all journals. Rather, a few journals will publish the majority of articles on a specific topic, while a greater number of journals will each contribute a smaller proportion of articles. In this context, the first two venues, Energy Policy and Energy Economics, clearly dominate the article count, highlighting their significant influence and central role in the field compared to other journals.

Figure 4.5 displays the top 10 authors with the highest local citations. The leading author, Chevallier J., has published 25 articles, with a fractional count of 13.90, reflecting their significant influence and active contribution to research on the European Emissions Trading System. Likewise, Hintermann, B. and Mansanet-Bataller, M. are significant contributors, with 8 and 7 articles, respectively, and fractional counts of 5.67 and 2.58.

Furthermore, other authors, such as Delarue E., Ren X. and Wei Y., each have 6 articles, with their fractional counts varying between 1.43 and 2.00. The remaining authors, including Alberola E., Böhlinger C., Feng Z., and Uhrig-Homburg M., each contribute 5 articles. Their fractional counts

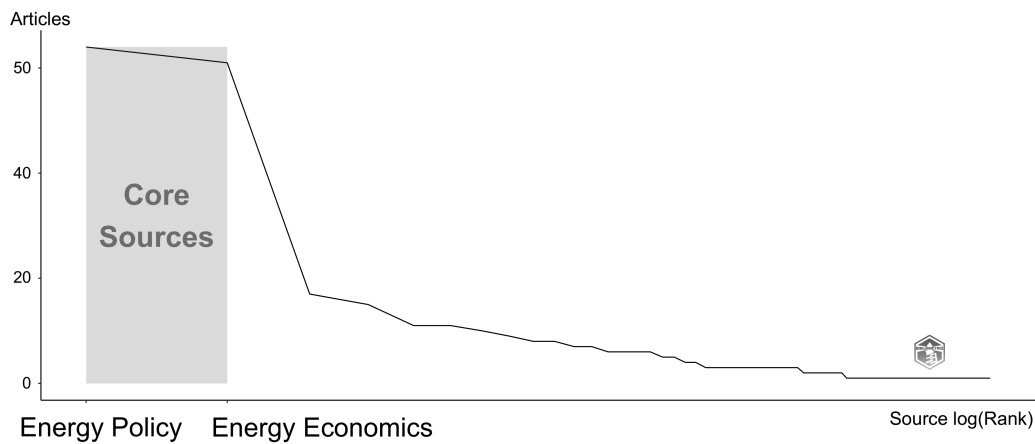


Figure 4.4: **Core Sources by Bradford’s Law.** This plot illustrates Bradford’s Law, which characterises the distribution of research articles across various sources. It highlights the concentration of publications in a small number of core sources, with the remainder distributed across a larger number of less prolific sources.

range from 1.57 to 2.57, indicating their relevant but comparatively modest influence in the field compared to the leading authors.

Lotka’s law, as described by [124], is shown in Figure 4.6. This metric highlights author productivity within a scientific field, stating that the number of authors publishing a specific number of papers is inversely proportional to the square of that number. In other words, a small group of authors is highly productive, while the majority contribute fewer works, leading to a power-law distribution. As a result, scientific output is concentrated among a few prolific authors, such as Chevallier J., while most authors contribute less.

Figure 4.6 presents a network analysis of author collaborations, emphasising key metrics such as cluster membership, betweenness centrality, closeness centrality, and PageRank for each author.




Figure 4.5: **Most Relevant Authors.** This plot highlights the most prolific authors in the field, based on their publication counts. It also showcases the impact of their work, as measured by the total number of citations, providing insight into both productivity and influence.

In Cluster 1, Edenhofer O. and Pahle M. have betweenness centrality values of 2, signifying their role as key intermediaries within the network. Their closeness centrality values of 0.25 indicate that they maintain strong connectivity with other authors in the cluster, facilitating the flow of information and collaboration within the network. These metrics highlight their central position and importance in bridging various research connections within this cluster.

In Cluster 2, Chevallier J. stands out with a notably high betweenness centrality of 62, underscoring their crucial role in connecting different segments of the network. However, their closeness centrality of 0.0588 indicates that they have fewer direct connections compared to other authors, suggesting that their influence is exerted more through indirect pathways. Despite this, Chevallier J. holds the highest PageRank in the dataset, with a value of




Figure 4.6: **Most Relevant Authors.** This plot highlights the most prolific authors in the field, based on their publication counts. It also showcases the impact of their work, as measured by the total number of citations, providing insight into both productivity and influence.

0.0809, reinforcing their significant influence within the network. Wei Y. and Huang W. exhibit betweenness centrality values of 22 and 0, respectively, with closeness values of 0.04 and 0.0435, indicating moderate connectivity and influence within the network. Lastly, Alberola E. has a betweenness centrality of 1 and a closeness of 0.0385, reflecting a more modest role in terms of network intermediation.

4.4 Document Analysis

As shown in figure 4.7, the most cited paper is [3] published in Energy Policy, with a total of 466 citations (TC) and a high citation rate of 27.41 TC per year. The next is [19] published in Energy Economics, with 446 total citations, the highest TC per year at 27.88, and a normalised TC of 4.25, demonstrating its significant influence in the field.

The paper by [66] in Energy Economics has the highest normalised TC of 4.62, indicating its strong influence relative to both the field and its publica-




Figure 4.7: **Most Globally Cited Documents.** This plot presents an overview of the top 10 most cited documents worldwide, showing the total number of citations each paper has received across all sources.

tion year. In addition, [109] published in the Journal of Cleaner Production boasts an impressive citation rate of 36.00 per year and a high normalised citation count of 3.94, highlighting its significant impact.

Not only that, research articles by [129] and [34] demonstrate consistent citation over the years, with TC per year values of 16.94 and 16.69, respectively, reflecting their ongoing recognition in the field. These citations span a variety of journals, including Energy Policy, Energy Economics, Journal of Environmental Economics and Management, and Ecological Economics.

The analysis in figure 4.8 emphasises the most prevalent terms in the dataset, shedding light on key themes in EU ETS research. Terms such as "European Union", "emissions trading," and "gas emissions" highlight a focus on regulations and environmental issues. The frequent appearance of "carbon," "costs" and "emission control" reflects an emphasis on the economic and practical dimensions of emissions management. Furthermore,

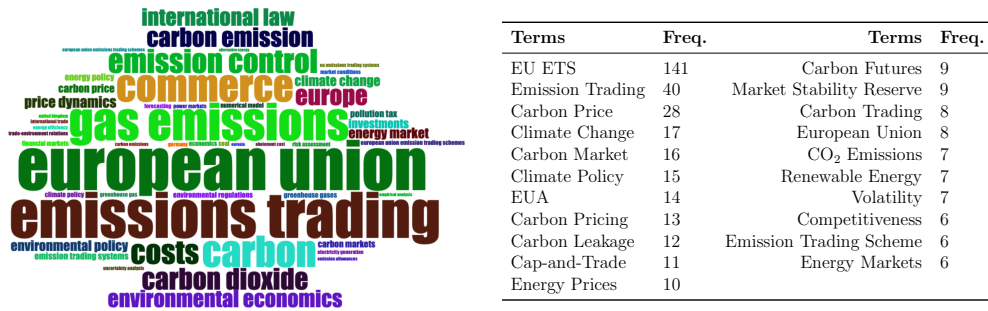


Figure 4.8: **Word Cloud and Term Frequency.** This figure consists of two parts: on the left, a word cloud that visually emphasises the most prominent themes and terms within the dataset, with larger words indicating higher frequency. On the right, a bar chart displays the frequency of specific terms, providing a quantitative overview of the term distribution across the textual data.

terms like "environmental economics," "international law," and "price dynamics" indicate an interdisciplinary approach linking economics, law, and market behaviour with climate-related topics.

Figure 4.10 illustrates the frequency of various terms associated with emissions and environmental policies during the study period. The term European Union has shown a consistent rise in frequency, starting with just one mention in 2004 and reaching 235 mentions by 2024.

Similarly, the term *Emission trading* has experienced a steady rise, increasing from only one occurrence in 2004 to 215 occurrences in 2024, reflecting the growing importance and discussion around this concept over the years. The terms *Gas emissions* and *Commerce* have followed a similar trajectory, beginning with very few or no mentions in the early years of the study period and gradually growing to 175 and 149 occurrences, respectively, by 2024. This pattern suggests a growing recognition of the role of gas emissions and commerce in broader environmental and policy discussions.

The term *Carbon* also shows a significant increase, starting from zero

occurrences in 2004 and increasing to 146 occurrences by 2024, indicating an increased focus on carbon-related issues, particularly in the context of climate change and emissions trading. Furthermore, terms such as *Costs*, *Emission control*, and *Carbon dioxide* have demonstrated substantial growth over the years. In 2024, these terms reached 113, 112, and 101 occurrences, respectively, highlighting the increasing emphasis on the economic and regulatory aspects of emissions reduction, as well as the growing recognition of the importance of controlling carbon dioxide emissions in global environmental policies. This general trend underscores the growing focus on the practical and economic dimensions of addressing climate change, as well as the regulatory and scientific efforts to reduce emissions.

The term *Europe* has experienced a steady increase, increasing from 1 occurrence in 2004 to 92 in 2024. Similarly, the term *Carbon emission* has grown from zero occurrences in 2004 to 90 occurrences in 2024.

4.5 Conceptual Structure

Figure 4.6 presents the co-occurrence network analysis based on author keywords. The plot reveals distinct groups of keywords that are frequently connected in the scientific literature, with the blue group showing a strong predominance. Two measures are discussed: the PageRank value, which assesses the influence of author keywords based on their connections, highlighting their centrality and importance within the network, and betweenness centrality, which quantifies how often a node acts as a bridge between other nodes, indicating its role in connecting different parts of the network.

In these networks, certain keywords, such as *EU ETS* and *Kyoto Protocol*, exhibit high betweenness values, highlighting their prominent role as

intermediaries connecting disparate areas of the network. These keywords serve as bridges, linking different thematic clusters and facilitating the flow of knowledge across various subfields.

Specifically, *EU ETS* stands out with a remarkably high betweenness value of 3437.60, positioning it as a central hub with significant influence on the flow of information within the network. This centrality is further reinforced by its substantial PageRank value of 0.273, highlighting its pivotal role in organising and shaping the structure of the network.

In contrast, nodes such as *carbon prices* and *energy markets* exhibit lower betweenness and PageRank values, indicating that while less central, they still play an important role in the overall network. This suggests a network characterised by dominant themes like *EU ETS*, with other topics acting as supporting elements that reinforce and expand the network's structure.

Figure 4.9 offers a clear visualisation of the main themes and their significance within the co-occurrence network, based on the author keywords. This thematic map highlights the distribution and centrality of key terms, helping to identify the relative importance of different research themes in the academic literature.

In Cluster 1 (*cap-and-trade*), key themes like *Cap-and-Trade*, *Carbon trading*, and *Emissions Trading Scheme* are central to discussions on carbon markets and regulations. The high betweenness centrality of *Cap-and-Trade* (190.19) highlights its key role in the connection of various ideas in this area. However, terms like *EU ETS reform* have less centrality, suggesting that they play a more specialised or less influential role within this cluster.

Cluster 1 focusses on cap-and-trade systems, with key terms acting as important connectors. Cluster 2 focusses on carbon emissions, carbon tax, and electricity markets, highlighting economic and policy aspects. Terms

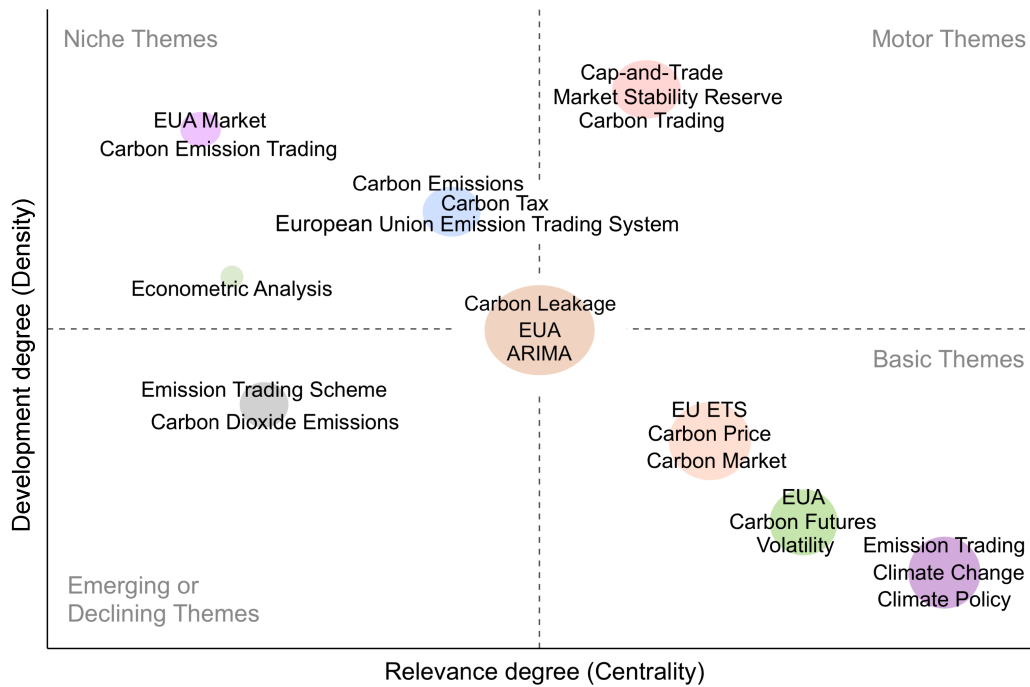


Figure 4.9: **Thematic Map Analysis.** This figure presents a thematic map that illustrates the distribution and centrality of key terms across various research clusters. It categorises the research areas into four main types of themes: niche themes, motor themes, basic themes, and emerging or declining themes. This classification allows for a clearer understanding of the relative importance, development, and dynamics of different thematic areas within the literature, highlighting how they evolve and interact over time.

such as *carbon emissions* and *carbon tax* play a key role in connecting ideas about carbon pricing and market impacts.

Cluster 2 also includes specialised terms like *Cost pass-through* and *Market power*, which add depth to the discussion on carbon economics. Cluster 3 focusses on the European Union Allowance (EUA) and related topics, with terms like *EUA*, *Carbon futures*, and *Volatility* playing central roles in the network. The term *EUA* has a high Betweenness Centrality of 43.37, highlighting its key role in EU carbon market discussions. Similarly, *Carbon*

futures and *Volatility* are important, as their centrality shows their relevance to market fluctuations and financial aspects of carbon trading.

In general, thematic map analysis shows the varying influence of terms across clusters. Cluster 1 focusses on cap-and-trade system policy, Cluster 2 on the economic aspects of carbon emissions, and Cluster 3 on European carbon market mechanisms.




Figure 4.10: **Word Frequency Over Time.** This plot shows the evolution of the frequency of key terms related to emissions, carbon, and environmental policy over a two-decade period. It highlights the growing emphasis on these critical themes, providing insight into how the focus of research has shifted and intensified in response to global environmental concerns.

4.6 Discussion

During the past two decades, there has been a noticeable increase in both the frequency and impact of certain key terms and topics that have played a central role in shaping the discussion on emissions trading. A strong focus on market mechanisms and regulations within the European Union reflects a shared interest among researchers and policymakers in the use of carbon pricing as a primary tool to combat climate change. The foundational works

of [3] and [19] have emerged as some of the most cited contributions, marking crucial advances in the field, particularly in terms of market dynamics and policy impacts. Their influence is evident in their high-citation metrics, which underscore their importance in shaping both academic and policy perspectives on the European carbon market. Likewise, word frequency analysis highlights a growing emphasis on terms such as *European Union*, *Emissions Trading*, and *Carbon*, reflecting the increasing prominence of carbon pricing mechanisms and the carbon market within the broader context of climate policy. This trend underscores the centralisation of climate policy discussions around carbon pricing as an economically viable solution for large-scale emissions reduction. The consistent appearance of these terms over time suggests a stable, if not expanding, alignment between economic and environmental policy-making communities, signalling a shared recognition of the potential of emissions trading systems in addressing climate change.

In parallel, there has been a growing focus of research on specific market behaviours and their implications for policy, as evidenced by the increasing frequency of terms such as *Carbon Price* and *Carbon Market*. This increasing interest highlights an intensifying examination of how carbon pricing influences not only emissions levels but also economic factors, including sectoral competitiveness, cost-effectiveness, and overall market behaviour under regulatory constraints. The evolving focus on pricing strategies and their economic effects underscores the complexity of balancing environmental objectives with economic stability, a recurring theme that has garnered significant attention in the recent literature.

Indeed, trend analysis has uncovered a gradual shift in research focus towards the financial aspects of the EU ETS. Topics such as *Carbon Leakage*, *Energy Prices*, and *Carbon Futures* have attracted increasing attention, indi-

cating that scholars are now more engaged in the financial and secondary effects of carbon trading, extending beyond just environmental or policy themes. This shift reflects a more nuanced understanding of the EU ETS, which now encompasses not only emissions control, but also complex financial dynamics, such as market volatility, investor behaviour, and risk assessment in carbon trading. These financial factors highlight the growing role of EU ETS in the economy, presenting emissions trading as both a regulatory tool and a financial instrument shaped by market trends and economic changes.

The network and thematic map analyses offer a clearer view of the conceptual structure within the EU ETS literature. Central themes such as *Carbon Leakage* and *EUA* serve as critical nodes, connecting diverse research themes and underscoring their importance in shaping the direction of research. Their high Betweenness Centrality and PageRank values highlight their significant role in linking different areas of study. This connectivity reflects a shared interest in understanding the indirect consequences of emissions trading, such as shifts in production patterns, competitiveness, and trade imbalances. In contrast, terms like *Carbon Prices* and *Energy Markets*, while influential, play a more supportive role in the network. This suggests potential opportunities for future research to further explore these secondary yet interconnected factors, providing a deeper understanding of their role within the broader emissions trading system.

Despite these important contributions, several research gaps remain unresolved.

There is a noticeable gap in the literature regarding the use of non-parametric methodologies, particularly in the areas of feature selection and forecasting models. This gap suggests that traditional, parametric approaches may be limited in capturing the non-linear and often complex relationships

inherent in emissions trading data. For instance, by employing non-parametric quantile regression based on copulas, [148] captures non-linear relationships and dependencies among variables while controlling for firm characteristics, addressing potential omitted variable biases. Non-parametric methods, in particular, may be more effective in modelling and capturing extreme market conditions, which have become more frequent due to recent economic events and policy shifts.

Another non-parametric approach is introduced by [145], who examines how the COVID-19 pandemic and energy crises have altered market dynamics, making EU ETS prices more sensitive to energy factors. The study also proposes a new non-parametric method for feature selection and forecasting, leveraging the information content of financial variables. Additionally, [121] introduces a novel secondary decomposition integration framework, validating the effectiveness of non-linear integration and optimisation methods to enhance model performance. Although non-parametric approaches offer flexible alternatives to parametric methods, they remain relatively unexplored in the context of emissions trading, representing a promising area for future research.

