Ollscoil na hÉireann, Corcaigh

National University of Ireland, Cork

Learning Short-Term Electricity Market Behaviours

Thesis presented by

Joseph Collins, MSc

(ORCID: 0009-0008-2521-2620)

for the degree of

Doctor of Philosophy

School of Mathematical Sciences

Head of School: Dr. Martin Kilian

Supervisors: Dr. Kieran Mulchrone, Dr. Andreas Amann

December 2025

Contents

Declaration

This is to certify that the work I am submitting is my own and has not been submitted for another degree, either at University College Cork or elsewhere. All external references and sources are clearly acknowledged and identified within the contents. I have read and understood the regulations of University College Cork concerning plagiarism and intellectual property.

Joseph Collins

Digital signature of the candidate

Acknowledgements

In the latter stages of this journey, it often occurs to me how much a PhD resembles running a marathon. To outside observers, the focus is often on the moment the runner crosses the finish line; the sense of relief and accomplishment is palpable. What often goes unseen is the immense support required just to reach the start line in the first place.

In my case, when I began this journey around seven years ago, my supervisors Kieran Mulchrone and Andreas Amann, had only a rough sense of my "running form," but were nonetheless willing to take a chance on an enthusiastic and motivated PhD student. Over several years, even during times when neither a clear start nor finish line was in sight, Kieran and Andreas generously gave their time, despite demanding teaching and research commitments. This journey would not have been possible without their guidance and support. I'm deeply grateful for the knowledge they shared, and profoundly indebted for their patience, encouragement, and endurance throughout.

I'm also thankful to my colleagues at capSpire for their flexibility and support in allowing me to take time off over the years to train for and pursue this passion project. Along the way, I've also encountered many fellow runners, too many to name, who offered encouragement and insight. In particular, I'd like to thank Ciaran O'Connor for acting as a sounding board and providing thoughtful feedback at various stages of the research.

As anyone who has trained for a marathon knows, there's no escaping the long training runs. While these are necessary, and even enjoyable at times, they often come at the expense of other exciting things. To Saoirse, Katelyn, and Anna-Mae: seeing the enthusiasm and energy that you brought to your own pursuits was an inspiration to me at various stages in my journey, I'm really looking forward to uninterrupted weekends and holidays with you from here on out. To Katie: when I first mentioned this endeavour all those years ago, all I heard from you then, and ever since, was how we'd make it work. Only with your support has this journey been possible.

Finally, I would like to thank the internal and external examiners for their careful reading of the manuscript and for the thoughtful and constructive feedback provided during the examination process. Their comments and suggestions have helped to improve both the clarity and presentation of the thesis.

Abstract

Electricity markets across many jurisdictions have undergone substantial restructuring in recent decades. In some regions, this has been accompanied by significant shifts in the electricity generation capacity landscape, for example the large-scale deployment of variable renewable generation or the closure of nuclear and coal assets in parts of Europe. These developments have contributed to increasingly dynamic trading environments, particularly in short-term electricity markets. It is within this context that our research is situated.

Our work focuses on two distinct but related areas. The first involves forecasting the evolution of a short-term electricity market using granular participant-level datasets. The second presents empirical analyses within the same market environment, complementing and contextualising the forecasting strand. In the forecasting component, which is conducted in a Day-Ahead market setting, we apply statistical, machine learning and deep learning techniques to predict individual participant commercial behaviours; subsequently we develop an aggregate view of electricity demand and supply. This granular approach addresses challenges associated with bottom-up (i.e. fundamental) models, namely the reliance on and validity of expert-driven input assumptions and the associated model calibration efforts. As part of this work we also develop an explainability framework to interpret differences between forecast and actual outcomes in addition to algorithms to manage high-dimensional participant order data. Our findings show that while the bottom-up approach effectively captures the complexity and evolution of participant behaviour, and by extension the market structure, its price forecasting accuracy is lower than that of top-down models in the literature (which typically focus solely on price). Using the explainability framework, we trace this underperformance to a cohort of speculative units exhibiting highly dynamic and difficult-to-forecast behaviours. Building on this, the second strand of our research explores how speculator participation in the market has evolved over time. We analyse factors such as marginality rate (the proportion of time a speculator sets the marginal price) and overall market share. We observe a clear increase in speculator participation over the study period and find that, relative to their market share, these units are marginal in a large number of trading periods. We also present an intuitive metric for analysing Day-Ahead market price inertia (the market price's ability to withstand small changes in demand or supply). We view the latter analysis as having conceptually similarities to order book analysis in continuous trading environments.

The relevance of the research is as follows. First, the modelling framework is, to our knowledge, the first of its kind applied at this scale offering a novel contri-

bution to the literature and a new perspective on how participant-level data can be used (along with its limitations) in electricity market analysis. This is important given the continued use of bottom-up models by both regulatory authorities and market participants. Second, the empirical analysis provides concrete metrics on both speculator activity within a European market context and market price inertia further highlighting the susceptibility of market prices to jumps given small changes in demand or supply.

The research is informed by both practical industry experience and engagement with the broader academic literature. As such, it is written for two audiences: practitioners, who may benefit from its applied insights, and academics, to whom it offers methodological contributions and new perspectives.

1 Introduction

1.1 The Commodity Industry

The commodity industry encompasses the sourcing, production, transportation, and delivery of raw materials, including agricultural products, metals, and energy commodities[1]. Given its wide reach, it is clear that it has an impact and influence on our everyday lives. Informally, those in the commodity industry often categorise it into three sectors:

- *Upstream* – Exploration, extraction, and initial processing of raw commodities.
- *Midstream* – Transportation, storage, and processing of commodities.
- *Downstream* – Marketing and distribution of commodities to end users, including retail and industrial consumers.

Across all three sectors, wholesale trading takes place, whether to hedge risk, optimise supply chains, or capitalise on market opportunities. Companies operate within one or multiple sectors, specialising in either a single commodity or engaging across multiple commodities. While the above classification covers a broad range of activities, the nuances and complexities of the industry become clearer when considering:

- **Commodity-specific differences** – The end-to-end value chain, from production to delivery, as well as the key market participants, varies across different commodities.
- **Regional differences** – Regulations, market structures, and fundamental drivers such as demand and supply create distinct dynamics for the same commodity across different regions.