Emissions trading involves complex and often incomplete datasets, and methodologies for effective imputation and aggregation are underdeveloped in the EU ETS literature. For example, [24] imputed missing values to include high-emitting installations likely to lose free allocations due to policy changes, resulting in more accurate regulatory impact assessments. This approach helps prevent the underestimation of the dynamics of carbon leakage, as discussed in [20]. By refining the imputation methods, future research could generate more reliable datasets, enabling more precise evaluations of market trends and policy impacts. Improved data quality would strengthen

both theoretical models and practical policy recommendations, improving the robustness and credibility of research results in the field.

Another potential research gap lies in examining how exchange rates interact with the dynamics of EU ETS, which could provide valuable information on cross-border trade effects and market competitiveness. In a recent study, [70] explores the influence of exchange rates on carbon prices, highlighting the Euro index as a significant driver, particularly during the Brexit crisis. Currency fluctuations can signal broader macroeconomic conditions that can affect the stability of the carbon market. Future research could investigate how global economic shifts, including exchange rate volatility, impact the carbon market, offering information on potential vulnerabilities and providing guidance to investors and policymakers on mitigating currency-related risks in emissions trading systems. Macroeconomic factors such as inflation, GDP growth, and interest rates can influence the pricing and trading of carbon allowances. [23] emphasises the significant impact of these conditions on carbon price dynamics, stressing the importance of integrating these variables when assessing the effectiveness of the EU ETS in meeting climate targets. By understanding the interplay between macroeconomic factors and carbon pricing, researchers could offer clearer insights into how emissions trading performs in different economic cycles. This could potentially lead to the development of more resilient policy frameworks that are better equipped to adapt to economic fluctuations, ensuring the stability and effectiveness of the EU ETS over time.

Another notable research gap lies in exploring the causal relationships between the EU ETS and other financial variables. In a recent study, [32] applied Granger causality tests in the frequency domain to distinguish between short- and long-term effects of climate policies. Financial instruments

such as derivatives, equities, and bonds may be influenced by or may influence the dynamics of the EU ETS. For instance, [171] found that the causality between crude oil and stock markets varies by equity index, remaining robust under normal conditions but weakening during extreme market scenarios. As emissions trading markets continue to mature, understanding these causal relationships could offer valuable insights into how external financial shocks and policy changes impact the EU ETS. This would be crucial for shaping future regulatory approaches and informing investment strategies in the carbon market.

In summary, these gaps and trends underscore key areas for future research to deepen our understanding of the EU ETS. Addressing issues related to methods, data quality, economic factors, and causal relationships within emissions trading is crucial for improving how the EU ETS can contribute to carbon reduction and climate policy objectives. Focusing on these areas can improve the effectiveness, resilience, and adaptability of emissions trading systems, ensuring that they are better equipped to respond to evolving global economic and environmental conditions.

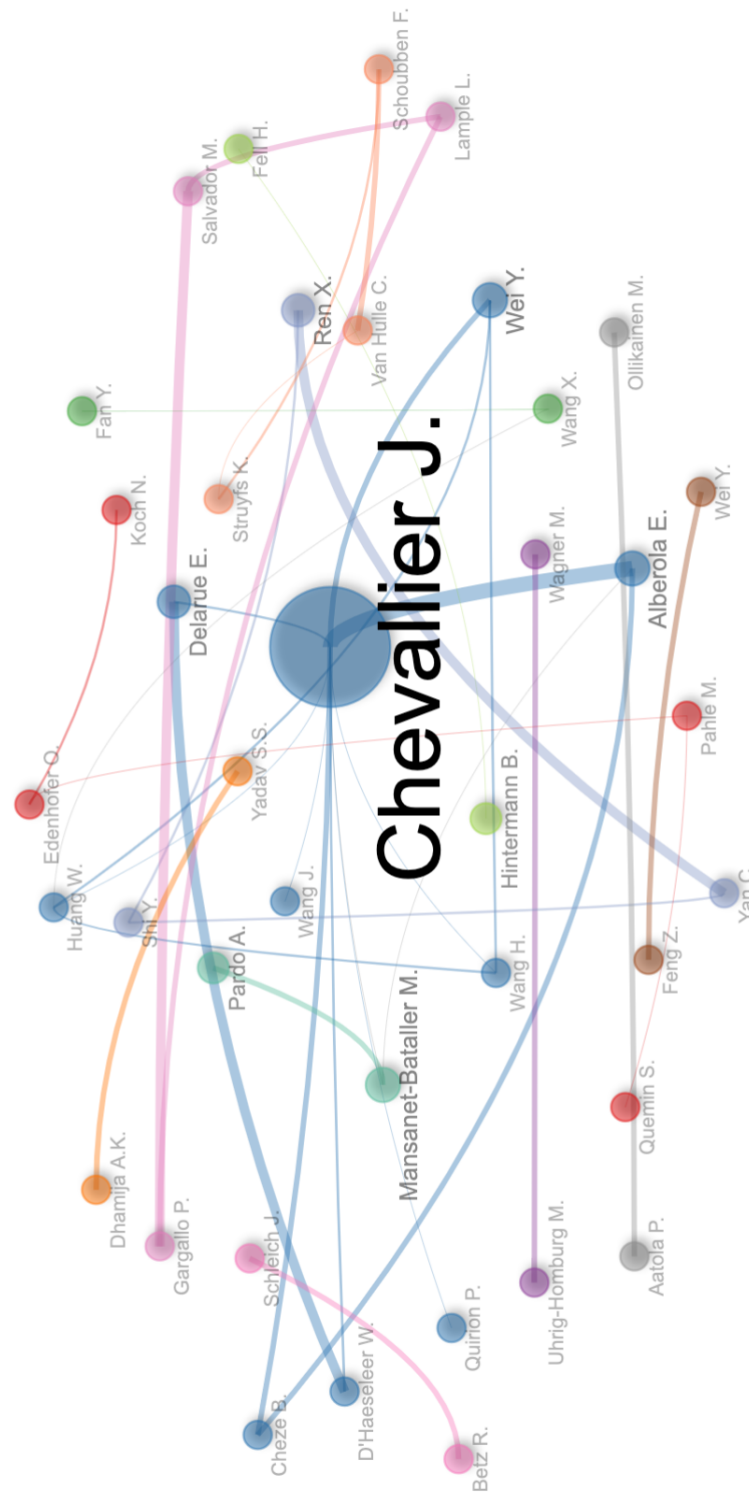


Figure 4.12: **Authors collaboration network.** This network visualises the research connections between authors, illustrating the strong collaborations and research groups formed among scholars. Each node represents an author, with lines connecting them to indicate co-authorships. The thickness of the lines reflects the frequency of their collaborations, emphasising key partnerships and the extensive network of ideas within the field.

Chapter 5

Empirical Results II: Price Determinants and Forecasting in the EU ETS

5.1 Data

The daily closing prices of European Union Allowances (EUA) are collected from Bloomberg[®], covering the time period from January 2014 to April 2023. This comprehensive dataset includes 237 observations, exceeding the temporal scope studied in [164]. In line with previous research focused on EUA price trends, our analysis deliberately omits data from Phase 1 and Phase 2 of the market. The rationale behind this exclusion is supported by evidence that indicates that during these phases, price variations were predominantly driven by regulatory changes and policy changes, since the market was in an experimental phase.

In line with previous research efforts, such as those detailed in [37], [30], [155], [140, 141], and [4], the 33 distinct predictors have been grouped into

ID	Category	Variables	Start to end	Database
0	T	EUA (ICEENDEX)	January 2014 - April 2023	Bloomberg [®]
1	UNC	GPR	January 2014 - April 2023	GPR website
2	UNC	VSTOXX (V2X)	January 2014 - April 2023	Bloomberg [®]
3	UNC	Uncertainty EUR/USD (CAFZUUEU)	January 2014 - April 2023	Bloomberg [®]
4	UNC	Uncertainty EUR/JPY (CAFZUEJP)	January 2014 - April 2023	Bloomberg [®]
5	UNC	Uncertainty EUR/GBP (CAFZUEGB)	January 2014 - April 2023	Bloomberg [®]
6	UNC	Uncertainty EUR/CHF (CAFZUECH)	January 2014 - April 2023	Bloomberg [®]
7	COM	ICE Dutch TTF Natural Gas (TTF0NXHR)	January 2014 - April 2023	Bloomberg [®]
8	COM	Electricity Prices Spain (OMLPDAH)	January 2014 - April 2023	Bloomberg [®]
9	COM	Electricity Prices Germany (EXAPBDHD)	January 2014 - April 2023	Bloomberg [®]
10	COM	Electricity Prices Italy (ELIODAHD)	January 2014 - April 2023	Bloomberg [®]
11	COM	Electricity Prices France (PWNXFRAV)	January 2014 - April 2023	Bloomberg [®]
12	COM	ICE Brent oil futures (CO1 Comdty)	January 2014 - April 2023	Bloomberg [®]
13	COM	ICE Coal Rotterdam futures (TMA Comdty)	January 2014 - April 2023	Bloomberg [®]
14	COM	Gold (GCZ3 Comdty)	January 2014 - April 2023	Bloomberg [®]
15	ER	EUR/USD spot (EUR/USD)	January 2014 - April 2023	Eikon Refinitiv [®]
16	ER	EUR/JPY spot (EUR/JPY)	January 2014 - April 2023	Eikon Refinitiv [®]
17	ER	EUR/GBP spot (EUR/GBP)	January 2014 - April 2023	Eikon Refinitiv [®]
18	ER	EUR/CHF spot (EUR/CHF)	January 2014 - April 2023	Eikon Refinitiv [®]
19	ENR	Bloomberg Energy price return index (EUNRJP)	January 2014 - April 2023	Bloomberg [®]
20	ENR	Solactive ESG Fossil Eurozone 50 index (S0ESG50N)	January 2014 - April 2023	Bloomberg [®]
21	ENR	S&P Eurozone 50 Environmental index (SPEENDET)	January 2014 - April 2023	Bloomberg [®]
22	ENR	MSCI Europe Energy Sector index (MXEU0EN)	January 2014 - April 2023	Bloomberg [®]
23	ENR	ERIX index	January 2014 - April 2023	Bloomberg [®]
24	ENR	EUROSTOXX Electricity index (SXEEELC)	January 2014 - April 2023	Bloomberg [®]
25	CTRY	EUROnext100 (N100)	January 2014 - April 2023	Bloomberg [®]
26	CTRY	IBEX35 (IBEX)	January 2014 - April 2023	Eikon Refinitiv [®]
27	CTRY	DAX	January 2014 - April 2023	Eikon Refinitiv [®]
28	CTRY	CAC	January 2014 - April 2023	Eikon Refinitiv [®]
29	CTRY	FTSE Mib	January 2014 - April 2023	Eikon Refinitiv [®]
30	MACRO	Euro-area 3-month bond yield (BN03)	January 2014 - April 2023	Bloomberg [®]
31	MACRO	Euro-area 10-year bond yield (BN10)	January 2014 - April 2023	Bloomberg [®]
32	MACRO	Euro-area inflation (HICP)	January 2014 - April 2023	Eurostat
33	MACRO	Euro-area GDP (current value)	January 2014 - April 2023	Eurostat

T: Target; UNC: Uncertainty variables; COM: Commodity related variables; ER: Exchange rates; ENR: Energy-related indexes/variables; CTRY: Country indexes; MACRO: Macro-economic variables.

Table 5.1: **Dataset Description.** This table lists 34 distinct time series used in this research, along with their respective sources, organised into seven categories. All time series are collected in euros (EUR).

six broad categories. These categories encompass elements related to geopolitical, economic and financial uncertainties, commodities, specific exchange rates, energy indices, national indices, and include European indexes, as well as broad macroeconomic variables.

All predictors were collected on a daily basis, with the exception of Euro-area inflation, which was obtained on a monthly frequency, and Euro-area GDP, collected quarterly. The reference period aligns precisely with that of the target variable, maintaining an equivalent number of observations for the

predictors collected daily. Specifically, there are 114 observations pertaining to monthly inflation data and 37 observations concerning quarterly GDP data.

Initially, six predictors related to uncertainty, specifically: (1) the GeoPolitical Risk (GPR) index; and (2 to 6) the uncertainty indexes pertaining to significant world exchange rates, including EUR / USD, EUR / JPY, EUR / GBP and EUR / CHF, are examined. Data for these uncertainty variables are sourced from Bloomberg, while the GPR index is obtained from its dedicated website [69]. Next, our analysis incorporates eight commodities that may not be directly related to energy: (7) ICE Dutch natural gas futures; (8 to 11) electricity prices for the countries of Spain, Germany, Italy, and France; (12) ICE Brent oil futures; (13) ICE Rotterdam coal futures, and (14) the gold index. In the third step, significant spot interest rates are considered to assess potential purely financial and highly volatile impacts, specifically for (15 to 18): EUR/USD, EUR/JPY, EUR/GBP, EUR/CHF. Fourth, six European energy indices are considered to assess the information content provided by energy indices of diverse characteristics: (19) Bloomberg Energy price return index; (20) Solactive ESG Fossil Eurozone 50 index; (21) S&P Eurozone 50 Environmental index; (22) MSCI Europe Energy Sector index; (23) ERIX index; and (24) EUROSTOXX Electricity index. Subsequently, the influence of four industrial nations and a European index is assessed to delve into the predictive capacities of financial activities: (25) EUROnext100; (26) IBEX35; (27) DAX; (28) CAC; and (29) FTSE Mib. Lastly, the impact of the economic cycle and European macroeconomic conditions is analysed, taking into account: (30) Euro-area 3-month bond yield; (31) Euro-area 10-year bond yield; (32) Euro-area inflation; and (33) Euro-area GDP.

Variables (32) and (33) are obtained at a frequency that does not match

the daily interval of the European Union Allowance (EUA) target variable.

Because inflation (monthly) and GDP (quarterly) are measured at different frequencies, analysing their impact on EUA prices would require extensive data imputation. However, such interpolation risks introducing errors or losing critical information, which could undermine the reliability of the results. To maintain robustness, we excluded these variables from the empirical analysis and focused only on predictors that did not require temporal transformation.

In Section 5.4, both Euro-area inflation and Euro-area GDP are incorporated into the analysis. Our process of selecting variables was guided by a thorough examination of the existing literature and an extensive understanding of market dynamics. For further elucidation, the appendix provides an illustrative example of imputation and aggregation techniques A.3.

The successful integration of new metrics such as GDP and Inflation, in conjunction with traditional ones, effectively tackles the challenge posed by mixed frequency in examining the macroeconomic dynamics associated with EUA. Not only that, incorporating key macroeconomic indicators such as GDP and inflation when investigating the determinants of EU ETS prices is crucial to acquire a comprehensive understanding of the wider economic context. This approach helps to discern market sentiments and anticipate the implications of policies related to emissions trading [117].

5.1.1 Descriptive Statistics

Table 5.2 offers an extensive overview of the descriptive statistics for all predictor variables along with the European Union Allowance (EUA). Within this table, a comprehensive summary is provided, detailing the mean value and the standard deviation (STD), which serves to portray the variation

ID	Variables	Mean	STD	Min	25%	50%	75%	Max
0	EUA	27.31	27.55	3.93	6.41	17.59	32.59	100.29
1	GPR	113.50	52.80	9.49	79.90	103.89	136.51	542.66
2	VSTOXX	20.87	7.38	10.68	15.88	19.45	24.03	85.62
3	Unc. EUR/USD	2.66	0.60	1.57	2.21	2.57	3.05	4.28
4	Unc. EUR/JPY	2.96	0.60	1.32	2.64	2.99	3.31	5.69
5	Unc. EUR/GBP	2.42	0.64	1.36	1.94	2.37	2.72	6.64
6	Unc. EUR/CHF	1.85	0.73	0.98	1.49	1.74	1.98	8.92
7	Natural Gas	33.90	41.47	3.63	14.90	19.15	24.57	311
8	Elec. Prices Spain	69.82	53.82	1.10	42.61	52.34	65.55	544.98
9	Elec. Prices Germany	69.23	83.68	-9.12	31.53	39	58.61	682.89
10	Elec. Prices Italy	92.18	96.61	10.66	46.95	55.67	74.10	718.71
11	Elec. Prices France	92.18	96.61	10.66	46.95	55.67	74.10	718.71
12	Brent oil	66.97	21.97	17.32	50.19	63.49	79.17	133.89
13	Coal futures	118.37	97.02	48.50	64.51	82.48	109.83	457.80
14	Gold	1,457.71	274.68	1,051.10	1,240.92	1,319.00	1,760.34	2,063.54
15	EUR/USD spot	1.15	0.08	0.96	1.10	1.13	1.18	1.39
16	EUR/JPY spot	129.74	8.11	111.15	123.42	129.68	135.82	149.18
17	EUR/GBP spot	0.84	0.05	0.69	0.83	0.86	0.88	0.94
18	EUR/CHF spot	1.10	0.06	0.95	1.06	1.09	1.14	1.24
19	Bloomberg Energy price return index	1.10	0.06	0.95	1.06	1.09	1.14	1.24
20	Solactive ESG Fossil Eurozone 50 index	101.58	17.01	48.01	91.17	106.05	113.87	130.32
21	S&P Eurozone 50 Environmental index	127.29	26.58	84.21	107.34	122.45	139.70	199.47
22	MSCI Europe Energy Sector index	1,477.89	152.24	1,059.38	1,350.86	1,467.14	1,560.97	1,870.90
23	ERIX index	1,325.94	617.13	567.78	840.82	1,028.58	1,949.09	3,106.55
24	EUROSTOXX Electricity index	314.53	67.26	212.46	253.71	284.09	379.67	471.78
25	EUROnext100	1,034.78	156.67	733.93	899.78	1,016.07	1,133.29	1,388.09
26	IBEX35	1,034.78	156.67	733.93	899.78	1,016.07	1,133.29	1,388.09
27	DAX	12,216.79	1,870.47	8,441.71	10,685.23	12,238.56	13,263.39	16,271.75
28	CAC	5,337.46	840.69	3,754.84	4,615.03	5,234.00	5,882.29	7,577
29	FTSE Mib	21,623.98	2,699.71	14,894.44	19,727.58	21,674.43	23,329.29	28,162.67
30	Euro-area 3-month bond yield	-0.38	0.68	-1.13	-0.71	-0.59	-0.29	4.55
31	Euro-area 10-year bond yield	0.39	0.74	-0.85	-0.21	0.32	0.68	2.75

Table 5.2: **Descriptive Statistics.** This table presents the descriptive statistics for a collection of 34 distinct time series. For each time series, the analysis includes the mean, standard deviation, minimum, and maximum values. Additionally, the values for three specific percentiles are provided to offer a deeper understanding of the data distribution and variability.

within the data. In addition, it enumerates the minimum value recorded. Also, the table delineates the 0.25, 0.5, and 0.75 percentiles, capturing the distribution spectrum of the dataset. Finally, it enumerates the maximum value observed, enhancing the completeness of the statistical description.

Figure 5.1, presents a detailed view of the correlations observed between each explanatory variable and the EUA price in Phase 3. Notably, it is evident that the variables categorised as commodities demonstrate a significant correlation with the EUA price. This outcome is anticipated since they belong to the same class of commodities. In contrast, in the course of Phase 3,

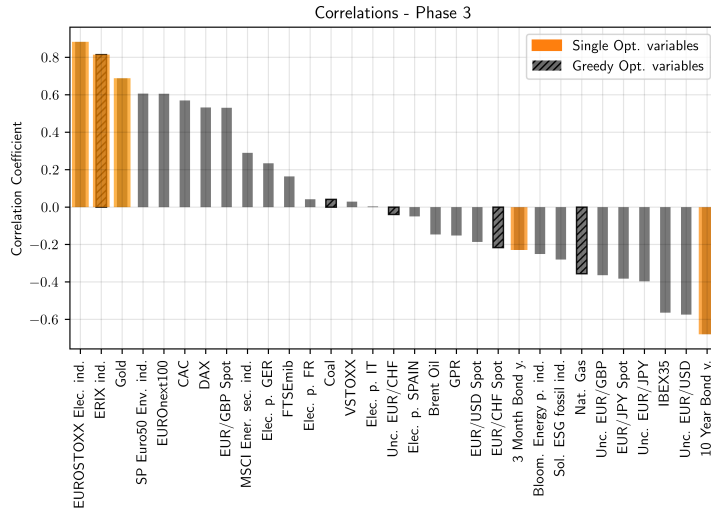


Figure 5.1: **Correlation Analysis - Phase 3.** This figure examines the correlations between the daily dataset and EUA. It highlights the variables that provide the greatest insights individually, as identified using the Information Imbalance (II) method, which are shown in orange in Fig. 5.3. Additionally, the dashed bars represent the top five most insightful variables, selected using a greedy algorithm based on the Information Imbalance approach. This variable selection process is further illustrated in Fig. 5.5, corresponding to Phase 3 of the analysis.

the variables associated with uncertainty display a minimal correlation with the target variable.

The Phase 4 correlations are presented in figure 5.2. In this case, most of the variables examined exhibit positive correlations, many of which are notably strong. However, exchange rates show limited correlation with the target variable during Phase 3, despite the fact that certain rates, such as the EUR / CHF spot and the EUR / USD spot, are markedly negatively correlated in Phase 4. Furthermore, a significant shift in correlation can be discerned for indices such as the EUROSTOXX Electricity prices index, the ERIX index, and the S&P Euro50 Environmental index. These indices transition from a strong positive correlation in Phase 3 to a strong negative




Figure 5.2: **Correlation Analysis - Phase 4.** This figure examines the correlations between the daily dataset and EUA. It highlights the variables that provide the greatest individual insights, as identified using the Information Imbalance (II) method, which are shown in orange in Fig. 5.4. Additionally, the dashed bars represent the top five most insightful variables, selected using a greedy algorithm based on the Information Imbalance approach. This variable selection process is further detailed in Fig. 5.6, corresponding to Phase 4 of the analysis.

correlation in Phase 4, marking a substantial change.

In this section, it is important to note that the correlation analysis performed serves a dual purpose. Firstly, it establishes a foundational benchmark for assessing the significance of the explanatory variables contained within our dataset. Secondly, it underscores the capability of Information Imbalance to identify non-linear relationships, which, due to their inherent nature, are overlooked by traditional correlation methods.

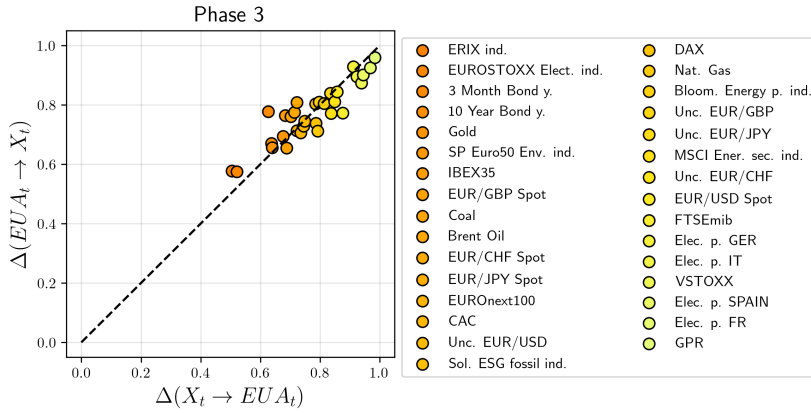


Figure 5.3: **Daily Information Imbalance Analysis - Phase 3.** This figure evaluates the informational value of each variable within the dataset. The legend ranks the variables from most to least informative, with darker orange indicating higher informativeness and lighter yellow-orange representing less informative variables. This analysis is specifically presented for Phase 3.

5.2 Information Imbalance Analysis

In contrast to conventional parametric econometric models, the concept of Information Imbalance is distinguished by its independence from assumptions regarding the underlying data-generating process. Rather than imposing predefined assumptions, it allows the use of data in its original form, thereby ensuring that the analytical results remain completely unmodified. This specific attribute stands out as a primary advantage intrinsic to the non-parametric approach being employed, as it inherently imposes fewer constraints compared to any form of parametric analysis.

Figures 5.3 and 5.4 displays the Information Imbalance (II) planes corresponding to Phase 3 and Phase 4, respectively. Within these graphs, the Information Imbalance metrics are illustrated, detailing the relationship between each individual explanatory variable X_t and the target variable EUA_t . The x -axis represents the direction of Information Imbalance from the set of predictors to the target set, denoted as $\Delta(X_t \rightarrow EUA_t)$, while the y -axis

represents the reverse flow, indicated as $\Delta(EUA_t \rightarrow X_t)$. Analysing these graphs allows for a fascinating comparison of the two phases, highlighting the significant distinctions regarding the informativeness of variables between these two stages.

With respect to Phase 3, the ERIX index emerges as the most insightful variable. It tracks the progress of European renewable energy firms engaged in one or more of six distinct investment clusters: biofuels, geothermal, marine, solar, hydro, and wind energy. In their study, [112] illustrated the independence of European Union Allowance (EUA) shares from the ERIX index by examining their time-dependent correlation via a GO-GARCH model. However, our research indicates that the ERIX index has a significant informational value in relation to EUA, much higher than other variables considered, suggesting potential novel interpretations.

Specifically, it is possible to assert that the price movements of the ERIX index exhibit a close similarity to those observed in the EUA, indicating analogous market dynamics between the two. This observation is noteworthy because, while renewable energy sources are not directly involved in emission trading schemes, their importance cannot be overlooked in influencing the larger market dynamics [44]. The adoption and encouragement of renewable energy sources play a critical role in achieving a reduction in greenhouse gas emissions within the energy sector [95].

As a result, these initiatives enhance the goals of the EU Emissions Trading Scheme, which is designed to reduce emissions originating from industrial sectors. In specific cases, renewable energy initiatives produce carbon credits or offsets, representing the decrease or prevention of greenhouse gas emissions. Entities involved in the EU emissions trading system have the opportunity to use these credits to counterbalance a segment of their emissions,

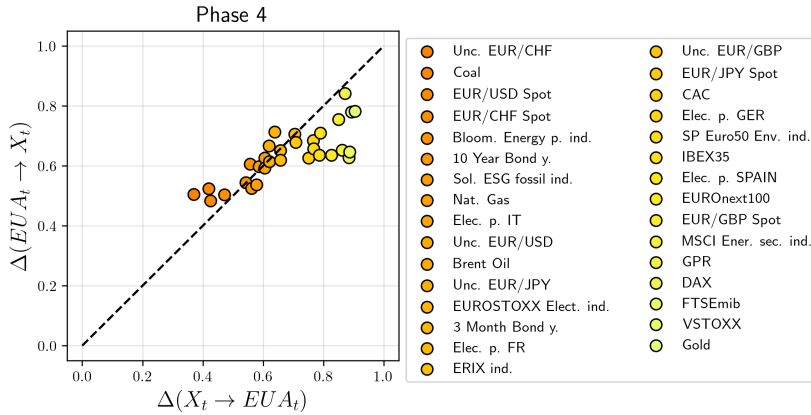


Figure 5.4: **Information Imbalance by Variable in Phase 4.** This figure quantifies the relative informational value of individual variables in predicting system dynamics during Phase 4. Variables are ranked from highest to lowest information content (left to right in the legend), with colour intensity indicating their relative importance—darker orange signifies higher information contribution, while lighter yellow-orange denotes lower contribution.

thereby achieving compliance requirements with greater efficiency.

The requirement for emission allowances within the European Union Emissions Trading System (EU ETS) is significantly affected by the mix of energy sources used [99]. An increased adoption of renewable energy sources can result in a noticeable decrease in emissions originating from the electricity generation sector. This, in turn, alters the dynamics of supply and demand in the emissions trading market. The European Union is actively involved in harmonising renewable energy policies with its goals of reducing emissions.

As an illustration, the directive outlined in [51] establishes legally binding objectives for the proportion of renewable energy within the total final energy consumption of the European Union. This strategic integration guarantees that initiatives aimed at fostering the adoption of renewable energy sources are harmonised with the overarching climate objectives, such as those

associated with the EU Emissions Trading Scheme. Also, the EU Emissions Trading Scheme provides economic motivations for investing in technologies that result in cleaner energy production, spanning renewable energy projects, thereby driving progress in sustainable energy development. Corporations that opt to invest in renewable energy sources have the opportunity to gain benefits not only from the commercial distribution of clean energy; they may also experience favourable results from potential profits derived from the transaction of emission allowances or the monetization of carbon credits [55]. Furthermore, during Phase 3, it is observed that the most insightful and significant individual variables are intricately linked with the intricacies of the European energy market infrastructure, including indices like the EUROSTOXX Electricity index.

Additionally, variables such as the yields on 3-month and 10-year bonds are examined, which signify the interest rates that investors demand to hold Eurozone government bonds maturing in either 10 years or 3 months, respectively. This particular metric is of considerable importance, as fluctuations in bond yields offer critical insights into market expectations regarding economic conditions, inflation rates, and monetary policy developments [6].

During Phase 4, a scenario is encountered that is distinct but shares similarities in certain respects. In particular, it is evident that the variables that represent financial variations assume considerably greater significance compared to their role in Phase 3.

Uncertainty in the context of exchange rates, such as the observed unpredictability in the EUR/CHF pair, is characterised by a diminished confidence or inability to accurately forecast future fluctuations in the currency pair's movements. The presence of high volatility in an exchange rate serves as an indicator of increased uncertainty. Instruments such as exchange rate op-

tions, with a specific emphasis on implied volatility, provide valuable insight into the future anticipated movements of currency pairs as perceived by market participants [15]. Global shifts in risk sentiment, frequently mirrored by shifts in stock market trends, have the potential to affect the demand for currencies considered as safe-havens. When there is uncertainty surrounding worldwide economic conditions, this often results in increased volatility within currency markets. This is particularly true, as substantial speculative trading activities or abrupt swings in market sentiment can exacerbate these uncertainties. Such quick alterations in market sentiment, driven by speculative actions, can consequently result in unpredictable fluctuations in currency values [84].

The distinct choice of informative variables highlights how the COVID-19 pandemic and the concurrent energy crisis have dramatically altered the pricing dynamics, thus establishing novel determinants for the price of EUA. In response, industries regulated by the Emissions Trading System (ETS) can investigate the potential for transitioning from coal to cleaner energy options, including natural gas or renewable sources. This shift is motivated by both economic factors and environmental objectives. The economic feasibility of these substitutions may experience significant impact due to variations in coal prices [22]. Although both markets are subjected to comparable influences, they are not directly interlinked. However, they can be indirectly influenced by larger-scale dynamics within the energy market. These broader dynamics encompass changes in supply and demand, geopolitical incidents, as well as prevailing economic conditions [7].

Changes in regulations concerning carbon emissions and the use of coal have the potential to impact both markets; particularly, stringent regulatory measures could increase costs and modify the demand for European Union

Allowance (EUA) permits. A pivotal element in this dynamic is the shift towards renewable energy sources, which is expected to decrease reliance on coal and, consequently, affect the carbon market. Beyond that, each market is subject to a wide array of global influences, including international trade dynamics, fluctuations in energy prices, and varying climate change policies [44].

The evolving energy framework is progressively focussing on cleaner and more sustainable energy sources, which has significant potential to alter the demand for coal and influence the evaluation of EUA permits. Renewable energy sources, including wind, solar, hydroelectric, and bio-energy, are characterised as low-carbon or carbon-neutral. The adoption and endorsement of these renewable sources plays a crucial role in the reduction of greenhouse gas emissions within the energy sector [95]. The transition to renewable energy sources has been a focal point in global energy policy discussions driven by the need to address environmental concerns and improve energy security. As [99] highlights, the integration of renewable energy into existing systems presents both opportunities and challenges, particularly in balancing economic growth with sustainability goals. This underscores the importance of strategic planning and innovation to maximise the benefits of the adoption of renewable energy.

5.3 Greedy Selection of Variables

The application of the greedy selection method in this analysis aims to pinpoint the most informative subset of explanatory variables to predict EUA prices. This approach, as depicted in figures 5.5 and 5.6, involves an iterative process.

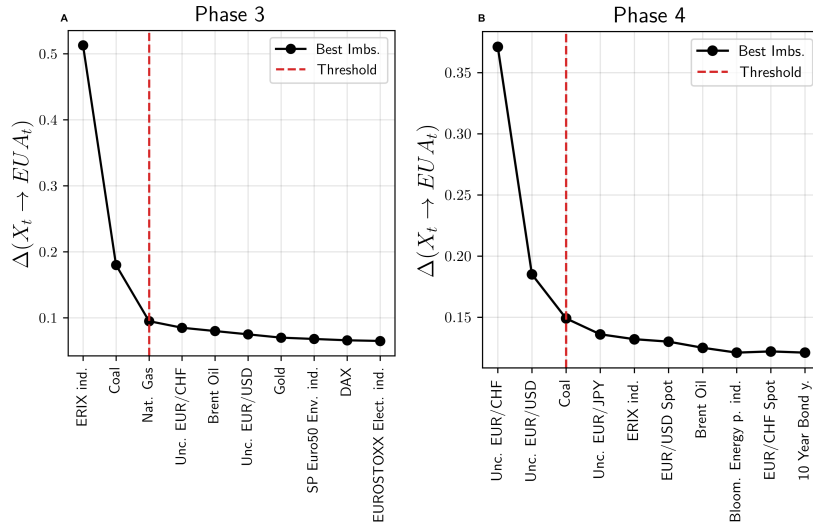


Figure 5.5: Information Imbalance Analysis of Phase 3 and Phase 4 EUA Price Determinants. This figure illustrates the Information Imbalance, denoted as $\Delta(X_t \rightarrow EUA_t)$, as it moves from the set of predictors to the European Union Allowance (EUA) price. The analysis tracks the changes in imbalance as an increasing number of variables are added to the set X_t . The x-axis labels correspond to each variable added to the set, providing a clear view of their cumulative effect on the EUA price.

In each step, the algorithm progressively constructs this subset, ensuring the inclusion of variables that improve predictive precision. The algorithm initially starts with an empty collection of explanatory variables, represented by $S_0 = \emptyset$. Eventually, this collection grows by sequentially incorporating variables that optimise the information content pertinent to the EUA price. During every iteration of this process, the algorithm evaluates the informational value contributed by each of the candidate variables still available that have not yet been included in the current set S_t . The content of information is generally assessed through measures such as mutual information, information gain, or Information Imbalance, which serve to quantify the degree of enhanced explanatory power a variable adds when integrated with the existing set denoted as S_t . These metrics offer insight into the extent to which the

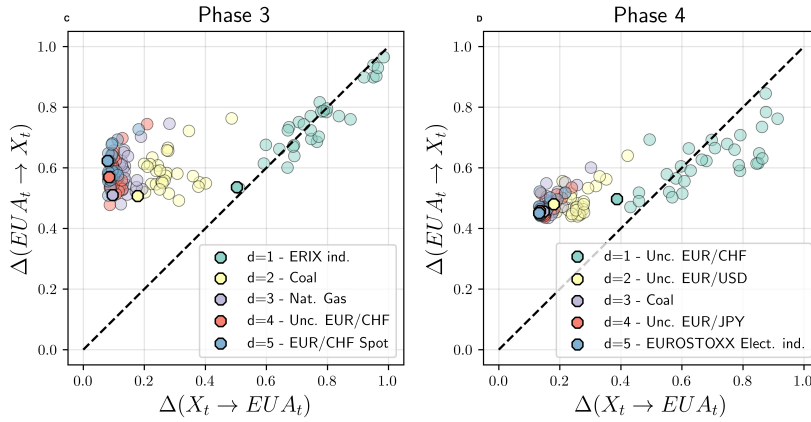


Figure 5.6: **Information Imbalance Greedy Optimization of Phase 3 and Phase 4 EUA Price Determinants.** This plot illustrates the greedy optimisation process used to select the variables that provide the most significant insights within the Information Imbalance framework. The approach aims to identify the most informative predictors for the European Union Allowance (EUA) price by iteratively adding the variables that maximise the imbalance.

variable contributes to a more thorough and comprehensive understanding of the dataset in question.

The variable that contributes the most significantly to the marginal increase in the content of the information is selected and subsequently integrated into the set, resulting in an updated set $S_{t+1} = S_t \cup \{X^*\}$. This procedure is repeated iteratively until each variable has been examined and ranked within the set. Identification of a sufficiently informative subset illustrated in 5.5 and 5.7 is carried out using the elbow rule. This rule is recognised at the point where the addition of further variables leads to only minor alterations in the gradient of the information content curve. To effectively implement the elbow rule, no predetermined threshold was relied on. Instead, the focus was on identifying the point at which supplementary variables provide only trivial improvements to the content of the information. Figure 5.5 shows that the most informative subset for Phase 3 ultimately

comprises the following explanatory variables: the ERIX index, Coal, and Natural Gas.

Extending the consideration beyond these three variables is inadvisable, as any additional gain in information content proves to be insignificant. This phenomenon is observable in figure 5.6, where the Information Imbalance is plotted against the number of variables (A, B) and can also be discerned from the Information Imbalance plane (C, D), with a noticeable denser cluster of points as more variables are incorporated into the set.

For Phase 4, the subset that provides the most valuable information is found to be different, with the notable exception of Coal. What is particularly noteworthy in this scenario is the incorporation of uncertainty indicators derived from the exchange rates evaluated in this analysis. Fluctuations in exchange rates can amplify uncertainties in the market. Investors often perceive currency volatility as a risk factor, and a spike in exchange rate uncertainty could increase investor risk aversion, possibly affecting their participation in the EU ETS market [130, 16]. Actions taken by central banks in response to exchange rate uncertainties and shifts in monetary policies have the potential to influence interest rates and broader economic conditions. These complex interconnections can, in turn, affect the regulatory framework and policy decisions regarding emission trading within the EUA [40].