The energy industry, encompassing electricity, oil, natural gas, and related products, represents a particularly significant component of the commodity sector in its own right. Although this has not always been the case historically (and it varies by energy product), the sector today is characterised by high trading volumes and substantial regulatory and operational complexity [1, 3, 4]. As such, energy markets are often analysed independently within the broader study of commodity markets.

From a European energy market perspective, there has been a clear trend toward regulatory alignment, particularly in electricity and gas markets over recent decades (section 1.3). However, despite this harmonisation, regional supply

[1] For a comprehensive introduction to commodity markets, see [1, 2].

and demand variations continue to drive distinct market dynamics. Even within the same geographic region, supply and demand conditions can vary significantly across different market timeframes (i.e. spot versus forward markets). The nuances and complexities of these distinct market environments make energy markets, particularly electricity markets, highly dynamic and they have proven to be a rich area for academic research. Our research is positioned within this evolving landscape.

1.2 Concepts and Terminology

Before providing an overview of European electricity markets (the domain in which this research is situated), this subsection introduces concepts and terminology used throughout the thesis to aid accessibility for readers without a background in energy or markets.

- *Participant*: an entity that submits bids and offers in a market. In short-term electricity markets, a company may operate multiple participants within the same market, each representing a distinct component of its portfolio (e.g. demand, renewable generation, thermal generation etc).

- *Bid, Offer (Ask)*: A *bid* specifies the price and quantity pairs at which a participant is willing to buy, while an *offer* (or *ask*) specifies the price and quantity pairs at which a participant is willing to sell. A bid or offer curve represents these pairs as a schedule of quantities indexed by price, typically visualised with quantity on the x-axis and price on the y-axis.

- *Order, Trade (Position)*: an *order* is a bid or offer indicating a willingness to buy or sell a given quantity at a specified price (or better). When an order is executed, it results in a *trade*. The term *position* is sometimes used interchangeably with *trade*.

- *Delivery period*: a discrete time interval over which electricity is delivered and settled, typically of fixed duration (e.g. one hour). Orders and trades may relate to a single delivery period or to multiple delivery periods. The term *trading period* is sometimes used to refer to the same concept.

- *Spot and forward markets*: *spot* markets are those in which delivery of the underlying commodity, or settlement of the trade, occurs over a short time horizon, typically within the next few days. *Forward* markets are those in which delivery or settlement occurs at a later date beyond this short-term horizon.

- *Balancing markets*: electricity markets operated close to real time in which the system operator procures adjustments to generation or demand to ensure that supply and demand remain in balance. In US jurisdictions, these are often referred to as *Real-Time* markets.

- *Day-Ahead* markets: an important construct in electricity markets, and one of the focus areas of the thesis. This is the market where participants submit orders on day D for each hour in delivery day D+1.

- *Marginal price*: the price associated with the last unit of supply required to meet demand.

- *Supply stack*: the ordering of available supply offers by price, indicating the quantities participants are willing to sell at different price levels. It is often referred to as the generation stack or merit order, with supply arranged from lowest to highest marginal cost.

- *Market-clearing prices, auction clearing*: In European Day-Ahead markets, the EUPHEMIA algorithm [5] determines accepted bids and offers and sets a single price for each delivery period and bidding zone. These resulting *zonal market-clearing prices* correspond to the prices at which accepted supply and demand bids are matched, while the matching process itself is referred to as *auction clearing*.

- *Negative electricity prices*: market outcomes in which the wholesale price of electricity is below zero, implying that generators pay to remain online or that consumers are effectively paid to consume electricity. Such outcomes can occur in short-term electricity markets.

- *Inelastic demand*: demand that responds weakly (or not at all) to price changes. In short-term electricity markets, a substantial share of demand is effectively inelastic.

- *European Energy Crisis*: a period of elevated prices and extreme volatility in European gas and electricity markets, beginning in mid 2021 and intensifying following Russia's invasion of Ukraine in February 2022 [6]. The crisis was driven by declining gas storage levels including reduced Russian gas flows to Europe resulting in heightened security of supply concerns. By mid to late 2023, market prices began to stabilise and return toward more typical levels.

- *Internal Energy Market*: a European Union policy initiative, emerging in the 1990s, aimed at integrating national gas and electricity markets into a single competitive market through liberalisation, harmonised regulation, and cross-border market integration [3].

- *Self-dispatch*: a market paradigm in which participants determine their own production or consumption schedules and submit these schedules to the system operator (subject to network and operational constraints). This contrasts with centralised dispatch where a market-clearing algorithm (or similar mechanism) determines dispatch based on participant bids and offers.

- *Algorithmic trading*: trading activity in which orders are generated, submitted, and managed automatically by computer algorithms based on predefined rules or models.

- *MW and MWh*: Units of measurement commonly used in electricity markets. A megawatt (MW) is a unit of power, representing the instantaneous rate of electricity generation or consumption. A megawatt-hour (MWh) is a unit of energy, representing one megawatt of power sustained over one hour.

- *Multivariate and univariate*: while these terms conventionally refer to models with multiple versus a single explanatory variable, some electricity price forecasting studies use them to distinguish between hour specific models (*multivariate*) and a single model applied across all hours (*univariate*).

- *Feedforward Neural Network (FFNN)*: A network architecture in which the input is transformed into the model output through a sequence of layer-wise transformations (i.e. the forward pass). This contrasts with neural network architectures that incorporate recurrence or state into the output computation, such as recurrent neural networks.

Simplified Illustrative Examples of Short-Term Electricity Market Trading The following simplified examples[2] are intended to reinforce the terminology introduced above. In presenting these examples, we introduce two additional concepts: Intraday auctions, which occur closer to the delivery period and are similar in structure to the Day-Ahead auction (see section 3.1), and speculators, who do not take physical positions but instead profit or incur losses from price differences across markets (see chapter 5).