5.4 Time-Scale Aggregation and Forecasting

5.4.1 Data Frequency Selection

In the process of performing imputation and aggregation via a Gaussian Process (GP)-based approach, four distinct datasets are generated, characterised

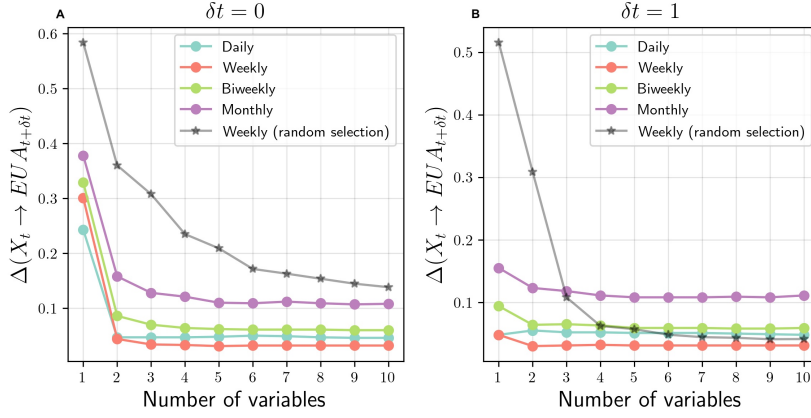


Figure 5.7: **Frequency identification through the Information Imbalance.** The figure illustrates the Information Imbalance, denoted by $\Delta(X_t \rightarrow EUA_t)$, highlighting how the transition from the predictor variable set X_t to the European Union Allowance (EUA) pricing structure changes as the number of variables in X_t increases and across different data frequencies.

by daily, weekly, biweekly, and monthly frequencies. Using these datasets as a foundation, the concept of Information Imbalance (II) is employed to discern the particular frequency at which the predictors exhibit the most substantial relevance to the EUA price. This process involves conducting a greedy and iterative stepwise selection of variables for each frequency as outlined in Section 5.3.

The objective of this selection process is to minimise the II, indicated as $\Delta(X_t \rightarrow EUA_{t+\delta t})$, from the ensemble of predictors at the time point t (X_t) to the EUA price at the subsequent time point $t + \delta t$ ($EUA_{t+\delta t}$). These calculations are executed under two scenarios: a ‘nowcasting’ scenario with $\delta t = 0$, and a ‘forecasting’ scenario with δt equating to 1. The value of δt is interpreted as one day, one week, two weeks, or one month, depending on the data set used.

The imbalances that arise from our analysis are illustrated in Figure 5.7 and 5.8. Through our findings, it becomes apparent that regardless of fre-

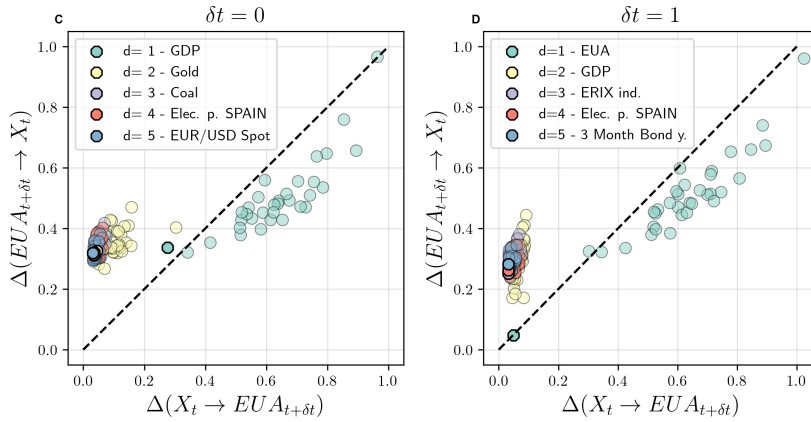


Figure 5.8: **Variable selection through the Information Imbalance.** The plot shows variable selection based on the Information Imbalance, highlighting how the predictor variables are ranked and selected in relation to the European Union Allowance (EUA) pricing structure.

quency and the two considered time lags, δt , the enhancement in the informational content of the predictor variables plateaus when more than three variables are included. This observation corresponds to the results shown earlier in Figures 5.7 and 5.8. Moreover, for both time lags, the data sampled at a weekly frequency have the most significant informational value to forecast EUA price movements.

This outcome implies that smoothing daily price fluctuations to a weekly frequency aids in enhancing the predictability of EUA prices. This is because such daily fluctuations are challenging to interpret with the features under consideration, whereas smoothing beyond a weekly time frame seems to obliterate substantial relationships, thereby diminishing the efficacy of predictions.

Due to the superior informational content associated with the weekly frequency, as well as considerations of simplicity and succinctness, the subsequent results in this section are exclusively presented for the weekly frequency.




Figure 5.9: **Information Imbalance plane for nowcasting EUA price.** The Information Imbalance simultaneously selection of each variable in relation to the target EUA. The analysis is based on weekly data. Variables with the highest informational value are placed at the top and highlighted in deep orange, as indicated in the legend.

5.4.2 Selection of Predictor Variables

As established in Section 5.4, the weekly time scale proved to be the most informative to predict the movements of EUA prices according to the time frequency selection carried out through Information Imbalance (II). This finding indicates that weekly data captures the essential dynamics of the carbon market more effectively than other temporal frequencies. Building on this result, we now apply this weekly time scale to examine both nowcasting ($\delta t = 0$) and forecasting ($\delta t = 1$) scenarios. The initial finding examined is the II computed for each individual predictor in Table 5.1, in relation to the target variable (EUA). This examination includes key predictors such as Euro-area Inflation and GDP, based on a dataset that encompasses both phase 3 and phase 4.

Figures 5.9 and 5.10 provide a comprehensive overview of the information content embodied by all predictive variables in relation to the target variable. Upon examining the observed imbalances, it becomes evident that, along

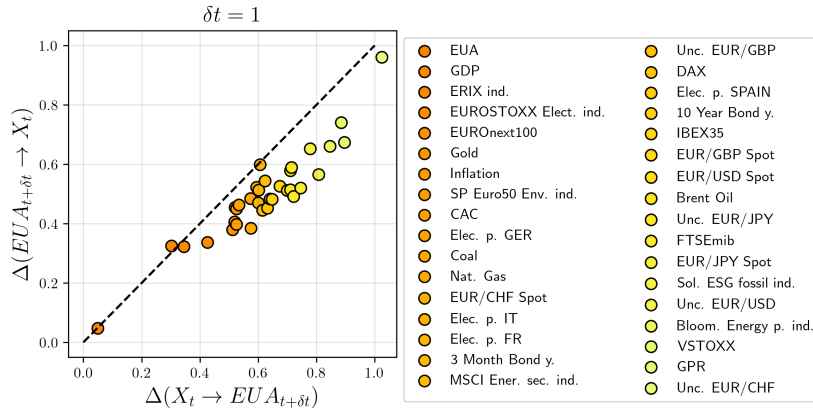


Figure 5.10: **Information Imbalance plane for forecasting EUA price.** The Information Imbalance 1-day ahead selection of each variable in relation to the target EUA. The analysis is based on weekly data. Variables with the highest informational value are placed at the top and highlighted in deep orange, as indicated in the legend.

with the Gross Domestic Product (GDP), other variables demonstrating substantial predictive capability include the ERIX index, the EUROSTOXX Electricity index, and the EUROnext100 index. Conversely, variables that are designed to quantify uncertainty, such as the Geopolitical Risk (GPR) index, and the VSTOXX index, exhibit markedly low predictive power when considered in this context.

In figure 5.10, the results derived from the computation of the Information Imbalance within a forecasting framework are further explained. In this context, the set of predictive variables is denoted as X_t , while the set of target variables is represented as EUA_{t+1} . In particular, the variable measured with a one-week delay, namely EUA_t , emerges as the most information-rich variable. Despite the presence of temporal lag, GDP still ranks among the most informative variables in the analysis.

Figures 5.7 and 5.8 illustrate the findings derived from the application of the iterative greedy algorithm for variable selection. The analysis reveals that

GDP stands out as the single most informative variable. This outcome corroborates the observations made when analysing Figures 5.9 5.10. Moreover, when considering a subset composed of three variables, the combination of GDP, Gold prices, and Coal prices provides the richest informational value.

Due to its status as a composite long-term metric, GDP serves as a reliable indicator for projecting the long-term dynamics of economic and financial variables [172]. In parallel, the European Union Allowance (EUA) fulfils a similar role within the scope of our research. Consequently, this variable can serve as a predictor of the future price trajectories of our target variable. Nevertheless, it is crucial to recognise that the model might have detected a possible distortion effect.

The path of GDP displays an almost consistent upward trajectory, reflecting the behaviour of EUA prices, which are initially low but progressively increase due to the impacts of COVID-19 and the subsequent energy crisis. In addition, gold is traditionally viewed as a secure investment option, which implies that investors frequently turn to it during times of economic turmoil or significant market fluctuations. Fluctuations in the price of gold often indicate changes in market sentiment, particularly during periods of economic instability or geopolitical uncertainty. Incorporating gold price metrics into a forecasting framework offers a robust tool to examine investor sentiment and predict future market trends [132]. In conclusion, the energy sector, with an emphasis on electricity production, is considerably dependent on coal pricing. Given that the European Union Emissions Trading Scheme (EU ETS) covers a range of industries, with a particular focus on the power generation sector where coal plays a pivotal role, fluctuations in coal prices emerge as critical indicators reflecting broader energy market dynamics. As a result, these variations can significantly influence both the demand and the price of

EU ETS allowances [125].

Figure 5.8 further illustrates the results of applying the greedy methodology when accounting for a time delay. In contrast to the prior nowcasting scenario, our approach discerns quite distinct informative subsets in this situation. Specifically, the currently identified optimal subset of three variables includes the lagged value of the target variable, which is GDP, in conjunction with the 3-month bond yield for the Euro-area. Although the link between the 3-month bond yield and the EUA price may not be immediately apparent, variations in short-term interest rates, as indicated by the 3-month bond yield, can indirectly affect economic circumstances and investor decisions. These effects have the potential to alter both the demand for and the pricing of EU ETS allowances, as noted in Chevallier’s study [33].

5.4.3 Prediction Performances

In this section, an in-depth verification is provided of how the variable selection approach, which is both information-driven and model-free, achieves accurate prediction results. A Gaussian Process (GP) regression model, as described in Section 3.3, is used for the baseline comparison. This GP model has previously been used for tasks such as imputation and aggregation, with further details in Appendix A.3. Analysis of Information Imbalance indicates that the additional predictive power contributed by the set of all predictors X_t is quite limited compared to using just three key variables that are carefully chosen. Therefore, in the comparison, the performance of a GP model constructed with these 3 most informative variables is examined against one with 3 randomly selected variables and, finally, against models using the complete set of 33 variables for the Nowcasting framework and 34 variables for the Forecasting framework, respectively.

Figure 5.11 illustrates the results of an investigation using Gaussian Process (GP) models to perform both nowcasting, denoted by $\sigma t = 0$, and forecasting, represented by $\sigma t = 1$. The study evaluates three distinct models: the first employs the three most informative variables determined by the Information Imbalance (II); the second incorporates all accessible variables; and the third uses three variables selected at random. The results shown are averaged across ten replicative trials. Within the figure, scatter plots depicting the residuals and the predicted European Union Allowance (EUA) time series are juxtaposed with the actual data for the contexts of nowcasting (on the left) and forecasting (on the right). A critical observation highlighted is that the model predicated on the three most informative variables delivers predictions of superior reliability and robustness compared to the model that encompasses all variables. This observation corroborates the pattern discerned in Figure 5.11, indicating that the incorporation of additional variables does not contribute to enhanced performance, but rather tends to introduce noise that undermines predictive precision. In conclusion, focussing on the most informative variables significantly improves prediction performance for both nowcasting and forecasting scenarios, whereas the inclusion of extraneous or randomly chosen variables generally diminishes predictive accuracy.

The same effect is summarised in Table 5.3, which reports the mean squared error (MSE) computed on the cross-validation sets along with their level of uncertainty. Additionally, our non-parametric methodology demonstrates exceptional prediction capabilities, as detailed in 5.3, with notably lower MSE values, by several orders of magnitude, than the ones reported in [100]. Comparing the prediction performance between nowcasting ($\sigma t = 0$) and forecasting ($\sigma t = 1$) using the same model reveals significant differ-

	$\delta t = 0$	$\delta t = 1$
<i>Inf. Imb.</i>	$0.8 \pm 0.3 \cdot 10^{-3}$	$0.1 \pm 0.11 \cdot 10^{-3}$
<i>All</i>	$1.1 \pm 0.6 \cdot 10^{-3}$	$0.6 \pm 0.4 \cdot 10^{-3}$
<i>Rand.</i>	$79.5 \pm 21.2 \cdot 10^{-3}$	$34.8 \pm 6.8 \cdot 10^{-3}$

Table 5.3: **Prediction performance (MSE)**. Mean Squared Error (MSE) for nowcasting ($\delta t = 0$) and forecasting ($\delta t = 1$) GP models built using all predictors (*All*), a set of 3 random predictors (*Rand.*) and the set of 3 variables selected via the iterative greedy optimisation of the Information Imbalance (*Inf. Imb.*). The GP model built on the 3 variables selected via Information Imbalance performs best.

ences. For the *Inf. Imb.* model, the MSE for nowcasting is $0.8 \pm 0.3 \cdot 10^{-3}$, while for forecasting it drops dramatically to $0.1 \pm 0.11 \cdot 10^{-3}$. This substantial reduction in MSE indicates that the *Inf. Imb.* model performs better in forecasting compared to nowcasting. In contrast, the *All* predictor model shows an MSE of $1.1 \pm 0.6 \cdot 10^{-3}$ for nowcasting and $0.6 \pm 0.4 \cdot 10^{-3}$ for forecasting. Although the *All* predictors model also performs better in forecasting, the improvement is less pronounced. However, the random predictor model shows a much higher MSE in both nowcasting $79.5 \pm 21.2 \cdot 10^{-3}$ and forecasting $34.8 \pm 6.8 \cdot 10^{-3}$, with less improvement observed between nowcasting and forecasting. Overall, the comparison underscores that while both the *Inf. Imb.* and *All* predictor models achieve better forecasting performance than nowcasting, and the former exhibits a particularly marked improvement, highlighting its effectiveness in longer-term predictions.

To quantitatively compare the residuals of the two forecasts (*All* and *Inf. Imb.*), we performed two statistical analyses: the Kullback-Leibler (KL) Divergence [119] and the Kruskal-Wallis (KW) test [118]. For $\delta t = 0$ (nowcasting), the KL Divergence, which measures the difference between two probability distributions, yielded a value of 0.2389, indicating that the distributions of the residuals are somewhat different. This suggests that the two

predictions produce residuals with distinct shapes or spreads. However, the Kruskal-Wallis test, which compares the medians of the residuals, resulted in a p-value of 0.0628, indicating that there is no statistically significant difference between the medians of the residuals. These results imply that while the central tendencies of the residuals are similar, the distributions differ in terms of their overall shape or spread, potentially due to how the methods handle specific cases or outliers. For $\delta t = 1$ (forecasting), the KL Divergence value of 0.7518 indicates a much larger difference between the residual distributions, and the Kruskal-Wallis test yielded a p-value of 9.7957^{-06} , indicating a statistically significant difference in the medians of the residuals. This suggests that the prediction performance of the two methods differs substantially, with significant differences in both the shape and the central tendencies of the residuals. This analysis provides a more nuanced understanding of the differences between the methods, complementing the visual representation in Figure 5.11.

5.5 Discussion

This chapter introduces a novel non-parametric approach to identify carbon price drivers in the EU ETS and perform mixed-frequency nowcasting and forecasting. The Information Imbalance (II) methodology eliminates the need for assumptions about functional forms, offering a flexible, data-driven alternative to traditional parametric models.

Previous studies have analysed carbon price drivers using a variety of parametric and semi-parametric methods. For example, [128, 3] used linear regression models to identify the influence of energy prices and weather conditions on carbon prices. [114] extended this work by examining causal

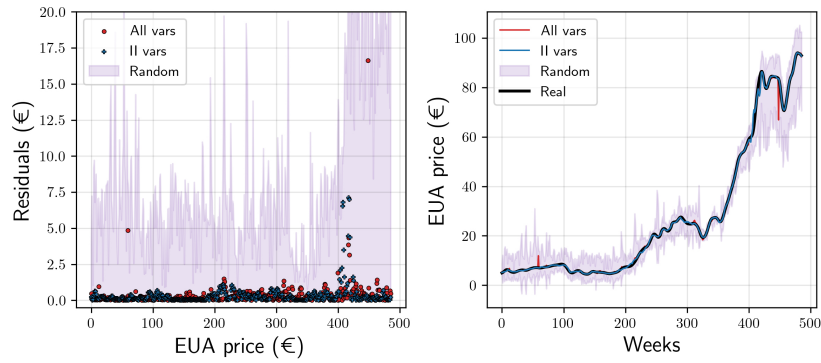
relationships between energy markets and carbon prices, while [102] applied cointegration techniques to study the long-term relationships between carbon prices and their drivers. Although these studies have provided valuable insights, they often rely on assumptions about the functional form of relationships between variables, which can limit their ability to capture complex, non-linear interactions. In contrast, our non-parametric Information Imbalance (II) approach operates without such constraints, allowing for a more flexible and accurate analysis of carbon price dynamics. This is particularly important in the context of the EU ETS, where carbon prices are influenced by a wide range of factors, including energy markets, macroeconomic conditions, and financial variables. By minimising the risk of model misspecification, our approach offers a potentially more reliable framework for examining carbon price drivers and their changes over time.

The II approach provides regulators with a powerful tool to monitor and understand the functioning of the EU ETS. By identifying the most informative drivers of carbon prices, such as energy indices, financial variables, and macroeconomic indicators, this methodology offers actionable insights into market dynamics. For example, the shift toward financially orientated variables in Phase 4 suggests that regulators should pay more attention to global financial markets, as these play a significant role in the formation of carbon prices during turbulent periods. Furthermore, the superior forecasting performance obtained by combining Gaussian Processes (GP) and the II metric enables regulators to anticipate price trends and assess the potential impact of policy changes or external shocks, such as energy crises or economic disruptions. This capability is critical for designing effective market stability mechanisms and ensuring the long-term efficiency of the EU ETS.

The observed differences between Phase 3 and Phase 4 of the EU ETS

reveal important trends in the evolution of carbon markets. In Phase 3, fundamental factors such as energy indices and macroeconomic variables were the primary drivers of carbon prices. In contrast, Phase 4 saw a significant shift towards financial variables, such as exchange rates and coal prices. This change likely reflects the impact of external shocks, including the COVID-19 pandemic and the energy crisis, which have increased the integration of the EU ETS with wider capital markets. The growing influence of financial variables suggests that carbon markets are no longer driven solely by energy and policy factors but are increasingly intertwined with global financial systems. This has implications for both market participants and policymakers, as it highlights the need to consider broader economic and financial contexts when analysing carbon price dynamics.