[2] These examples abstract from several operational details, including the full range of order types, half-hourly delivery granularity, continuous trading, and settlement processes. They are intended solely to illustrate core concepts and trading logic.

For this hypothetical example, the following timelines inform participant actions. The Day-Ahead (DA) auction for each delivery period in day $D{+}1$ closes[3] at approximately 11am on day D, with auction results communicated roughly one hour later. Similarly, gate closure for the first intraday auction for delivery on day $D{+}1$ occurs at approximately 5:30pm on day D, with results published shortly thereafter. Within this simplified setting, participants may be active in one or both auctions, and the orders they submit are informed by forecasts, operational constraints, and commercial objectives. Consider the following participants:

- At approximately 9am on day D, a **demand participant**, representing a retail customer portfolio, produces a forecast of electricity consumption for each delivery period in day $D{+}1$. Based on this forecast, prior to the 11am gate closure on day D, the participant submits 24 simple orders (one per delivery period) to the Day-Ahead auction. Each order is a bid with quantity equal to forecast demand and price set at the market cap (e.g. €3000/MWh effectively ensuring execution). At approximately 12pm on day D, the participant is informed that all day $D{+}1$ orders have been fully matched at the market-clearing prices. Later in day D, the demand forecast for day $D{+}1$ is updated; as the demand forecast remains unchanged, no further trading is required as the participant's portfolio remains hedged (i.e. procured supply equals expected demand).

- Prior to the 11am gate closure on day D, a **thermal generation participant**, operating a gas-fired power plant, forecasts both its available capacity and marginal running costs for day $D{+}1$ based on prevailing gas and carbon input prices. Based on these estimates the participant submits 24 simple sell orders, with quantities reflecting available capacity and prices set at estimated marginal cost plus a commercial markup (e.g. 10%). At approximately 12pm on day D, the generator is informed that its orders have been fully matched in the Day-Ahead auction. As a result, the plant is scheduled to run at its available output and has no remaining capacity to offer in the subsequent Intraday auction.

- A **speculator participant**, observing that Day-Ahead prices have recently exceeded corresponding Intraday prices, expects this pattern to persist. On this basis, the participant submits 24 simple sell orders to the Day-Ahead auction, each for 50 MW at a price of −€500/MWh, and 24 simple buy orders to the first Intraday auction, each for 50 MW at a price of €3000/MWh. The Day-Ahead orders are submitted prior to the 11am gate closure on day

[3] i.e. the time by which participants must submit their orders; submissions after gate closure are invalid.

D, with results announced at approximately 12pm, while the Intraday orders are submitted before the 5:30pm gate closure and cleared shortly thereafter. In this scenario, given that the speculator is fully matched in both auctions, its payoff is given by the price spread between the Day-Ahead and Intraday markets for each delivery period.

1.3 European Electricity Markets

Below, we provide a high-level overview of the evolution of European wholesale electricity markets and the key factors shaping their development. But first, why have electricity markets historically attracted such sustained academic and industry attention and why do they remain so relevant today? As electricity markets liberalised from the 1990s onward and expanded in both scale and complexity, they presented new forecasting challenges related to generation, demand, and market prices, alongside increased concerns around market power and competition [7, 8, 9, 10, 11]. These factors have long driven interest from researchers and practitioners alike. More recently, the European energy crisis, marked by extreme volatility in wholesale electricity and gas prices, has brought these markets into sharper public and political focus [6, 12, 13]. In this context, analyses that deepen our understanding of electricity market dynamics and their underlying complexities is both timely and important.

Factors that have helped influence the structure of European wholesale electricity markets include market deregulation and privatisation, the development of the *Internal Energy Market (IEM)* and the EU 2020 and 2030 energy targets.

Market deregulation & privatisation In an electricity and gas context market deregulation largely refers to the introduction of competition in the generation and supply sectors. Historically, many countries operated under a monopoly model where a single state-owned entity was responsible for the sourcing or generating of energy and the subsequent supply to end users at fixed prices or regulated tariffs. With deregulation, these sectors were gradually opened to competition, this was often accompanied by the divestiture of state-owned energy assets [14, 15, 3]. This transition took place in Britain during the 1980s and 1990s for example while Ireland underwent a similar process in the early 2000s.

Internal Energy Market The Internal Energy Market is a fundamental pillar of the European energy sector established through legislative packages starting in the 1990s and evolving over time to strengthen market integration and competition [16]. A key development in European electricity market integration is the *European Target Model*. As outlined in [17], *'the European Target model is the* *blueprint for market integration across the internal energy market. Key features of the Target Model are*

- *A common price coupling algorithm for scheduling day-ahead markets and determining flows between geographic regions.*

- *Energy trading within regions and across borders close to real time...'.*

Day-ahead (DA) markets are the primary focus of our research. In a European context, these markets operate as a two-way auction conducted at a set time on day D, where participants submit bids and offers for each hourly delivery period on day D+1. The auction is cleared using a market coupling algorithm that maximises total economic surplus across bidding zones, determining market clearing prices and cross border electricity flows subject to transmission constraints. Figure 1 shows countries that participate in DA market price coupling in Europe. For additional context on market coupling (in European and other settings) and how it has evolved over time see [18].

Figure 1: European Countries participating in Single Day-Ahead Coupling (in blue). As of 1st January 2021 Great Britain is no longer coupled in the Day-Ahead market. Image source: ENTSO-e [19].

2020 and 2030 Targets In 2007 the EU set three key targets for 2020 [20], namely

- '20% *cut in greenhouse gas emissions (from 1990 levels).*

- 20% *of EU energy from renewables.*

- 20% *improvement in energy efficiency'.*

Member states were given flexibility in how they achieved these targets. In Ireland, for example, the electricity sector was expected to play a leading role with a target of 40% of electricity demand being met by renewable energy sources by 2020 rising to 80% by 2030 [21]. Across Ireland and other European electricity markets, these policies have led to significant changes including:

- A substantial expansion of variable renewable electricity generation capacity [22].
- The closure or planned closure of a significant portion of coal fired electricity generation assets [23, 24].
- A noticeable rise in instances of negative wholesale electricity prices [25].