$\delta t = 0$



$\delta t = 1$

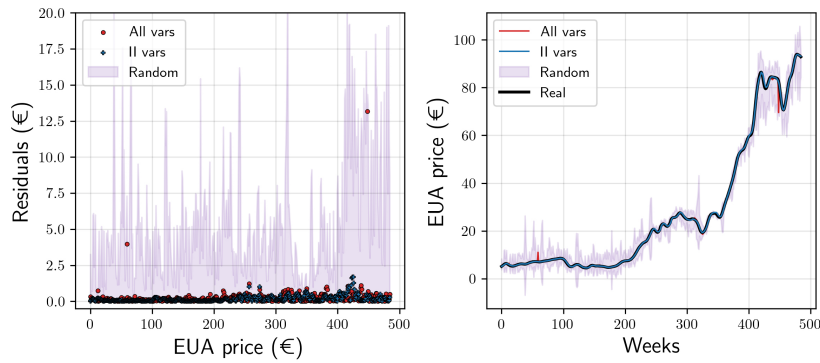


Figure 5.11: Performance of a GP model built on: the 3 most informative variables; all variables; 3 randomly selected predictors (average over 10 replications). In the contexts of nowcasting (on the left) and forecasting (on the right), a scatter plot is presented to illustrate the relationship between the predicted and actual EUA (European Union Allowance) time series. The figure includes both the residuals and the observed time series for a comprehensive comparison.

Chapter 6

Empirical Results III: Volatility Causality and Interaction with Financial Variables in the EU ETS

6.1 Data

The daily dataset analysed in this study covers a variety of market categories, including environmental markets, commodities, exchange rates, energy indices, and country-specific indices. The dataset of financial returns extends from January 2013 to April 2024, with a total of 2902 observations. Table 6.1 displays the dataset, highlighting variables related to uncertainty indicators and commodity prices, particularly the VSTOXX volatility index, as well as futures for ICE Brent oil and LME Copper.

Previous research has demonstrated the importance of emissions trading schemes in influencing market dynamics and price formation [113]. Also,

ID	Category	Variables	Abbreviations	Database
0	T	EUA (ICEEINDEX)	EUA	Bloomberg®
1	UNC	GPR	GPR	GPR website
2	UNC	VSTOXX (V2X)	VSTOXX	Bloomberg®
3	UNC	Uncertainty EUR/USD (CAFZUEU)	UncEURUSD	Bloomberg®
4	UNC	Uncertainty EUR/JPY (CAFZUEJP)	UncEURJPY	Bloomberg®
5	UNC	Uncertainty EUR/GBP (CAFZUEGB)	UncEURGBP	Bloomberg®
6	UNC	Uncertainty EUR/CHF (CAFZUECH)	UncEURCHF	Bloomberg®
7	COM	ICE Dutch TTF Natural Gas (TTF0NXHR)	NatGas	Bloomberg®
8	COM	Electricity Prices Spain (OMLPAHD)	ElecES	Bloomberg®
9	COM	Electricity Prices Germany (EXAPBDHD)	ElecDE	Bloomberg®
10	COM	Electricity Prices France (PWNXFRAV)	ElecFR	Bloomberg®
11	COM	ICE Coal Rotterdam futures (TMA Comdty)	CoalFut	Bloomberg®
12	COM	LME Copper futures (LMCADS03 Comdty)	CuFut	Bloomberg®
13	COM	ICE Brent oil futures (CO1 Comdty)	Brent	Bloomberg®
14	COM	Silver (XAG Comdty)	AgFut	Bloomberg®
15	COM	Gold (GCZ3 Comdty)	Gold	Bloomberg®
16	ER	EUR/USD spot (EUR/USD)	EURUSD	Eikon Refinitiv®
17	ER	EUR/JPY spot (EUR/JPY)	EURJPY	Eikon Refinitiv®
18	ER	EUR/GBP spot (EUR/GBP)	EURGBP	Eikon Refinitiv®
19	ER	EUR/CHF spot (EUR/CHF)	EURCHF	Eikon Refinitiv®
20	ENR	WilderHill New Energy Global Innovation index (Nex)	WHNewEnergy	Bloomberg®
21	ENR	Bloomberg Energy TR index (BCOMENTR)	BbgEnergy	Bloomberg®
22	ENR	Solactive CEA Future index (SOLAFCEA)	SolCEA	Bloomberg®
23	ENR	EUROSTOXX Electricity index (SXEELC)	ESTXElect	Bloomberg®
24	ENR	Solactive ESG Fossil EU 50 index	SEF EU50	Bloomberg®
25	ENR	Low Carbon 100 index (LC100)	LC100EU	Bloomberg®
26	ENR	MSCI Europe Energy Sector index (MXEU0EN)	MSCIEng	Bloomberg®
27	ENR	ERIX index	ERIX	Bloomberg®
28	CTRY	Euronext100 (N100)	Euronext100	Bloomberg®
29	CTRY	IBEX35	IBEX35	Eikon Refinitiv®
30	CTRY	DAX	DAX	Eikon Refinitiv®
31	CTRY	CAC	CAC	Eikon Refinitiv®
32	CTRY	FTSEmib	FTSEmib	Eikon Refinitiv®
33	MACRO	Euro-area 10-year bond yield (BN10)	Bund10y	Bloomberg®
34	MACRO	Euro-area 3-month bond yield (BN03)	Bond3m	Bloomberg®

T: Target; UNC: Uncertainty variables; COM: Commodity-related variables;
ER: Exchange rates; ENR: Energy-related indexes/variables; CTRY: Country indexes; MACRO: Macro-economic variables.

Table 6.1: **Variable descriptions and data sources.** The variables used in the analysis with their full names, abbreviations, codes, categories, and data sources. All data is collected in euros (EUR).

uncertainty indices such as the GPR and VSTOXX have shown strong correlations with market volatility and investment decisions [12]. This dataset also includes spot rates for major currency pairs such as EUR/USD, EUR/JPY, EUR/GBP, EUR/CHF, as well as energy indices, such as the WilderHill New Energy Global Innovation Index. Standardised returns are used in the subsequent analyses, as the variances of the 35 original dataset variables vary by up to three orders of magnitude. This standardisation ensures comparability and reduces the influence of larger-scale variables.

The dataset is sourced from Bloomberg and Eikon Refinitiv.

Tables 6.1 and 5.1 share similar variables organised into the same categories, but differ in selection and presentation. The former introduces additional variables, such as LME Copper futures, Silver, and the WilderHill

New Energy Global Innovation index. In contrast, the latter includes specific macroeconomic variables like inflation in the Euro area and GDP, which are omitted in Table 6.1 to avoid the issue of mixed frequency data and maintain a focus on empirical results. Both tables cover exchange rates, country indexes and uncertainty variables, although Table 6.1 replaces certain energy-related indexes (e.g. Bloomberg Energy price return index) with alternatives such as the Bloomberg Energy TR index and the Low Carbon 100 index. In general, 6.1 builds on 5.1 by incorporating new variables and abbreviations while excluding GDP and inflation to focus on the empirical analysis.

6.1.1 Computation and Descriptive Analysis of Financial Returns

The financial returns at time t , R_t are calculated as:

$$R_t = \frac{P_{t+1} - P_t}{P_t}, \quad (6.1)$$

where P_t is the price of the asset at time t . Financial returns are often used as an indicator of volatility, as they capture price fluctuations, and larger returns generally reflect higher volatility. This relationship is fundamental to models such as the ARCH model, where historical returns are used to estimate current volatility [78].

In this analysis, we standardise financial returns to ensure a fair and unbiased comparison within the DII framework, which operates in a multidimensional space. Although converting prices to returns provides some level of normalisation, variations in return volatility across assets can still lead to distortions in distance-based calculations. In a multidimensional setting, features with significantly higher variance can dominate the distance metric,




Figure 6.1: **EUA futures financial returns.** European Union Allowance (EUA) futures returns are analysed alongside significant events such as the EU Banking Crisis (2012/11 - 2014/09), Brexit (2016/01 - 2020/12), and the Energy Crisis (2021/03 - 2024/04). Both the returns and their 30-day moving averages are plotted, with key EU ETS phases marked. Periods of increased market uncertainty, reflected by volatility spikes, align with major European crises, which are shaded in light blue for context.

skewing the model and potentially leading to incorrect feature selection. By standardising the returns, we mitigate this risk, ensuring that all financial variables contribute equally to the distance measurement. This process enhances the robustness of our approach and prevents distortions caused by differences in variance from influencing our results.

Figure 6.1 illustrates EUA futures returns in conjunction with key economic and geopolitical events that could have impacted them. For example, during the EU Banking Crisis (2009–2014), financial instability may have contributed to reduced liquidity and market confidence, potentially leading firms to reassess their emissions and compliance costs.

Table 6.2 presents summary statistics for the 35 time series of returns

ID	Variables	Mean	STD	Min	Max	25%	50%	75%	Skewness	Kurtosis
0	EUA	0.0013	0.0316	-0.3511	-0.0143	0.0010	0.0177	0.2703	-0.2903	10.5662
1	GPR	0.0999	0.5408	-0.9500	-0.2355	-0.0018	0.3031	9.4341	3.5417	37.5951
2	VSTOXX	0.0024	0.0725	-0.3526	-0.0398	-0.0049	0.0332	0.6256	1.3944	7.1025
3	UncEURUSD	0.0001	0.0207	-0.1033	-0.0078	0.0000	0.0054	0.1850	0.9916	8.9573
4	UncEURJPY	0.0002	0.0291	-0.3153	-0.0076	0.0000	0.0064	0.4380	1.6756	44.7733
5	UncEURGBP	0.0003	0.0294	-0.4078	-0.0081	0.0000	0.0065	0.6793	5.7195	145.6690
6	UncEURCHF	0.0011	0.0638	-0.3430	-0.0065	0.0000	0.0057	3.0505	38.3296	1814.1401
7	NatGas	0.0009	0.0437	-0.2970	-0.0155	-0.0004	0.0147	0.5110	1.4312	16.1099
8	ElecES	0.1398	2.1517	-0.9918	-0.0755	-0.0028	0.0723	70.6610	26.3826	776.1512
9	ElecDE	0.0674	3.2588	-72.7895	-0.1358	-0.0198	0.1348	150.2692	29.4670	1648.0117
10	ElecFR	0.0467	0.7004	-0.8609	-0.1047	-0.0126	0.0970	32.3447	35.1401	1574.0379
11	CoalFut	0.0005	0.0259	-0.3103	-0.0093	0.0001	0.0097	0.3746	0.9355	36.1277
12	CuFut	0.0001	0.0120	-0.0776	-0.0065	0.0000	0.0067	0.0712	-0.0431	2.2281
13	Brent	0.0334	0.6771	-0.8609	-0.0386	-0.0022	0.0249	32.3447	38.8817	1805.5031
14	AgFut	0.0001	0.0176	-0.1161	-0.0080	0.0000	0.0081	0.0930	-0.2531	5.5203
15	Gold	0.0001	0.0092	-0.0907	-0.0046	0.0003	0.0050	0.0509	-0.4578	6.2462
16	EURUSD	-0.0001	0.0050	-0.0278	-0.0030	0.0000	0.0026	0.0318	0.0964	2.9530
17	EURJPY	0.0001	0.0060	-0.0510	-0.0030	0.0001	0.0034	0.0445	-0.1842	5.4367
18	EURGBP	0.0001	0.0057	-0.1046	-0.0026	-0.0001	0.0025	0.1164	1.2510	101.7974
19	EURCHF	-0.0001	0.0045	-0.1689	-0.0016	-0.0001	0.0016	0.0290	-18.0536	688.2632
20	WHNewEnergy	0.0004	0.0134	-0.1023	-0.0063	0.0007	0.0071	0.0948	-0.2600	5.5705
21	BbgEnergy	-0.0001	0.0189	-0.1353	-0.0095	0.0001	0.0097	0.1055	-0.3188	4.7386
22	SolCEA	0.0012	0.0314	-0.3508	-0.0141	0.0005	0.0174	0.2690	-0.3441	10.8382
23	ESTXElect	0.0003	0.0115	-0.1594	-0.0056	0.0003	0.0066	0.0593	-1.2169	15.6282
24	SEF EU50	0.0002	0.0114	-0.1321	-0.0049	0.0003	0.0059	0.0881	-0.7391	11.3356
25	LC100EU	0.0003	0.0010	-0.1100	-0.0041	0.0006	0.0052	0.0768	-0.7697	9.9159
26	MSCIEng	0.0002	0.0164	-0.1808	-0.0071	0.0004	0.0076	0.1938	-0.1182	17.7366
27	ERIX	0.0007	0.0163	-0.1217	-0.0076	0.0010	0.0092	0.1055	-0.1282	3.8237
28	Euronext100	0.0003	0.0108	-0.1197	-0.0045	0.0006	0.0058	0.0818	-0.7453	10.1859
29	IBEX35	0.0002	0.0123	-0.1406	-0.0060	0.0004	0.0066	0.0857	-0.9574	12.7946
30	DAX	0.0004	0.0118	-0.1224	-0.0047	0.0006	0.0063	0.1098	-0.3781	9.6570
31	CAC	0.0003	0.0117	-0.1228	-0.0049	0.0006	0.0061	0.0839	-0.6053	9.6292
32	FTSEmib	0.0003	0.0140	-0.1692	-0.0062	0.0007	0.0077	0.0892	-1.1029	13.0261
33	Bund10y	-0.0080	0.7693	-32.6667	-0.0348	0.0000	0.0353	13.0000	-25.8877	1165.6719
34	Bond3m	0.0002	0.7113	-27.0000	-0.0125	0.0000	0.0155	8.2857	-18.4459	742.9117

Table 6.2: **Descriptive statistics of financial returns.** Some key descriptive statistics for 35 financial time series, including the mean, standard deviation (STD), minimum (Min), maximum (Max), skewness, kurtosis, and percentile values (25%, 50%, and 75%).

included in our analysis. These statistics highlight key characteristics in various markets. For example, EUA futures exhibit a slightly positive mean return (0.0013) along with relatively high volatility (0.0316) and negative skewness (-0.2903), suggesting a tendency for occasional extreme negative returns.

Energy commodities, including natural gas and Brent oil, show significant volatility, likely influenced by geopolitical tensions and supply constraints during the ongoing energy crisis. The Spanish electricity market shows an even higher dispersion (STD 0.38361), reflecting the increased uncertainty

in the regional energy markets. In contrast, financial indices such as EUROSTOXX and DAX display lower volatility. These differences highlight the diverse return behaviours across asset classes. The elevated volatility in energy commodities, the extreme kurtosis observed in currency markets, and the asymmetric distributions in financial indices underscore the relevance of these characteristics for risk assessment and portfolio management. These patterns also emphasise the sensitivity of markets to geopolitical shocks, which present both challenges and opportunities for investors [14]. Furthermore, the increase in volatility in currency pairs, especially EUR/USD and EUR/GBP, following Brexit, points to the economic instability arising from new trade agreements and regulatory changes [25]. Such observations are important for understanding market dynamics and formulating investment strategies.

In the analysis, the causal sufficiency hypothesis is adopted, assuming that all economic and geopolitical influences on the EUA returns are directly captured by the variables listed in Table 6.2 or indirectly mediated by one or more variables in the dataset.

6.2 Empirical Analysis

The use of the VAR(1) model in our analysis is motivated by its ability to capture linear causal relationships between variables in the Granger sense, where a variable is said to Granger-cause another if its past values provide statistically significant information about the future values of the target variable. The choice of 1-day time lag was determined by the Akaike Information Criterion (AIC), which identified it as the optimal lag length to model temporal dependencies in our data. This ensures that the VAR(1) model

is both parsimonious and effective for our purposes. We employ VAR(1) as a benchmark to compare with our proposed DII (Differentiable Information Imbalance) model, which is designed to capture non-linear relationships through distance weighting. By using the same lag order as the VAR(1), we maintain a consistent framework for comparison, allowing us to evaluate the advantages of the DII model in identifying non-linear causal relationships that may be overlooked by the linear VAR(1) approach. This comparative analysis highlights the importance of considering both linear and non-linear methodologies when investigating causal dependencies, particularly in complex systems where non-linear interactions may play a significant role.

6.2.1 Granger’s F-statistic vs. Imbalance Gain: A Comparative Analysis

Figure 6.2 presents a comparison between the Information Imbalance Gains (IGs) and Granger’s F-statistics, calculated for all predictors in relation to the European Union Allowance (EUA) returns. In the first panel, the IG metric identifies IBEX35 and CoalFut as having the strongest causal influence. Notably, the IG and Granger’s F-statistic both attribute similar significance to these variables. However, a more detailed analysis reveals discrepancies in the ranking of causal effects between the two methods. Among the next eight most influential variables, only Gold and CuFut are identified as causally relevant by both the linear and non-linear approaches.

In particular, in the scatter plot of Fig. 6.2, several variables that display a non-zero Imbalance Gain are also assigned a negligible F-statistic. Following the findings of the first example in Sec. 3.2.4, it is suggested that the causal impact of these variables might be underestimated by multivariate GC due to their non-linear relationships with respect to EUA. It is important to note




Figure 6.2: **The Imbalance Gain as a non-linear alternative to the GC F-statistic.** The outcomes derived from the multivariate Granger causality (GC) and DII methods, with EUA as the target variable, are presented. The left panel shows the sorted F-statistics from the GC analysis, calculated using VAR models of order 1. The central panel illustrates the sorted Imbalance Gain (IG) measures. Bold highlights indicate the top four variables among the first ten identified by both methods as having a causal influence on EUA. The right panel provides a scatter plot that correlates the F-statistics with the Imbalance Gain values.

that a key difference between the two metrics lies in the types of relationships they can detect. The F-statistic specifically captures linear relationships between the target variable and each predictor, which are constrained by the linear structure of the VAR model. In contrast, the IG metric captures non-linear, model-free causal contributions. Conversely, fewer variables show null IG estimates while having a non-zero F-statistic. As demonstrated in the second example of Sec. 3.2.4, this situation can arise when a common driver influences both EUA and the candidate causal variable identified by GC, where similar non-linear relationships exist that linear methods cannot

directly detect. If non-linear relationships are significant in the complex dynamics of financial returns, particularly in the factors influencing EUA's evolution, the Imbalance Gain could serve as a valuable tool for avoiding both false negatives and false positives produced by traditional non-linear techniques.

6.2.2 Estimating Causal Weights: VAR and Information Imbalance Approaches




Figure 6.3: **Beyond VAR causal weight estimation.** The first panel ranks the estimated weights from a VAR(1) model in descending order. The second panel shows the causal weights, arranged in descending order according to the DII measure. Bold highlights emphasize the top 5 weights among the initial 10 that both methods identify as causal for EUA. The third panel features a scatter plot that correlates the VAR weights with those from the DII measure.