The combined impact of these factors (i.e. market deregulation and privatisation, the internal energy market, increasing build out of renewable generation assets) has made the environment in which electricity market participants operate today far more dynamic than in previous decades. Our research is conducted within this evolving context.

1.4 Thesis Focus and Background

The primary focus of this thesis is to investigate whether participant-level order data, specifically the granular bids and offers submitted by participants to Day-Ahead electricity auctions, can be leveraged to gain deeper insights into the electricity market price formation process. Bottom-up[4] electricity market models commonly encountered in both the literature and industry often operate at a coarse level of granularity, for example by omitting financial trader participants or treating demand as a single inelastic cohort (i.e. demand that is assumed not to respond to price changes). In addition, they often rely on expert-led input assumptions. In contrast, this research explores the feasibility of modelling all participants in a Day-Ahead electricity market and, in doing so, forecasting demand, supply, and price dynamics with sufficient accuracy to offer both explanatory and predictive value. The approach aims to address practical challenges associated with bottom-up modelling particularly around automation and calibration (i.e. the need to manually configure model inputs and iteratively adjust them until the outputs are reflective of observed market outcomes). A secondary focus of the thesis is a historical empirical analysis of a Day-Ahead electricity market, aimed at developing a deeper understanding of both the behaviours of a particularly dynamic participant cohort and price inertia levels. The latter is conceptually similar to liquidity analysis in continuous trading environments.

This research focus aligned with over a decade of professional experience in energy markets. Since the late 2000s, I have held roles spanning renewable power

[4]Discussed in chapter 2.

purchase agreement origination, portfolio hedging for electricity supply and demand, and competitor and market analysis. More recently, I have worked in advisory and digital transformation roles including assisting with the design and build of a Python based electricity market modelling platform and involvement in the design and specification of an algorithmic trading platform. These roles surfaced several industry trends, such as an awareness of:

- Increasing importance of algorithmic trading in commodity markets [26].
- Short-term electricity markets remain prone to price spikes and in some cases the participants influencing a spike can be identified (e.g. [27, 28]).
- Bottom-up modelling approaches continue to be used in industry.
- Increasing availability of public electricity market datasets.

Recognising these developments, and appreciating the relevance of the research questions, I saw an opportunity to utilise my industry and domain experience within a rigorous research programme. This also allowed me to pursue my growing interest in statistical computing and machine learning while contributing to a domain of practical significance.

The research officially began in October 2018, with an initial focus on two fronts: conducting a comprehensive literature review (chapter 2) to ground the work in current academic and industry thinking and developing robust data pipelines (chapter 3) to support the forecasting and empirical analysis that followed.

1.5 Publications

The following papers form the core of this thesis:

- Chapter 4 is based on *A New Approach to Fundamental Electricity Market Modelling: Exploring Market Dynamics and Speculator Influence* [29], published in the *Energy and AI* journal in December 2025.
- Chapter 5 is based on *Speculators and Price Inertia in a Day-Ahead Electricity Market: An Irish Case Study* [30], published in the *Energies* journal in September 2025.

Beyond the thesis-related publications, research collaborations during the PhD produced additional commentaries and academic articles on short-term electricity markets. While not part of the thesis, these contributions include:

- *Why Gas and Electricity Prices in Ireland Are So Closely Linked*, 2022, RTÉ Brainstorm [31].
- *Electricity Price Forecasting in the Irish Balancing Market*, 2024, Energy Strategy Reviews [32].
- *Optimising Quantile-Based Trading Strategies in Electricity Arbitrage*, 2025, Energy and AI [33].

1.6 Thesis Structure

The remainder of this thesis is structured as follows. In Chapter 2, we summarise the relevant literature and provide a brief overview of the research in chapters 4 and 5 discussing any overlaps with, and gaps in, existing research. In chapter 3 we describe the data pipelines and datasets used in the analysis, we also present a concise overview of the main modelling techniques employed in the research. In chapter 4 we use granular participant datasets in conjunction with statistical, machine learning, and neural network modelling techniques to forecast the evolution of a DA electricity market. We also provide custom metrics which provide an explainability framework to understand forecast performance, in addition we benchmark our approach against publicly available (top-down) models. Motivated in part by findings from chapter 4, in chapter 5 we present an empirical analysis of the short-term electricity market setting focusing on two key aspects: the presence of financial traders (who we refer to as *speculators*) and a market price inertia analysis. In chapter 6 we conclude.

2 Contextualising the Domain

2.1 Introduction

This chapter begins by discussing a selection of *Electricity Price Forecasting (EPF)* review publications followed by an overview of seminal and innovative contributions to the field. We then discuss broader trends within the EPF domain before turning to additional literature that while not in the EPF domain is tangentially related to the empirical research presented in chapter 5. The chapter concludes by situating our research questions and contributions within the relevant research domains. Before proceeding, we note that forecast horizons in the EPF domain often lack standardised terminology or formal definitions [34]. For context, some common forecast horizons within the EPF domain and examples of interested parties are the following:

- Intraday and DA: Market participants use price forecasts for portfolio optimisation.
- Short-Term (Balance of Week): Longer duration storage asset operators may consider forecasts beyond intraday and DA to decide when to act as a source of demand or supply.
- Medium-Term (1 month to 1 year): Generators often consider forecasts when planning maintenance related outages.
- Long-Term (1 to 10 years): Investors use long-term price forecasts when assessing whether to dispose of, acquire or develop generation assets.

In this thesis our primary focus is on the DA market and hence we largely limit our EPF literature review to studies that deal with this timeframe.

2.2 EPF: Review Literature

In [11] Weron provides a classification of EPF models. It categorises these models into five groups: *multi-agent* models, which simulate market interactions among heterogeneous agents; *fundamental* methods, which model price dynamics based on physical and economic factors; *reduced-form* models which focus on the statistical properties of electricity prices for risk management and derivatives pricing (often using stochastic differential equations, SDEs); *statistical* approaches, which apply econometric techniques and methods from load forecasting for example; and *computational intelligence* techniques which leverage artificial intelligence, non-parametric and non-linear statistical techniques. While the above high-level summary may be somewhat ambiguous, the references Weron provides to notable studies within each category offer useful clarification and context. Weron also points out how many EPF models in the literature adopt hybrid approaches that combine elements from multiple categories, a characteristic also found in a selection of papers discussed later in this chapter, as well as in our own research. Additional contributions include a discussion on forecast error metrics and reflections on the future trajectory of EPF research. Albeit tangential to our research [35] provides an insightful overview of probabilistic forecasting in the EPF domain. It presents a detailed discussion of a global forecasting competition[5], which included an EPF component. The paper also provides an in-depth analysis of forecast error metrics and statistical tests.

In [37] Ventosa et al. present an overview of multi-agent approaches in electricity price forecasting identifying three primary modelling trends: optimisation models, simulation models, and equilibrium[6] models. Optimisation models focus on profit maximisation, with two main types: those treating price as an exogenous variable and those modelling price as a function of firm-specific supply. These models, often based on linear programming techniques, effectively account for operational constraints but face calibration challenges in competitive markets where strategies evolve dynamically and demand is elastic. Simulation models offer an alternative. They represent strategic decision making through sequential rules, governing actions like generation scheduling or offer curve adjustments in response to market conditions. While they allow for greater flexibility, their realism depends on the theoretical justification of underlying agent assumptions, [37] highlights their wide applicability across different market contexts and forecast horizons.

Additional relevant review-oriented literature includes [38] and [39], both of which also incorporate analytical contributions. In [38] the authors provide a review of deep learning and statistical forecasting methods, in addition to emphasising the need for greater rigour and standardisation in EPF research. A key contribution is the provision of open-source datasets and *'state of the art'* forecasting models (one statistical and one deep learning) enabling researchers to benchmark competing forecasting approaches. In [39] Lago et al. present a rigorous empirical comparison of statistical and computational intelligence models in EPF addressing their relative performance. Using Belgian day-ahead market

[5] GEFCom2014, see [36].

[6] Equilibrium models attempt to represent overall market dynamics, [37] outlines two key types: *Cournot competition* models, where firms compete on quantity, and *Supply Function Equilibrium (SFE)* models, where they compete via offer curves. These models rely on a formal mathematical equilibrium, which according to [37] makes them difficult to solve and often too rigid to capture the dynamic strategies of market participants.

data, the authors find that deep learning models generally outperform statistical approaches, though several statistical models still demonstrate strong predictive capabilities.

Summary: While Weron and Ventosa et al. provide valuable references and broader context for the EPF domain, the most directly relevant studies for our research are [38] and [39]. These works offer insights into more recent EPF approaches, specifically a class of methods we will shortly refer to as top-down, and we draw on similar techniques to benchmark the electricity market modelling framework developed in chapter 4.

2.3 EPF: Bottom-Up and Top-Down

Before continuing with the EPF literature review, we introduce a simplified taxonomy that serves as a useful guide for the discussion that follows. Rather than adopting the more detailed classification proposed by Weron [11], we distinguish between two broad methodological approaches: *bottom-up* and *top-down*. In this categorisation, an approach is considered bottom-up if it focuses on modelling the components that underlie or drive the price formation process. In contrast, top-down approaches forecast electricity prices directly, without explicitly modelling the underlying mechanisms. As will be seen in following subsections, our bottom-up category can be further subdivided in which case the subdivisions largely align with a subset of Weron's original classifications. Table 1 presents a comparison of the two taxonomies, with rows having conceptually similar approaches.

Table 1: EPF taxonomies.

Simplified	**Weron**	**Description**
• Top-Down	• Statistical approaches • Computational intelligence methods • Reduced-form models	Price is modelled directly; includes statistical, AI, and SDE-based models.
• Bottom-Up: – Fundamental – RL & agent-based – Hybrid	• Fundamental methods • Multi-agent models	Models physical or economic drivers of price formation.

2.3.1 Bottom-Up

We start by discussing bottom-up EPF studies grouped according to the modelling technique employed. The first group consists of reinforcement learning (RL) [40] and agent-based models which simulate market dynamics and price formation using multi-agent systems. The second consists of fundamental models which rely on cost driven principles like the *merit order*[7] mechanism and often use linear programming (LP) or mixed-integer programming (MIP) for unit commitment, generation dispatch, and price determination. The third category comprises hybrid approaches which blend fundamental and top-down modelling techniques.

RL & Agent-Based Approaches In [42], the authors develop an agent-based simulation model of the England and Wales electricity pool to assess proposed market design changes. Their bottom-up approach represents each generating firm as an autonomous adaptive agent at the power plant level. These agents adjust their bidding strategies based on plant constraints, market demand, and clearing mechanisms, using what the authors describe as a naive reinforcement learning algorithm to learn and optimise strategies for maximising profit and market share over time. Similarly, [43] uses an agent-based simulation model to assess the impact of asset divestiture on market power in the England and Wales electricity market. The authors propose an alternative to the Hirschman-Herfindahl Index (HHI)[8] and validate their modelling framework by showing that the simulation model closely matches results from a supply function equilibrium model[9]. They suggest that the validation exercise provides support for the use of simulation approaches to analyse complex market conditions that may be intractable for analytical methods. More recently [44] introduces a scalable multi-agent deep reinforcement learning model to simulate the German wholesale electricity market. The paper presents a helpful overview of different reinforcement learning paradigms and their suitability in electricity market modelling contexts. As with the studies referenced above, the approach is proposed as a tool to assess the impact of changes in market regulations, generation fleet composition, or strategic behaviours on market dynamics and price outcomes.

A general observation on the above contributions is that they tend to emphasise relative price over absolute price levels with the focus instead being on how market dynamics (including price) and participant behaviours evolve under different market arrangements and trading strategies. In addition, as noted earlier by [37] these approaches are highly flexible and hence have a wide applicability

[7]Conceptually in an electricity market context the merit order is an arrangement of generation resources in terms of their marginal cost from least expensive to most expensive. As per [41] *'pricing in power markets is explained by the merit order model. This model assumes that the spot prices at electricity exchanges are based on the marginal generation costs of the last power plant that is required to cover the demand'*.