The VAR(1) model is adopted here as a linear benchmark to assess causality between EUA prices and the selected financial variables. Its role is not

to provide definitive causal inference, but to offer a standard point of comparison with the non-linear DII framework. The order of the VAR model, specifically the use of a single lag, was determined using the Akaike Information Criterion (AIC), which identified lag 1 as the optimal choice based on model fit and parsimony. This choice ensures coherence with standard practices in multivariate time series analysis, while keeping the structure consistent for fair comparison. In this context, the VAR(1) model enables us to contrast the behaviour of linear weights with the data-driven, potentially non-linear weights derived from the DII methodology, helping to highlight the added value of the latter in capturing complex dependencies.

In figure 6.3, a comparison of the weight estimates obtained from the VAR(1) model and the DII model is presented. The variables highlighted in bold represent the top five weights, chosen from the ten largest weights, which are common to both estimation methods. This highlights a certain level of consistency between the two approaches, as reflected in the weight estimates. As demonstrated in the applications to synthetic data (Sec. 3.2.4), while the optimal weights can offer valuable insights, they do not always accurately represent the true causal impact of the variables on the target. This is because non-zero weights may not necessarily reflect significant causal contributions, particularly when associated with large uncertainties that render them practically irrelevant. For this reason, the results obtained from the Imbalance Gain are considered to be more reliable. The optimal DII weights are presented primarily for the purpose of comparison with their linear counterparts from the VAR model.

6.3 Discussion

The comparative analysis presented in this chapter highlights the strengths and limitations of both the VAR(1) model and the DII (Differentiable Information Imbalance) approach in identifying causal relationships, particularly in the context of the European Union Allowance (EUA) returns. The primary goal of this analysis was to evaluate the performance of the DII model, which captures non-linear relationships, against the traditional VAR(1) model, which is limited to linear dependencies, also presented in Section 3.2.4. By using the same lag order, 1-day time lag, for both models, we ensured a consistent and fair comparison, allowing us to isolate the effects of non-linearity on causal inference.

The results reveal several key insights. First, the Imbalance Gain (IG) metric, derived from the DII model, demonstrates a clear advantage in identifying non-linear causal relationships that are overlooked by the linear VAR(1) approach. For instance, while both methods identify IBEX35 and CoalFut as significant predictors of EUA returns, the IG metric uncovers additional variables with non-linear causal influences that are missed by the Granger causality F-statistic. This is evident in the scatter plot of Figure 6.2, where several variables exhibit non-zero IG values but negligible F-statistics. These findings suggest that the linear constraints of the VAR model may lead to false negatives in causal inference, particularly in complex systems where non-linear interactions are prevalent.

Conversely, the VAR(1) model occasionally identifies variables with non-zero F-statistics but negligible IG values. This discrepancy can arise when a common driver influences both the target variable and the predictor, creating spurious linear relationships that the VAR model detects but the DII model does not. This underscores the importance of using non-linear methods like

DII to avoid false positives that may arise from linear assumptions.

Comparison of causal weight estimates further reinforces the advantages of the DII approach. Although both methods show some consistency in identifying the main predictors (e.g., IBEX35 and CoalFut), the DII weights are more robust in distinguishing true causal contributions from spurious associations. This is particularly important in financial systems, where non-linear dynamics are often significant, and traditional linear methods may fail to capture the true underlying relationships.

The findings of this study have important implications for policy makers, market participants, and researchers. For policy makers, the ability to accurately identify causal relationships in financial markets, particularly in the context of carbon markets like the EUA, is crucial for designing effective regulatory frameworks. The DII model's ability to capture non-linear dependencies can provide deeper insights into the factors driving EUA price fluctuations, such as the influence of energy commodities (e.g. coal futures) and equity indices (e.g. IBEX35).

This is particularly relevant in the context of stress testing and scenario analysis, where accurate identification of causal drivers is essential for assessing the potential impact of adverse events.

To conclude, the DII model, with its ability to capture non-linear dependencies, provides a more comprehensive framework for causal inference compared to the linear VAR(1) model. The Imbalance Gain metric, in particular, serves as a powerful tool for identifying non-linear causal relationships, reducing both false negatives and false positives that may arise from traditional linear methods. These findings highlight the importance of incorporating non-linear methodologies in causal analysis, especially in complex systems like financial markets, where linear models may fall short. Future

work could explore the application of the DII model to other domains with complex, non-linear interactions, further validating its utility as a robust tool for causal inference.

Chapter 7

Conclusions

The European Union Emissions Trading System (EU ETS) is a cornerstone of Europe's climate policy, implementing a cap-and-trade mechanism to regulate greenhouse gas emissions. The system's current phase (2021-2030) continues to tighten emission limits while expanding its coverage to additional sectors, aligning with the EU's climate targets. Our bibliometric review of 367 papers spanning from 2004 to 2024, sourced from the Scopus database, provides a comprehensive analysis of the academic research surrounding the EU ETS. Advanced tools such as citation analysis, co-authorship networks, and keyword analysis are used to examine the evolving trends, collaborative networks, and shifts in research topics over time.

The descriptive analysis reveals a consistent annual growth rate of 12.99%, with notable spikes corresponding to significant policy changes and regulatory updates. Countries such as China, Germany, and France have emerged as key contributors to this field, indicating the global importance of emissions trading systems. Germany, in particular, leads in total citations, while the Netherlands stands out with the highest average citation rates, underscoring the significant influence of research from these countries. Leading academic

journals like *Energy Policy* and *Energy Economics* have been essential in disseminating pivotal findings on the subject.

The review highlights an increasing focus on the complexities of carbon pricing, market volatility, and the financial dimensions of carbon trading. Keywords such as carbon price and carbon market have gained prominence, signalling a heightened interest in pricing strategies, market behaviours, and the broader economic ramifications of emissions trading systems. Network and thematic map analyses further reveal central concepts as carbon leakage, European Union Allowance (EUA), and ARIMA (AutoRegressive Integrated Moving Average), which serve as critical intersections in the literature, linking various areas of research. These key themes reflect a growing emphasis on understanding the economic and financial dimensions of cap-and-trade systems, especially in light of their impacts on market stability and financial outcomes.

However, several research gaps remain. Notably, there is a lack of exploration into non-parametric methods, particularly concerning feature selection and forecasting in emissions trading models. In addition, there is a shortage of research on the application of mixed-frequency data, which could provide a more robust analysis by incorporating a broader range of economic factors. Another important gap is the lack of examination into how exchange rates influence the EU ETS market dynamics. Moreover, macroeconomic variables such as inflation, interest rates and GDP growth have not been fully integrated into the analysis of EU ETS price fluctuations, highlighting an important area for future exploration.

In light of these gaps, this study proposes an innovative methodology for forecasting EU ETS prices, leveraging a completely non-parametric approach. The focus is on identifying exogenous variables that significantly

influence the price of the EU ETS market. Our data set, which spans from January 2014 to April 2023, includes 33 exogenous variables categorised into uncertainty-related variables, commodities, exchange rates, energy indexes, country-specific indexes, and macroeconomic variables. The use of the Information Imbalance technique stands out for its ability to process variables in their original form without relying on any assumptions about the underlying model or the relationships between time series. This provides a distinct advantage over conventional parametric models.

Our empirical analysis reveals that the most informative variables influencing EUA prices vary significantly between Phase 3 and Phase 4 of the EU ETS. In Phase 3, energy indices, commodities, and macroeconomic variables, such as the ERIX index and EUROSTOXX Electricity index, played a pivotal role. However, in Phase 4, following the disruptions caused by the COVID-19 pandemic and the ensuing energy crisis, the focus shifted towards financially-oriented variables, with uncertainty in the EUR/CHF exchange rate, coal prices, and the EUR/USD spot rate becoming more informative. The iterative greedy selection process further confirmed these findings, demonstrating the shifting dynamics of the market during these periods. Besides, this study introduces an innovative non-parametric methodology to enhance both now-casting and forecasting performance. Longer-term macroeconomic variables, such as Euro-area GDP and inflation, are integrated, with data captured at quarterly and monthly frequencies, respectively. This addition allows for a better capture of the movements of the economic cycle that can influence EUA prices over an extended period. By using Gaussian processes to aggregate variables and ensure consistency across temporal scales, the results show that the weekly scale is the most informative for forecasting EU ETS prices. Importantly, it is demonstrated that focussing on high-information

variables significantly improves both nowcasting and forecasting accuracy.

Non-linear causal relationships are also explored using a novel technique, the Differentiable Information Imbalance (DII), to identify key drivers of EU ETS market movements. The DII approach is a non-parametric, model-free measure designed to detect non-linear relationships, a crucial advancement over traditional linear models like VAR (Vector AutoRegression). Our empirical analysis of financial returns from January 2013 to April 2024 illustrates the effectiveness of DII in capturing causal relationships that VAR models often fail to identify. Notably, both the DII and VAR models identified significant causal influences from IBEX35 and coal futures on EUA returns, with the IBEX35 showing a unidirectional causal relationship with EUA prices, particularly during Phase II of the EU ETS, which coincided with the financial crisis.

Our work also highlights the limitations of traditional linear methods in detecting non-linear relationships. By applying the Imbalance Gain measure, we demonstrate how this technique can provide a more accurate representation of causal contributions, distinguishing genuine causal links from spurious ones. This is especially valuable in the context of energy and carbon markets, where non-linear dynamics are prevalent.

In conclusion, our findings contribute to the ongoing evolution of EU ETS research by addressing key gaps, proposing innovative methodologies, and offering new insights into the drivers of carbon market fluctuations. Although much progress has been made in understanding the financial aspects of emissions trading, our study emphasises the need for continued development in non-parametric analysis, macroeconomic integration, and the exploration of causal relationships, which will ultimately aid in refining policies aimed at combating climate change.

Chapter 8

Further Research

8.0.1 Causal Graph Reconstruction Algorithm

The ongoing development of a non-parametric methodology for constructing causal networks represents a promising avenue for advancing our understanding of complex dynamic systems. This approach, currently under refinement, aims to provide a more adaptable and data-driven framework, bypassing the limitations imposed by strict parametric assumptions. Its flexibility makes it particularly well-suited for capturing intricate causal interactions in high-dimensional and non-linear systems. As the methodology evolves, it emphasises the integration of innovative measures, such as Differential Information Imbalance (DII), to quantify predictive relationships with greater precision. Through continued optimisation and the incorporation of regularisation techniques, the approach aspires to achieve reliable and interpretable causal graphs that accurately reflect the system dynamics. These developments have potential to transform the way causal relationships are explored and visualised, offering significant insights into key drivers of behaviour in a range of scientific and practical applications.

The construction of a causal graph is based on the principles of dynamic system analysis, specifically targeting the relationships among variables over time. This process involves several key steps aimed at understanding the interactions within the system under investigation. The initial step in constructing the causal graph involves representing the dynamical system through a time series dataset, denoted as $T \in \mathbb{R}^{N_{\text{times}} \times D}$. Here, N_{times} represents the total number of time steps, while D signifies the number of dynamical variables. This representation captures the temporal evolution of the system, allowing for an analysis of how these variables influence each other at various points in time.

To explain the causal relationships among the variables, the concept of nearest neighbours is used. For a given variable at a specific time t , the nearest neighbours in the feature space are identified. This involves calculating a distance metric, the Euclidean distance, which quantifies how similar or dissimilar the states of the variables are at different time points. The choice of distance metric is critical, as it influences the subsequent analysis of interactions. The nearest neighbour search is conducted over a set of time lags, τ , which represent the temporal intervals at which the relationships between variables are assessed. By examining the nearest neighbours at these different time lags, a more comprehensive view of the temporal dynamics of the system can be constructed. A central concept in the analysis is the Differential Information Imbalance (DII), which serves as a measure of the predictive relationships among the variables. DII quantifies the extent to which knowledge about one variable can improve the prediction of another variable in a future time step. Mathematically, the DII can be represented as:

$$\text{DII}(A, B) = \sum_{i=1}^D w_i \cdot L(d_A, d_B),$$

where w_i are weights that denote the influence of each variable i and L is a loss function that measures the discrepancy between the observed and predicted states. The optimisation process aims to minimise DII through iterative adjustment of the weights w_i . This is achieved using gradient descent techniques, which allow for the systematic updating of weights based on the gradient of the loss function [26]. By minimising the DII, the understanding of how variables interact causally is improved. The optimisation process can incorporate regularisation techniques, such as L1 regularisation, to mitigate over-fitting. This is particularly important when dealing with high-dimensional data, where many variables may exhibit weak correlations. Regularisation helps to ensure that only the most significant relationships are retained in the causal graph. Upon convergence of the optimisation algorithm, the final weights w_i are extracted to characterise the causal influences among the variables.

Specifically, the maximum weight associated with each variable across the different time lags and embedding is retained, reflecting the strongest causal relationships identified during the analysis. This process culminates in the construction of a causal graph, which visually represents the relationships between the variables. Each node in the graph corresponds to a variable, while directed edges indicate the strength and direction of causal influence. This graphical representation provides a powerful tool for interpreting the dynamics of the system and identifying key drivers of behaviour.

In comparison to the VAR (Vector Autoregression) spillover methodology, which relies on linear dependencies and predefined structures to capture the transmission of shocks across variables, the non-parametric approach

discussed here offers greater flexibility by avoiding stringent parametric assumptions. While VAR spillover analysis provides a clear framework for understanding linear interdependencies and forecasting based on historical relationships, it may struggle with high-dimensional or non-linear systems where interactions are more complex. The incorporation of Differential Information Imbalance (DII) in the non-parametric method enables a more nuanced analysis of causal interactions, particularly in capturing subtle and non-linear dependencies that VAR models might overlook. This makes the non-parametric approach more adaptable to dynamic and intricate systems, potentially providing deeper insights into the underlying causal mechanisms.

Bibliography

- [1] Piia Aatola, Markku Ollikainen, and Anne Toppinen. Price determination in the eu ets market: Theory and econometric analysis with market fundamentals. *Energy Economics*, 36:380–395, 2013.
- [2] Oluwasegun B Adekoya. Predicting carbon allowance prices with energy prices: A new approach. *Journal of Cleaner Production*, 282:124519, 2021.
- [3] Emilie Alberola, Julien Chevallier, and Benoît Chèze. Price drivers and structural breaks in european carbon prices 2005–2007. *Energy policy*, 36(2):787–797, 2008.
- [4] Carlos Aller, Lorenzo Ductor, and Daryna Grechyna. Robust determinants of co2 emissions. *Energy Economics*, 96:105154, 2021.
- [5] Matteo Allione. Towards the estimation of causal networks by the information imbalance approach, 2024.
- [6] Carlo Altavilla, Raffaella Giacomini, and Riccardo Costantini. Bond returns and market expectations. *Journal of Financial Econometrics*, 12(4):708–729, 2014.
- [7] Carl-Philipp Anke, Hannes Hobbie, Steffi Schreiber, and Dominik Möst. Coal phase-outs and carbon prices: Interactions between eu

- emission trading and national carbon mitigation policies. *Energy Policy*, 144:111647, 2020.
- [8] Massimo Aria and Corrado Cuccurullo. bibliometrix: An r-tool for comprehensive science mapping analysis. *Journal of Informetrics*, 11(4):959–975, 2017.
- [9] Mohamed El Hédi Aroui, Fredj Jawadi, and Duc Khuong Nguyen. Nonlinearities in carbon spot-futures price relationships during phase ii of the eu ets. *Economic Modelling*, 29(3):884 – 892, 2012. Cited by: 71.
- [10] Aydin Aslan and Peter N. Posch. Does carbon price volatility affect european stock market sectors? a connectedness network analysis. *Finance Research Letters*, 50, 2022. Cited by: 14.
- [11] Mustafa H. Babiker. Climate change policy, market structure, and carbon leakage. *Journal of International Economics*, 65(2):421–445, 2005.
- [12] Scott R. Baker, Nicholas Bloom, and Steven J. Davis. Measuring Economic Policy Uncertainty*. *The Quarterly Journal of Economics*, 131(4):1593–1636, 07 2016.
- [13] Anca Claudia Balietti. Trader types and volatility of emission allowance prices. evidence from eu ets phase i. *Energy Policy*, 98:607 – 620, 2016. Cited by: 42.
- [14] Christiane Baumeister and Lutz Kilian. Forty years of oil price fluctuations: Why the price of oil may still surprise us. *Journal of Economic Perspectives*, 30(1):139–60, February 2016.

- [15] Joscha Beckmann and Robert Czudaj. Exchange rate expectations and economic policy uncertainty. *European journal of political economy*, 47:148–162, 2017.
- [16] Geert Bekaert, Marie Hoerova, and Marco Lo Duca. Risk, uncertainty and monetary policy. *Journal of Monetary Economics*, 60(7):771–788, 2013.
- [17] Eva Benz and Stefan Trück. Modeling the price dynamics of co2 emission allowances. *Energy Economics*, 31(1):4–15, 2009.
- [18] Eva Benz and Stefan Trück. Modeling the price dynamics of co2 emission allowances. *Energy Economics*, 31(1):4–15, 2009.
- [19] Eva Benz and Stefan Trück. Modeling the price dynamics of co2 emission allowances. *Energy Economics*, 31(1):4 – 15, 2009. Cited by: 429.
- [20] F. Biancalani, G. Gnecco, and R. et al. Metulini. The impact of the european union emissions trading system on carbon dioxide emissions: a matrix completion analysis. *Scientific Reports*, 14:19676, 2024.
- [21] Bloomberg L.P. *Bloomberg Professional Services: Futures & Commodities Reference Guide*, 2023. Accessed via Bloomberg Terminal. Documentation available via ‘HELP HELP‘ or by searching ”MO1 COMDTY”.
- [22] Christoph Böhringer and Knut Einar Rosendahl. Europe beyond coal—an economic and climate impact assessment. *Journal of Environmental Economics and Management*, 113:102658, 2022.

- [23] C. Kaan Bolat, Ugur Soytaş, Bulent Akinoglu, and Saban Nazlioglu. Is there a macroeconomic carbon rebound effect in eu ets? *Energy Economics*, 125:106879, 2023.
- [24] Massimo Bordignon and Duccio Gamannossi degl’Innocenti. Third time’s a charm? assessing the impact of the third phase of the eu ets on co2 emissions and performance. *Sustainability*, 15(8), 2023.
- [25] Simone Borghesi and Andrea Flori. With or without u(k): A pre-brexit network analysis of the eu ets. *PLOS ONE*, 14(9):1–17, 09 2019.
- [26] Léon Bottou, Frank E. Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311, 2018.
- [27] Samuel C Bradford. Sources of information on specific subjects. *Engineering*, 137:85–86, 1934.
- [28] Frédéric Branger and Philippe Quirion. Would border carbon adjustments prevent carbon leakage and heavy industry competitiveness losses? insights from a meta-analysis of recent economic studies. *Ecological Economics*, 99:29–39, 2014.
- [29] Don Bredin and Cal Muckley. An emerging equilibrium in the eu emissions trading scheme. *Energy Economics*, 33(2):353 – 362, 2011. Cited by: 111.
- [30] Suk Joon Byun and Hangjun Cho. Forecasting carbon futures volatility using garch models with energy volatilities. *Energy Economics*, 40:207–221, 2013.