[8]The HHI is an approach to measuring market concentration within an industry; the paper argues that its usage in this context is misleading due to demand inelasticity, as a result they use the market price premium over short-run marginal cost as a proxy.

[9]In a simplified setting with symmetric firms and continuous supply functions.

but rely on valid/realistic agent assumptions.

Fundamental Approaches Fundamental models are widely applied in the broader electricity literature (e.g. [45, 46, 47]); however, as noted by [48], they are less prevalent in short-term EPF. As previously noted, these models aim to replicate the merit order by modelling the physical and economic structure of the electricity system. In [49] the authors use a fundamental model to forecast DA and Intraday prices in the German wholesale electricity market emphasising its ability to account for non-linearities in the supply stack. The overall forecast performance, which is positive in the main, indicates that forecasted prices tend to be flatter and less volatile than actual prices. To explore these discrepancies, they employ a regression model, identifying contributing factors such as avoided start-up costs, market conditions, and trading behaviour. Building on [49], [50] develop a parsimonious fundamental model to model German wholesale Day-Ahead prices (in addition to other output variables). While their model effectively captures overall trends, it also produces forecasts that are flatter (i.e. less peaky) and less volatile than observed prices. Furthermore, they also point out that fundamental models do not account for effects, such as strategic bidding, underscoring a key limitation in their ability to fully capture market dynamics.

The recently published preprint by Ghelasi and Ziel, [48], is a novel take on fundamental models. While we discuss it within the context of fundamental modelling approaches, it endeavours to present a framework that integrates the strengths[10] of what it terms *'data-driven'* models (statistical and econometric approaches that learn from historical patterns) and fundamental models. As such, it could also be discussed with the next group of bottom-up models, hybrid models[11]. The model introduces a mechanism to partially or fully automate the calibration of day-ahead fundamental models and *'outperforms both classical fundamental and purely data-driven models'* in price forecast accuracy for the German wholesale electricity market. While the paper appears to be in draft form and does not explicitly address this, it may be the case that its data-driven component implicitly[12], or directly[13], helps mitigate the challenges fundamental models face in incorporating strategic bidding into the forecasting framework. A notable assumption in the framework is that of inelastic demand[14].

[10]Suggested strengths of data-driven models being *'they effectively capture the high volatility of power prices and run efficiently'*, in contrast strengths for fundamental approaches include *'interpretability of fundamental models and their ability to handle regime switches'.*

[11]Since the approach *'retains the interpretability of fundamental models, offering insights into marginal technologies, fuel switches, and dispatch patterns'*, we include it alongside fundamental models as these outputs are typically associated with such approaches.

[12]i.e. via the models structural parameters.

[13]i.e. via modelling the generation fleets efficiency and behaviours

[14]i.e. *'...the price is determined by the intersection of the power plant supply curve and*

Hybrid Approaches References [41, 51, 52] present bottom-up modelling approaches that forecast supply and demand curves, and hence electricity market prices, using statistical methods. Given the use of a statistical approach within a bottom-up context, we classify these as hybrid bottom-up approaches. In [41] Liebl introduces a functional data perspective on modelling and forecasting electricity spot prices that implicitly accounts for the merit order model. The approach follows a two-step process: first, cubic splines are fit to price and demand observations to construct smooth daily price-demand functions; second, a set of basis functions are derived to approximate these daily functions, capturing common underlying structures in price formation. In [51] Ziel and Steinert note that while other studies model demand and supply curves and predict prices based on their intersection, their approach takes an additional step by calibrating the model to publicly available exchange data. Using historical data from the EPEX spot market of Germany and Austria they apply an autoregressive LASSO model (section 3.4) to forecast the evolution of demand and supply curves. The model demonstrates a strong forecasting performance and is also capable of producing probabilistic forecasts. In [52], there is significant overlap with [51], with the main difference being that machine learning techniques are used instead of statistical/econometric forecasting methods.

We note that all three references effectively capture key stylised characteristics of electricity market prices with price forecast accuracy being a primary focus. The ability to generate probabilistic forecasts, as demonstrated in [51], is also an attractive feature. However, what appears to be absent from these novel approaches is an explainability framework that can help analyse and interpret forecast errors. Additionally, a critique raised by [48] in relation to data-driven models is likely applicable here as well, namely that such models are *'highly sensitive to the historical data used for training and can fail under regime changes*[15]*, even if the market mechanism remains unchanged'*. This limitation arises when models are trained on historical data that predominantly reflect a prior regime, resulting in predictions that may not generalise well when market conditions change. We would suggest that this limitation is also likely to apply to the majority of EPF approaches referenced in this thesis and indeed in our own research.

Summary: RL and agent-based bottom-up models offer flexibility but often focus on exploring hypothetical market evolution rather than forecasting. Fundamental and hybrid bottom-up approaches are more closely aligned with our work in chapter 4. In that chapter we address common challenges in fundamen-

inelastic demand'.

[15] Here a regime change has the broad meaning of a shift in underlying market fundamentals or dynamics that results in materially different market outcomes.

tal models (e.g. input assumptions) and examine the impact of greater forecast granularity (of the underlying components) in addition to introducing an explainability framework, these areas are underexplored in existing studies.

2.3.2 Top-Down

In [53], Zareipour presents a chapter on EPF in the Ontario market. The approach primarily employs statistical methods such as *autoregressive integrated moving average (ARIMA)* but also briefly examines a *feedforward neural network* (section 3.4). The chapter highlights key stylised characteristics of electricity prices, including frequent price spikes, daily/weekly seasonality, and non-stationarity. Another contribution is its discussion of EPF model development, covering model selection, feature selection, data transformation, model fitting, and diagnostics. The analysis of exogenous and endogenous variables is also noteworthy. A limitation of the study is its small test set, though this issue is addressed in later EPF studies and again in our own research.

In [54] a *Self-Organizing Map (SOM)*[16] is used to cluster similar observations before applying *Support Vector Machines (SVMs)*. The intuition behind the SOM filtering step is an awarness of a lack of stationarity in the underlying data with the expectation that in applying this step it will improve the overall forecast accuracy relative to stand-alone SVM models. While approaches that apply a filtering or transformation step prior to forecasting are common in the EPF literature, the use of a SOM in this initial stage remains relatively rare.

The study by Keles et al. [56] is a significant contribution in that it challenges earlier claims of neural network underperformance in EPF by providing a more thorough evaluation. It improves forecasting outcomes by including additional explanatory variables and also applying feature selection techniques. It also examines neural network architectures testing different layer depths, node counts, activation functions, learning rates, momentum settings, and calibration windows. They emphasise that no single neural network configuration is universally optimal and point out that extensive experimentation is often required. The paper also highlights limitations including difficulty in forecasting price spikes and challenges in interpreting network parameters.

In a Turkish electricity market setting [57] compares neural networks and statistical time series models. The study explores shallow and deep *recurrent neural networks (RNNs)* for day-ahead price forecasting, specifically *Long Short-Term Memory (LSTM)* and *Gated Recurrent Units (GRU)*. The results show that deeper architectures improve forecast accuracy, with deep LSTM and GRU

[16] An *unsupervised* clustering technique, see [55].

models outperforming both statistical models and shallow neural networks. The paper also finds that deep LSTM and GRU perform better than deep feedforward networks in this market context.

Other relevant EPF research includes Uniejewski et al. [58] where they review various shrinkage methods[17] and provide models which serve as the basis for one of the two publicly available models in the EPF toolbox accompanying [38]. It also offers one of the more comprehensive analyses of explanatory variables in the EPF literature to date. In [59] Ziel and Weron distinguish between *univariate* and *multivariate* modelling frameworks in EPF. The univariate approach uses a single model to forecast the entire horizon while the multivariate approach employs separate models for different delivery periods. They find no clear evidence that one approach consistently outperforms the other. As will be seen in chapter 4, our day-ahead electricity market framework adopts a univariate approach.

Summary: For the most part, EPF publications over the past decade have utilised top-down modelling approaches with neural networks and statistical methods showing strong performance, the former often outperforming the latter. Some studies also explore shrinkage techniques and the role of explanatory variables. While some of these findings inform aspects of our bottom-up modelling in chapter 4, the explainability of forecast performance, specifically identifying the causes of deviations between forecasts and realised outcomes, as in much of the existing bottom-up literature, remains largely unaddressed in this strand.

2.4 EPF: Common Themes

In this section we highlight some of the recurring trends or topics that appear in the EPF literature, particularly within the top-down EPF papers. The market environment plays a crucial role in EPF modelling as each electricity market has distinct physical infrastructure, demand patterns, and market participant composition. Consequently, a model or approach that works well in one environment may not necessarily transfer, or be as relevant, in another (i.e. the approach may need to be tailored to the market environment or structure under consideration). This need for tailored approaches is referenced in various EPF studies.

Explanatory variables in EPF models typically fall into two categories: endogenous (e.g. price) or a combination of endogenous and exogenous variables. The inclusion of exogenous factors (e.g. forecast system demand, generator outages etc) is influenced by domain expertise, data availability, and market structure. Recent EPF studies show a degree of convergence in the choice of explanatory variables for day-ahead forecasting, with commonly used inputs including

[17] Including *LASSO* and *ELASTIC NET* which are discussed in section 3.4.

forecast system demand, forecast renewable generation, plant availability, and historical day-ahead prices. Additionally, some studies use autocorrelation and partial autocorrelation functions when determining lags for both endogenous and exogenous variables. Given the potentially large number of explanatory variables, some studies apply feature selection techniques to improve forecasting accuracy by removing redundant explanatory variables. These feature selection techniques fall into two categories: *filter* methods which operate independently of the model fitting process (e.g., *mutual information* and *k-nearest neighbours* in [56]), and *wrapper* methods which incorporate feature selection within model training. An example of the latter is the LASSO regression used in [58] which performs implicit variable selection during model fitting.

Data preprocessing is a common step in EPF studies, though its necessity depends on the model choice. For time series models like ARIMA, stationarity is a key assumption, so detrending and differencing are typically applied to the data before model fitting or calibration. Similarly for neural networks feature scaling (dependent on the activation function) is often recommended to improve convergence. In contrast, *random forests (RF)* (section 3.4) do not require feature scaling making preprocessing less critical for such models.

Perhaps as a result of the increasing maturity of the EPF field and the growing availability of large datasets, there has been a sustained shift toward using training, validation, and test sets. This trend may also be influenced by the increasing prevalence of computational intelligence techniques (i.e. one of the EPF model classes referenced by [11]) where such practices are standard. Training and validation sets play a key role in selecting model hyperparameters, while an independent test set helps provide a more realistic assessment of the predictive capability of the calibrated model. A sizeable test set of contiguous observations is preferred as it allows for performance evaluation across potentially varying market conditions. Another key consideration in model fitting is the length of the training data (i.e. the calibration window) and whether the model is recalibrated at each forecast step.

The distinction between univariate and multivariate frameworks was described in section 2.3.2. We simply note that for the more recent deep learning EPF publications, the univariate approach seems to be favoured, outside of this context though, both approaches seem to be equally prevalent.

As expected, error metrics for evaluating forecasting performance are encountered throughout the EPF literature. The most frequently referenced are root mean squared error (RMSE) and mean absolute error (MAE), though other metrics such as mean absolute percentage error (MAPE) and symmetric mean absolute percentage error (sMAPE) also appear. Recent studies increasingly compare model forecast performance relative to that of a naïve model (section 3.4). Additionally some EPF publications utilise statistical tests to distinguish between competing forecasting models.

Summary: Over time, the EPF literature (particularly top-down models) has demonstrated increasing rigour with common practices including the use of training/validation/test splits and evaluation over extended forecast horizons. Several studies also acknowledge the need to adapt models to different market environments. In our modelling framework in chapter 4 we aim to match this level of rigour while contributing additional insights.