- [31] Christoph Böhringer, Edward J. Balistreri, and Thomas F. Rutherford. The role of border carbon adjustment in unilateral climate policy: Overview of an energy modeling forum study (emf 29). *Energy Economics*, 34:S97–S110, 2012.
- [32] Bertrand Candelon and Jean-Baptiste Hasse. Testing for causality between climate policies and carbon emissions reduction. *Finance Research Letters*, 55:103878, 2023.
- [33] Julien Chevallier. Carbon futures and macroeconomic risk factors: A view from the eu ets. *Energy Economics*, 31(4):614–625, 2009.
- [34] Julien Chevallier. Carbon futures and macroeconomic risk factors: A view from the eu ets. *Energy Economics*, 31(4):614 – 625, 2009.
- [35] Julien Chevallier. Carbon price drivers: An updated literature review. *International Journal of Applied Logistics*, 4(2):1–19, 2011.
- [36] Julien Chevallier. Detecting instability in the volatility of carbon prices. *Energy Economics*, 33(1):99 – 110, 2011. Cited by: 85; All Open Access, Green Open Access.
- [37] Julien Chevallier. Macroeconomics, finance, commodities: Interactions with carbon markets in a data-rich model. *Economic modelling*, 28(1-2):557–567, 2011.
- [38] Julien Chevallier. A model of carbon price interactions with macroeconomic and energy dynamics. *Energy Economics*, 33(6):1295–1312, 2011.
- [39] Julien Chevallier. Carbon price drivers: an updated literature review. *International Journal of Applied Logistics (IJAL)*, 4(4):1–7, 2013.

- [40] Julien Chevallier, Yannick Le Pen, and Benoît Sévi. Options introduction and volatility in the eu ets. *Resource and Energy Economics*, 33(4):855–880, 2011.
- [41] Julien Chevallier, Yannick Le Pen, and Benoît Sévi. Options introduction and volatility in the eu ets. *Resource and Energy Economics*, 33(4):855 – 880, 2011. Cited by: 39; All Open Access, Green Open Access.
- [42] Julien Chevallier, Duc Khuong Nguyen, and Juan Carlos Reboredo. A conditional dependence approach to co2-energy price relationships. *Energy Economics*, 81:812–821, 2019.
- [43] Julien Chevallier and Benoît Sévi. On the realized volatility of the ecx co2 emissions 2008 futures contract: Distribution, dynamics and forecasting. *Annals of Finance*, 7(1):1 – 29, 2011. Cited by: 40; All Open Access, Green Open Access.
- [44] Dohyun Chun, Hoon Cho, and Jihun Kim. The relationship between carbon-intensive fuel and renewable energy stock prices under the emissions trading system. *Energy Economics*, 114:106257, 2022.
- [45] Ronald H. Coase. The problem of social cost. *Journal of Law and Economics*, 3(1):1–44, 1960.
- [46] European Commission. Kyoto 1st commitment period (2008–12). Accessed 2024-11-14.
- [47] European Commission. Directive 2003/87/ec of the european parliament and of the council of 13 october 2003 establishing a scheme for greenhouse gas emission allowance trading within the community and amending council directive 96/61/ec, 2005. Accessed 2024-11-14.

- [48] European Commission. Proposal for a directive of the european parliament and of the council amending directive 2003/87/ec to improve and extend the eu emissions trading system, 2011. Accessed: 2024-11-17.
- [49] European Commission. Eu ets and aviation sector. *European Commission Journal*, 2012. Accessed: 2024-11-14.
- [50] European Commission. Eu ets phase 2 overview, 2015. Accessed: 2024-11-14.
- [51] European Commission. Renewable energy directive, 2018. Accessed: April 6, 2023.
- [52] European Commission. Eu ets phase 4, 2020. Accessed: 2024-11-14.
- [53] European Commission. The european green deal. https://ec.europa.eu/info/strategy/priorities-2019-2024/european-green-deal_en, 2020. Accessed: 2024-11-14.
- [54] European Commission. Fit for 55 package: Delivering the european green deal, 2020. Accessed: 2024-11-17.
- [55] European Commission. European commission - eu ets, 2023. Accessed: April 6, 2023.
- [56] European Commission. Auctioning of allowances under the eu ets), 2024. Accessed: 2024-11-17.
- [57] European Commission. Carbon border adjustment mechanism (cbam), 2024. Accessed: 2024-11-17.
- [58] European Commission. Eu ets - buildings, road transport, and additional sectors, 2024. Accessed: 2024-11-17.

- [59] European Commission. Eu ets and sectors covered, 2024. Accessed: 2024-11-17.
- [60] European Commission. Monitoring, reporting, and verification under the eu ets, 2024. Accessed: 2024-11-17.
- [61] European Commission. Regulatory oversight of the eu ets, 2024. Accessed: 2024-11-17.
- [62] U.S. Congress. *Clean Air Act Amendments of 1990*. U.S. Government Printing Office, Washington, D.C., 1990.
- [63] Christian Conrad, Daniel Rittler, and Waldemar Rotfuß. Modeling and explaining the dynamics of european union allowance prices at high-frequency. *Energy Economics*, 34(1):316 – 326, 2012. Cited by: 105; All Open Access, Green Open Access.
- [64] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, 2nd edition, 2006.
- [65] Anna Creti, Pierre-André Jouvét, and Valérie Mignon. Carbon price drivers: Phase i versus phase ii equilibrium? *Energy Economics*, 34(1):327–334, 2012.
- [66] Anna Creti, Pierre-André Jouvét, and Valérie Mignon. Carbon price drivers: Phase i versus phase ii equilibrium? *Energy Economics*, 34(1):327 – 334, 2012.
- [67] Anna Cretí and Marc Joëts. Multiple bubbles in the european union emission trading scheme. *Energy Policy*, 107:119 – 130, 2017. Cited by: 38.

- [68] Thomas D. Crocker. The structuring of atmospheric pollution control systems. In Harold Wolozin, editor, *The Economics of Air Pollution*, pages 61–86. W.W. Norton & Company, 1966.
- [69] Caldara D. and Iacoviello M. Geopolitical risk index, 2022. Accessed: April 24, 2023.
- [70] Amir HM Daghigh and Majid M Ghazani. Analyzing the drivers of co2 allowance prices in eu ets under the covid-19 pandemic: Considering memd approach with a novel filtering procedure. *Journal of Cleaner Production*, 427:139043, 2023.
- [71] Peng-Fei Dai, Xiong Xiong, Toan Luu Duc Huynh, and Jiqiang Wang. The impact of economic policy uncertainties on the volatility of european carbon market. *Journal of Commodity Markets*, 26, 2022. Cited by: 43; All Open Access, Green Open Access.
- [72] John H. Dales. *Pollution, Property and Prices: An Essay in Policy-Making and Economics*. University of Toronto Press, Toronto, Canada, 1968.
- [73] Ajay K. Dhamija, Surendra S. Yadav, and Pk Jain. Forecasting volatility of carbon under eu ets: a multi-phase study. *Environmental Economics and Policy Studies*, 19(2):299 – 335, 2017. Cited by: 20.
- [74] Ajay K. Dhamija, Surendra S. Yadav, and Pk Jain. Volatility spillover of energy markets into eua markets under eu ets: a multi-phase study. *Environmental Economics and Policy Studies*, 20(3):561 – 591, 2018. Cited by: 19.

- [75] A. Denny Ellerman and Barbara K. Buchner. *The European Union Emissions Trading System: Origins, Allocation, and Early Results*, volume 1. Review of Environmental Economics and Policy, 2007.
- [76] A. Denny Ellerman, Paul L. Joskow, Richard Schmalensee, Juan Pablo Montero, and Elizabeth M. Bailey. *Markets for Clean Air: The U.S. Acid Rain Program*. Cambridge University Press, Cambridge, UK, 2000.
- [77] A. Denny Ellerman, Claudio Marcantonini, and Aleksandar Zaklan. The european union emissions trading system: Ten years and counting. *Review of Environmental Economics and Policy*, 10(1):89–107, 2016.
- [78] Robert F Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica: Journal of the econometric society*, pages 987–1007, 1982.
- [79] María Eugenia Sanin, Francesco Violante, and María Mansanet-Bataller. Understanding volatility dynamics in the eu-ets market. *Energy Policy*, 82(1):321 – 331, 2015. Cited by: 71; All Open Access, Green Open Access.
- [80] European Commission. Eu and switzerland link carbon markets, 2019. Accessed: 2024-12-19.
- [81] European Securities and Markets Authority. European securities and markets authority (esma), 2024. Accessed: 2024-12-16.
- [82] Zhen-Hua Feng, Yi-Ming Wei, and Kai Wang. Estimating risk for the carbon market via extreme value theory: An empirical analysis of the eu ets. *Applied Energy*, 99:97 – 108, 2012. Cited by: 75.

- [83] Zhen-Hua Feng, Le-Le Zou, and Yi-Ming Wei. Carbon price volatility: Evidence from eu ets. *Applied Energy*, 88(3):590 – 598, 2011. Cited by: 154.
- [84] Laurent Ferrara and Joseph Yapi. Measuring exchange rate risks during periods of uncertainty. *International Economics*, 170:202–212, 2022.
- [85] Ronald A. Fisher. *The Design of Experiments*. Oliver and Boyd, Edinburgh, 1935.
- [86] Eric Ghysels. Macroeconomics and the reality of mixed frequency data. *Journal of Econometrics*, 193(2):294–314, 2016.
- [87] Eric Ghysels, Virmantas Kvedaras, and Vaidotas Zemlys-Balevičius. Mixed data sampling (midas) regression models. In *Handbook of Statistics*, volume 42, pages 117–153. Elsevier, 2020.
- [88] Luis A. Gil-Alana, Rangan Gupta, and Fernando Perez de Gracia. Modeling persistence of carbon emission allowance prices. *Renewable and Sustainable Energy Reviews*, 55:221–226, 2016.
- [89] Aldo Glielmo, Claudio Zeni, Bingqing Cheng, Gábor Csányi, and Alessandro Laio. Ranking the information content of distance measures. *PNAS nexus*, 1(2):pgac039, 2022.
- [90] Xu Gong, Rong Shi, Jun Xu, and Boqiang Lin. Analyzing spillover effects between carbon and fossil energy markets from a time-varying perspective. *Applied Energy*, 285:116384, 2021.
- [91] C. W. J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, 37(3):424–438, 1969.

- [92] C.W.J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, 1969.
- [93] C.W.J. Granger. Testing for causality: A personal viewpoint. *Journal of Economic Dynamics and Control*, 2:329–352, 1980.
- [94] Michael Grubb. The earth summit agreements: A guide and assessment. *Energy Policy*, 23(4–5):341–345, 1995.
- [95] Abebe Hailemariam, Kris Ivanovski, and Ratbek Dzhumashev. Does r&d investment in renewable energy technologies reduce greenhouse gas emissions? *Applied Energy*, 327:120056, 2022.
- [96] James D. Hamilton. Time series analysis. *Princeton University Press*, 1994. Chapter 16 on Granger Causality Tests.
- [97] Shawkat Hammoudeh, Duc Khuong Nguyen, and Ricardo M Sousa. Energy prices and co2 emission allowance prices: A quantile regression approach. *Energy policy*, 70:201–206, 2014.
- [98] Shawkat Hammoudeh, Duc Khuong Nguyen, and Ricardo M Sousa. What explain the short-term dynamics of the prices of co2 emissions? *Energy Economics*, 46:122–135, 2014.
- [99] Waqas Hanif, Jose Arreola Hernandez, Walid Mensi, Sang Hoon Kang, Gazi Salah Uddin, and Seong-Min Yoon. Nonlinear dependence and connectedness between clean/renewable energy sector equity and european emission allowance prices. *Energy Economics*, 101:105409, 2021.
- [100] Aron Denes Hartvig, Aron Pap, and Peter Palos. EU Climate Change News Index: Forecasting EU ETS prices with online news. *Finance Research Letters*, 54, 2023.

- [101] Sean Healy, Katja Schumacher, Anamaria Stroia, and Stephan Slingerland. Review of literature on eu ets performance. *A literature review and gap analysis of policy evaluations. Öko-Institut Working Paper*, 2, 2015.
- [102] Beat Hintermann. Allowance price drivers in the first phase of the eu ets. *Journal of Environmental Economics and Management*, 59(1):43–56, 2010.
- [103] Martin Hintermayer. A carbon price floor in the reformed eu ets: Design matters! *Energy Policy*, 147, 2020.
- [104] Steffen Hitzemann, Marliese Uhrig-Homburg, and Karl-Martin Ehrhart. Emission permits and the announcement of realized emissions: Price impact, trading volume, and volatilities. *Energy Economics*, 51:560 – 569, 2015. Cited by: 12.
- [105] Paul W. Holland. Statistics and causal inference. *Journal of the American Statistical Association*, 81(396):945–960, 1986.
- [106] Yu-Jie Hu, Lishan Yang, Fali Duan, Honglei Wang, and Chengjiang Li. A scientometric analysis and review of the emissions trading system. *Energies*, 15(12):4423, 2022.
- [107] Chang-Jing Ji, Xiao-Yi Li, Yu-Jie Hu, Xiang-Yu Wang, and Bao-Jun Tang. Research on carbon price in emissions trading scheme: a bibliometric analysis. *Natural Hazards*, 99:1381–1396, 2019.
- [108] Qiang Ji, Dayong Zhang, and Jiang-bo Geng. Information linkage, dynamic spillovers in prices and volatility between the carbon and energy markets. *Journal of Cleaner Production*, 198:972–978, 2018.

- [109] Qiang Ji, Dayong Zhang, and Jiang-bo Geng. Information linkage, dynamic spillovers in prices and volatility between the carbon and energy markets. *Journal of Cleaner Production*, 198:972 – 978, 2018. Cited by: 240.
- [110] Jun-Jun Jia, Jin-Hua Xu, and Ying Fan. The impact of verified emissions announcements on the european union emissions trading scheme: A bilaterally modified dummy variable modelling analysis. *Applied Energy*, 173:567 – 577, 2016. Cited by: 25.
- [111] Rebeca Jiménez Rodríguez. What happens to the relationship between eu allowances prices and stock market indices in europe? *Energy Economics*, 81:13–24, 2019.
- [112] Memoona Kanwal and Hashim Khan. Does carbon asset add value to clean energy market? evidence from eu. *Green Finance*, 3(4):495–507, 2021.
- [113] Andreas Karpf, Antoine Mandel, and Stefano Battiston. Price and network dynamics in the european carbon market. *Journal of Economic Behavior & Organization*, 153:103–122, 2018.
- [114] Jan Horst Keppler and Maria Mansanet-Bataller. Causalities between co₂, electricity, and other energy variables during phase i and phase ii of the eu ets. *Energy policy*, 38(7):3329–3341, 2010.
- [115] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- [116] Nicolas Koch, Sabine Fuss, Godefroy Grosjean, and Ottmar Edenhofer. Causes of the eu ets price drop: Recession, cdm, renewable policies or a bit of everything?—new evidence. *Energy Policy*, 73:676–685, 2014.

- [117] Maximilian Konradt, Thomas McGregor, and Frederik G Toscani. Carbon prices and inflation in the euro area. *IMF Working Papers*, 2024(031):A001, 2024.
- [118] William H. Kruskal and W. Allen Wallis and. Use of ranks in one-criterion variance analysis. *Journal of the American Statistical Association*, 47(260):583–621, 1952.
- [119] S. Kullback and R. A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.
- [120] Timothy Laing, Misato Sato, Michael Grubb, and Claudia Comberti. The effects and side-effects of the eu emissions trading scheme. *Wiley Interdisciplinary Reviews: Climate Change*, 5(4):509 – 519, 2014.
- [121] Yuqiao Lan, Yubin Huangfu, Zhehao Huang, and Changhong Zhang. Breaking through the limitation of carbon price forecasting: A novel hybrid model based on secondary decomposition and nonlinear integration. *Journal of Environmental Management*, 362, 2024.
- [122] Ming-Fang Li, Hui Hu, and Lu-Tao Zhao. Key factors affecting carbon prices from a time-varying perspective. *Environmental Science and Pollution Research*, 2022. Cited by: 8; All Open Access, Bronze Open Access, Green Open Access.
- [123] Jian Liu, Shuai Tang, and Chun-Ping Chang. Spillover effect between carbon spot and futures market: evidence from eu ets. *Environmental Science and Pollution Research*, 28(12):15223 – 15235, 2021. Cited by: 12.
- [124] Alfred J. Lotka. The frequency distribution of scientific productivity. *Journal of the Washington Academy of Sciences*, 16(12):317–323, 1926.

- [125] Yuliya Lovcha, Alejandro Perez-Laborda, and Iryna Sikora. The determinants of co2 prices in the eu emission trading system. *Applied Energy*, 305:117903, 2022.
- [126] Julio J. Lucia, Maria Mansanet-Bataller, and Angel Pardo. Speculative and hedging activities in the european carbon market. *Energy Policy*, 82(1):342 – 351, 2015. Cited by: 31.
- [127] Benjamin Johannes Lutz, Uta Pigorsch, and Waldemar Rotfuß. Non-linearity in cap-and-trade systems: The eua price and its fundamentals. *Energy Economics*, 40:222 – 232, 2013. Cited by: 100; All Open Access, Green Open Access.
- [128] Maria Mansanet-Bataller, Angel Pardo, and Enric Valor. Co2 prices, energy and weather. *The Energy Journal*, 28(3), 2007.
- [129] Maria Mansanet-Bataller, Angel Pardo, and Enric Valor. Co2 prices, energy and weather. *Energy Journal*, 28(3):73 – 92, 2007.
- [130] Lukas Menkhoff, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf. Currency momentum strategies. *Journal of Financial Economics*, 106(3):660–684, 2012.
- [131] Junlong Mi, Xing Yang, Feifei Huang, and Yufa Xu. Geopolitical, economic risk and the time-varying structure of extreme risk in the carbon emissions trading market. *Frontiers in Environmental Science*, 12, 2024.
- [132] Mahtab Mohtasham Khani, Sahand Vahidnia, and Alireza Abbasi. A deep learning-based method for forecasting gold price with respect to pandemics. *SN Computer Science*, 2(4):335, 2021.

- [133] Roger B. Nelsen. *An Introduction to Copulas*. Springer, New York, 2nd edition, 2006.
- [134] Karsten Neuhoff, Markus Åhman, Regina Betz, Johanna Cludius, Federico Ferrano, Kristina Holmgren, Gabriella Pal, Michael Grubb, Felix Matthes, Karoline Rogge, et al. Implications of announced phase ii national allocation plans for the eu ets. In *National Allocation Plans in the EU Emissions Trading Scheme*, pages 411–422. Routledge, 2010.
- [135] Scott J. Niblock and Jennifer L. Harrison. Carbon markets in times of vuca: a weak-form efficiency investigation of the phase ii eu ets. *Journal of Sustainable Finance and Investment*, 3(1):38 – 56, 2013. Cited by: 17.
- [136] European Parliament and the Council of the European Union. Directive 2003/87/ec of the european parliament and of the council establishing a scheme for greenhouse gas emission allowance trading within the community and amending council directive 96/61/ec, 2003. Accessed: 2024-11-17.
- [137] Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, 2nd edition, 2009.
- [138] Karl Pearson. Mathematical contributions to the theory of evolution. ii. skew variation in homogeneous material. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 187:253–318, 1896. Pearson’s development of the formalized correlation coefficient method.

- [139] Lixin Qiu, Lijun Chu, Ran Zhou, Haitao Xu, and Sai Yuan. How do carbon, stock, and renewable energy markets interact: Evidence from europe. *Journal of Cleaner Production*, 407:137106, 2023.
- [140] Xiaohang Ren, Kun Duan, Lizhu Tao, Yukun Shi, and Cheng Yan. Carbon prices forecasting in quantiles. *Energy Economics*, 108:105862, 2022.
- [141] Xiaohang Ren, Yiying Li, Fenghua Wen, Zudi Lu, et al. The interrelationship between the carbon market and the green bonds market: Evidence from wavelet quantile-on-quantile method. *Technological Forecasting and Social Change*, 179:121611, 2022.
- [142] Reuters. Cop29: What is a carbon credit? what is article 6? <https://www.reuters.com/business/environment/cop29-what-is-carbon-credit-what-is-article-6-2024-11-09/>?, 2024. Accessed: 2025-01-14.
- [143] Stephen Roberts, Michael Osborne, Mark Ebden, Steven Reece, Neale Gibson, and Suzanne Aigrain. Gaussian processes for time-series modelling. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371(1984):20110550, 2013.
- [144] J. Runge. Causal network reconstruction from time series: From theoretical assumptions to practical estimation. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 28(7):075310, 07 2018.
- [145] Cristiano Salvagnin, Aldo Glielmo, Maria Elena De Giuli, and Antonietta Mira. Investigating the price determinants of the european emission trading system: a non-parametric approach. *Quantitative Finance*, 2024.

- [146] Hannah Scheuing and Johanna Kamm. The eu on the road to climate neutrality—is the ‘fit for 55’ package fit for purpose? *Renewable Energy Law and Policy Review*, 10(3-4):4–18, 2022.
- [147] Richard Schmalensee, Paul L. Joskow, A. Denny Ellerman, Juan Pablo Montero, and Elizabeth M. Bailey. An interim evaluation of sulfur dioxide emissions trading. *Journal of Economic Perspectives*, 12(3):53–68, 1998.
- [148] Sara Segura, Luis Ferruz, Pilar Gargallo, and Manuel Salvador. Environmental versus economic performance in the eu ets from the point of view of policy makers: A statistical analysis based on copulas. *Journal of Cleaner Production*, 176:1111–1132, 2018.
- [149] Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.
- [150] Jiaxin Shi, Michalis Titsias, and Andriy Mnih. Sparse orthogonal variational inference for gaussian processes. In *International Conference on Artificial Intelligence and Statistics*, pages 1932–1942. PMLR, 2020.
- [151] A. Sklar. Fonctions de répartition à n dimensions et leurs marges. *Publications de l’Institut de Statistique de l’Université de Paris*, 8:229–231, 1959.
- [152] Skander Soltani. On the use of the wavelet decomposition for time series prediction. *Neurocomputing*, 48(1-4):267–277, 2002.
- [153] Alessandro Spelta and Maria Elena De Giuli. Does renewable energy affect fossil fuel price? a time–frequency analysis for the europe. *Physica A: Statistical Mechanics and its Applications*, 626:129098, 2023.

- [154] Xueping Tan, Kavita Sirichand, Andrew Vivian, and Xinyu Wang. How connected is the carbon market to energy and financial markets? a systematic analysis of spillovers and dynamics. *Energy Economics*, 90:104870, 2020.
- [155] Xueping Tan, Kavita Sirichand, Andrew Vivian, and Xinyu Wang. Forecasting european carbon returns using dimension reduction techniques: Commodity versus financial fundamentals. *International Journal of Forecasting*, 38(3):944–969, 2022.
- [156] Bao-jun Tang, Pi-qin Gong, and Cheng Shen. Factors of carbon price volatility in a comparative analysis of the eua and scer. *Annals of Operations Research*, 255(1-2):157 – 168, 2017. Cited by: 25.
- [157] Vittorio Del Tatto, Gianfranco Fortunato, Domenica Bueti, and Alessandro Laio. Robust inference of causality in high-dimensional dynamical processes from the information imbalance of distance ranks. *Proceedings of the National Academy of Sciences*, 121(19):e2317256121, 2024.
- [158] J. Teixidó, Stefano F. Verde, and Francesco Nicolli. The impact of the eu emissions trading system on low-carbon technological change: The empirical evidence. *Ecological Economics*, 164, 2019.
- [159] Nader Trabelsi, Aviral Kumar Tiwari, Shawkat Hammoudeh, and Noureddine Benlagha. Extreme linkages of carbon futures, energy markets, and economic indicators: A copula approach. *Energy Sources, Part B: Economics, Planning, and Policy*, 18(1):2165738, 2023.

- [160] United Nations Framework Convention on Climate Change. United nations framework convention on climate change, 1992. Accessed: 2024-11-17.
- [161] Border Tax Adjustments versus Rebates. Comparing policies to combat emissions leakage, 2009.
- [162] Elena Villar-Rubio, María-Dolores Huete-Morales, and Federico Galán-Valdivieso. Using egarch models to predict volatility in unconsolidated financial markets: the case of european carbon allowances. *Journal of Environmental Studies and Sciences*, 13(3):500 – 509, 2023. Cited by: 0; All Open Access, Green Open Access, Hybrid Gold Open Access.
- [163] Jingyuan Wang, Ze Wang, Jianfeng Li, and Junjie Wu. Multilevel wavelet decomposition network for interpretable time series analysis. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2437–2446, 2018.
- [164] Jiqian Wang, Xiaozhu Guo, Xueping Tan, Julien Chevallier, and Feng Ma. Which exogenous driver is informative in forecasting european carbon volatility: Bond, commodity, stock or uncertainty? *Energy Economics*, 117, 2023. Cited by: 6.
- [165] Xiang-Yu Wang and Bao-Jun Tang. Review of comparative studies on market mechanisms for carbon emission reduction: A bibliometric analysis. *Natural Hazards*, 94:1141–1162, 2018.
- [166] Yudong Wang and Zhuangyue Guo. The dynamic spillover between carbon and energy markets: new evidence. *Energy*, 149:24–33, 2018.

- [167] Zi Jie Wang and Lu Tao Zhao. The impact of the global stock and energy market on eu ets: A structural equation modelling approach. *Journal of Cleaner Production*, 289:125140, 2021.
- [168] Romina Wild, Vittorio Del Tatto, Felix Wodaczek, Bingqing Cheng, and Alessandro Laio. Automatic feature selection and weighting using differentiable information imbalance, 2024.
- [169] Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.
- [170] Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P Xing. Deep kernel learning. In *Artificial intelligence and statistics*, pages 370–378. PMLR, 2016.
- [171] Kangyin Dong Xiaohang Ren, Yue Dou and Cheng Yan. Spillover effects among crude oil, carbon, and stock markets: evidence from nonparametric causality-in-quantiles tests. *Applied Economics*, 55(38):4486–4509, 2023.
- [172] Zhenhui Xu. On the causality between export growth and gdp growth: an empirical reinvestigation. *Review of International Economics*, 4(2):172–184, 1996.
- [173] Nannan Yuan and Lu Yang. Asymmetric risk spillover between financial market uncertainty and the carbon market: A gas-dcs-copula approach. *Journal of Cleaner Production*, 259:120750, 2020.
- [174] Yue-Jun Zhang and Yi-Ming Wei. An overview of current research on eu ets: Evidence from its operating mechanism and economic effect. *Applied Energy*, 87(6):1804 – 1814, 2010.

- [175] Xin Zhao, Meng Han, Lili Ding, and Wanglin Kang. Usefulness of economic and energy data at different frequencies for carbon price forecasting in the eu ets. *Applied Energy*, 216:132–141, 2018.
- [176] Meirui Zhong, Rui Zhang, and Xiaohang Ren. The time-varying effects of liquidity and market efficiency of the european union carbon market: Evidence from the tvp-svar-sv approach. *Energy Economics*, 123, 2023.
Cited by: 5.

A Appendix

A.1 Gradient-base Optimisation

Gradient-based optimisation methods are foundational for many machine learning and statistical inference techniques. These methods rely on computing gradients of an objective function to iteratively update model parameters in the direction that minimises or maximises the function. The core idea is that by following the negative gradient, convergence towards an optimal solution that best fits the data or model requirements can be achieved. This subsection delves into the principles of gradient-based optimisation, with a particular focus on *Gradient Descent*, one of the most widely used optimisation algorithms.

Principles of Gradient-Based Optimisation

The basic principle behind gradient-based optimisation is to find the minimum (or maximum) of a differentiable function by iteratively moving towards the steepest descent (or ascent) direction. Mathematically, if an objective function $f(\boldsymbol{\theta})$ depends on parameters $\boldsymbol{\theta}$, the aim is to find the set of parameters $\boldsymbol{\theta}^*$ that minimises the function, i.e.:

$$\boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} f(\boldsymbol{\theta}). \quad (1)$$

In each step of the optimisation process, the parameter vector $\boldsymbol{\theta}$ is updated based on the gradient of the function with respect to $\boldsymbol{\theta}$. The update rule is:

$$\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} - \eta \nabla f(\boldsymbol{\theta}^{(t)}), \quad (2)$$

where $\boldsymbol{\theta}^{(t)}$ is the parameter vector in iteration t , η is the learning rate (a small positive scalar), $\nabla f(\boldsymbol{\theta}^{(t)})$ is the gradient of the function f with respect to the parameters $\boldsymbol{\theta}$ in iteration t .

The learning rate η controls the size of the steps taken in the direction of the gradient. A high learning rate can lead to overshooting, while a very small learning rate can make the optimisation process very slow.

Gradient Descent Algorithm

Gradient Descent is the most basic form of gradient-based optimisation and can be understood as a simple iterative procedure for minimising a differentiable objective function. The algorithm consists of the following steps:

Algorithm 1 Gradient Descent Optimization Process. This figure illustrates the iterative process of gradient descent optimisation, highlighting how the algorithm updates model parameters to minimise the loss function over successive iterations.

- 1: **Input:** Initial parameters $\boldsymbol{\theta}^{(0)}$, learning rate η , maximum iterations T , convergence tolerance ϵ
- 2: **Output:** Optimized parameters $\boldsymbol{\theta}^*$
- 3: Initialize $\boldsymbol{\theta}$ randomly or based on prior knowledge
- 4: **for** $t = 1$ to T **do**
- 5: Compute the gradient $\nabla f(\boldsymbol{\theta}^{(t)})$
- 6: Update the parameters:

$$\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} - \eta \nabla f(\boldsymbol{\theta}^{(t)}) \quad (3)$$

- 7: **if** Convergence criterion is met, i.e., $\|\boldsymbol{\theta}^{(t+1)} - \boldsymbol{\theta}^{(t)}\| < \epsilon$ **then**
 - 8: **break**
 - 9: **end if**
 - 10: **end for**
 - 11: **Return** $\boldsymbol{\theta}^{(t+1)}$
-

Gradient Descent is widely used in many applications, especially when dealing with large-scale data and high-dimensional parameter spaces. However, the basic form of gradient descent may be computationally expensive,

due to the need to compute gradients over the entire dataset (which can be large).

Variants of Gradient Descent

There are several important variants of Gradient Descent that help improve its efficiency and convergence speed, particularly in large-scale problems:

Stochastic Gradient Descent (SGD). In *Stochastic Gradient Descent (SGD)*, instead of computing the gradient using the entire dataset, the gradient is computed based on a single randomly selected data point (or a small mini-batch) at each iteration. This reduces the computational cost significantly and introduces noise into the optimisation process, which can help escape local minima. The update rule for SGD is:

$$\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} - \eta \nabla f(\boldsymbol{\theta}^{(t)}, x_i), \quad (4)$$

where x_i represents a single data point (or a mini-batch of data points).

Mini-Batch Gradient Descent. *Mini-batch Gradient Descent* is a compromise between full-batch gradient descent and SGD. In mini-batch gradient descent, the gradient is computed using a small subset (mini-batch) of the dataset rather than using the entire dataset or just one point. This method provides a balance between computational efficiency and the ability to escape local minima, combining the benefits of batch and stochastic methods. In the analysis, this approach was employed to optimise the performance of the model. The update rule is:

$$\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} - \eta \nabla f(\boldsymbol{\theta}^{(t)}, \mathcal{B}), \quad (5)$$

where \mathcal{B} represents the mini-batch of data.

Convergence and Challenges

Although Gradient descent is widely used and often effective, it comes with several challenges that must be considered:

Local Minima and Saddle Points. For non-convex objective functions, such as those found in deep learning models, Gradient Descent may get stuck in *local minima* or *saddle points*. These are points where the gradient is zero but are not the global minima of the objective function. The use of *Stochastic Gradient Descent (SGD)* or *Mini-Batch Gradient Descent* can help alleviate this problem by introducing noise into the optimisation process, making it more likely to escape local minima.

Learning Rate Selection. Selecting an appropriate learning rate is crucial to the success of gradient-based optimisation. A large learning rate can lead to overshooting, where the algorithm diverges and fails to converge, while a very small learning rate can result in slow convergence. A popular approach to mitigate this issue is the use of learning rate schedules, which adapt the learning rate during the optimisation process, or adaptive learning rate algorithms such as *Adam*.

Over-fitting and Regularisation. When using gradient descent to fit a model, over-fitting can occur if the model is too complex relative to the amount of data available. To combat this, regularisation techniques such as L1 or L2 regularisation can be incorporated into the objective function, penalising overly large weights and promoting simpler models. The updated objective function becomes:

$$f(\boldsymbol{\theta}) + \lambda \|\boldsymbol{\theta}\|_2^2, \quad (6)$$

where λ is a regularisation parameter.

A.2 Details on the DII optimisation

Technical details about DII optimisation are reported in this section.

A relevant hyperparameter of the method is the parameter λ appearing in the softmax coefficients of Eq. (3.16), which defines the size of the neighbourhoods in the first distance space. Although the classical Information Imbalance is recovered in the limit $\lambda \rightarrow 0$, too small values of λ make the DII optimisation inefficient, as decreasing λ also decreases the magnitude of the DII derivatives employed in the gradient descent updates [168].

In this work, λ is computed using a point-adaptive scheme. Specifically, a distinct value λ_i is used for each set of coefficients c_{ij}^N . This approach allows handling data sets where points are not homogeneously distributed, and where a single distance scale would result in different numbers of neighbours for points in different regions of the data manifold.

Specifically, each λ_i is computed as:

$$\lambda_i = 0.1 d_{ij(k)}^{\mathbf{x}}(\mathbf{w})^2, \quad (7)$$

where 0.1 is an empirical prefactor and $j(k)$ is the k -th nearest neighbour of i according to distance $d^{\mathbf{x}}(\mathbf{w})$. In all the optimisations carried out in this work, k was fixed to 5% of the points entering the calculation of the DII (i.e., $k/N = 5\%$). To keep into account the variation of $d^{\mathbf{x}}(\mathbf{w})$ during the DII optimisation, the set of $\{\lambda_i\}$ is recomputed at each gradient descent update.

Other relevant aspects that may affect the convergence of the DII to its

global minimum are the choice of the optimiser and the use of mini-batches.

The use of mini-batches, namely the computation of the gradient on random subsets of points at each gradient descent update, is a well-known strategy that can improve convergence speed and stability of training in optimisation algorithms, especially when the loss function features several local minima. In this work, all the optimisations were carried out by selecting $N = 2800$ distinct frames from the original time series, and randomly splitting the data set into 28 mini-batches with $N' = 100$ points in each training epoch. Using this approach, the value of k defining the set of $\{\lambda_i\}$ was fixed to 5 in each mini-batch. Beyond improving the convergence properties, in our application the use of small mini-batches allows to sample time frames which are likely to be uncorrelated, satisfying one of the requirements of our approach. In contrast, computing the DII without mini-batches would require an undersampling of the original time series to fulfil this requirement, reducing the statistical significance of the results when dealing with a short time series. Rather than using the vanilla gradient descent approach provided in Eq. (3.17), in this work the Adam optimiser [115] is used, which is widely employed by the machine learning community and well-known for its good convergence properties. All optimisations were carried out for 2000 training epochs, setting the initial learning rate η to 10^{-3} , and progressively decreasing it to zero according to a cosine decay schedule. After learning the optimal weights, the DII terms appearing in the IG expression of Eq. (3.21) were computed over all the $N = 2800$ time frames originally extracted.

In the analysis of the financial returns, in order to satisfy the requirement of independent initial conditions in the DII calculation, the final DII estimates were carried out by discarding, for each frame t , distances with points within a time window $[t - 1, t + 1]$. Similarly, in both systems of

Sec. 3.2.4 the final IGs were computed by discarding distances between each frame t and points within $[t - 5, t + 5]$.

A.3 Mixed frequency forecasting with Gaussian Processes

A non-parametric method is proposed, leveraging Gaussian Processes (GPs) in conjunction with the Information Imbalance to optionally aggregate data from different time frequencies.




Figure A1: **Imputation and aggregation using GPs.** In the left panel, a GP is used to impute the GDP time series variable from quarterly to weekly frequency. In the right panel, a GP is used to aggregate the target time series of EUA price from daily to weekly frequency.

Figure A1 illustrates an example of imputation and aggregation of time series using GPs. In the imputation process (left panel of the figure), a GP is fitted to the data with a low noise level of $\sigma_n^2 = 10^{-3}$, sufficient to regularise the inversion in Eq. 3.26. The posterior GP mean is then used to estimate the value of the time series at the required frequencies. For aggregation (right

panel of the figure), the GP noise level is set to the average rolling variance of the time series, computed with the target period as the rolling window, and the posterior GP mean is again used to determine the value of the time series at the needed frequencies. GPs were also employed to conduct experiments in nowcasting and forecasting of allowance prices.

Besides, a k -fold cross-validation with $k = 5$ was applied. In all experiments, the smoothness parameter, ν , was set to 1.5, the highest degree of smoothness compatible with all available time series, while the length scale, l , was selected separately for each fit by maximising likelihood [169]. GPs are highly flexible stochastic processes and have been extensively used in the literature to model non-Gaussian data [169]. For example, Figure A1 demonstrates how a GP effectively fits the clearly non-Gaussian EUA price data.