2.5 Literature Relating to Chapter 5

For clarity, in this thesis the term *speculator* refers to a financial trader who does not intend to take physical delivery, but instead seeks to profit from price differentials across market stages (while participants with physical positions may also engage in speculative behaviour, they do not match our conceptual definition of speculators). In U.S. short-term electricity markets, this activity is typically conducted via *virtual bidding*, which allows participants to arbitrage differences between Day-Ahead and Real-Time prices without physical delivery. While European electricity markets do not feature an identical instrument, economically equivalent speculative behaviour can arise through trading across sequential short-term markets.

Chapter 5 examines financial traders and market price inertia in a short-term electricity market (the latter is the extent to which prices respond, or fail to respond, to small changes in demand or supply). As these topics lie outside the core EPF domain, this section provides a brief overview of the broader literature for context.

As noted in [60], speculation in commodity futures markets is not a new phenomenon, having been examined in various forms since the late nineteenth century, and becoming the subject of rigorous theoretical analysis in the second half of the twentieth century. A notable finding of this metastudy of approximately 100 related papers is the absence of a clear consensus regarding the effects of speculation in commodity futures markets, which may in part reflect inconsistencies across studies in how speculative activity is defined and measured. From time to time, journalistic commentary such as [61, 27], both relating to short-term electricity markets, draws public attention to speculative activity and may contribute to a negative public perception of its effects on commodity and energy markets, an issue also discussed in a historical context by [60]. More recent academic studies, including [62, 63, 64, 65], provide in-depth analyses of speculative trading in

U.S. short-term electricity markets, where such activity is typically conducted via *virtual bidding* to capture price spreads between Day-Ahead and Real-Time (RT) markets[18]. The evidence, while again not conclusive and with notable caveats[19], indicates that Virtual Bidding helps to narrow the spread whilst also reducing its volatility without any noticeable deterioration in overall system performance in addition to increasing consumer welfare. Some of the main contributions from these studies are the nuanced insights and discussions. For instance, [62] argues that Virtual Bidding's theoretical basis relies on the *'strong assumption that the different stages of organized wholesale markets operate identically'*, an assumption that does not always hold in practice. The authors also highlight that market pricing algorithms are considerably more complex than simple intersections of demand and supply curves. This latter observation, rarely highlighted elsewhere, aligns with our view that findings from some bottom-up studies (section 2.3.1) including our own research need to be interpreted with a degree of caution. Models that assume price formation is solely the result of demand and supply curve intersection ignore underlying complexities such as the existence of more complicated order types. Even when these are incorporated, the full details of the DA market pricing algorithm, particularly in a European context, remain unavailable forcing researchers to rely on approximations. This lack of transparency, combined with complex market dynamics explains some of the challenges in establishing clear cause-and-effect outcomes in short-term electricity market research.

In other research various studies have examined electricity market price volatility, price spikes, and stability (e.g., [66, 67, 68]). For instance, [66] investigates how demand and supply shocks are associated with price jumps in the Nordic short-term electricity market[20]. Their two-stage approach, first identifying shocks and jumps, then exploring correlations, reveals that shocks are more likely to cause price jumps when the market operates near capacity constraints. Similarly, [67] finds that speculative shocks from trading activities contribute significantly

[18] As explained in section 1.2, Real-Time markets in the United States serve a role broadly analogous to balancing markets in Europe in that both resolve supply and demand imbalances close to delivery. However, U.S. Real-Time markets are typically centrally dispatched and cleared using locational marginal pricing, whereas European balancing markets generally rely on balancing energy bids activated by the system operator and are zonal in nature. The empirical methods used in Chapter 5, applied in a European balancing market context, are not market-specific and could in principle be applied to other short-term electricity market settings, including U.S. markets.

[19] For example, [65] show that when financial traders are constrained in their ability to arbitrage price differences between markets, large participants may instead find it more profitable to place bids strategically so as to influence prices and thereby increase the value of related financial positions. Elsewhere [62] provide logical examples which would lead to a narrowing of the DA and RT spread without any overall system benefit.

[20] The study period overlapped with structural market changes associated with market integration.

to forecast error volatility compared to fundamental shocks. In a separate study, [68], in a stability analysis that is rare in the electricity literature, shows that as the market expanded and competition increased electricity prices became more stable and less volatile indicating that the integrated Nordic short-term electricity market is less sensitive to shocks. Though the three studies differ in terms of underlying modelling techniques, all three use daily average spot prices in the analysis. We also note the existence of other literature such as [69, 70, 9] which use more granular electricity price data (e.g. hourly) in the volatility and price spike analysis. Finally, defining market price inertia as the responsiveness of price to small changes in demand or supply, to our knowledge no existing studies directly address market price inertia in short-term (e.g. auction-based) electricity markets.

Summary: This subsection referenced two distinct research strands. The first examined speculators in short-term electricity markets, with the referenced studies focusing on US markets and the impact on DA and RT price spreads. The second explored short-term electricity price volatility. As explained in the following section, these studies have little direct overlap with our empirical analysis in chapter 5 but are included to help frame the research.

2.6 Research in Context

This section provides an overview of the thesis research questions and contributions while also highlighting conceptual and methodological overlaps (and differences) with some of the previously referenced literature. Although some of these aspects are covered in chapters 4 and 5, this section offers a more in-depth discussion to provide additional context.

Chapter 4 This research focuses on two key areas. First, we examine the application of a fundamental modelling approach in a DA electricity market and assess whether it can overcome typical limitations, that is the challenge of incorporating strategic behaviours and reducing or eliminating the reliance on expert input or predefined assumptions. Second, we explore whether an explainability framework can be developed to analyse and interpret forecast performance thereby addressing a common shortcoming in the EPF literature.

Our forecasting methodology is essentially a hybrid bottom-up approach integrating elements of both fundamental and top-down models. We forecast individual participant order data (i.e. the fundamental element), the output is then aggregated to construct forecast demand and supply curves. This enables the prediction of market dynamics, prices, and other output. The participant order forecasts are generated using statistical, machine learning, and deep learning techniques (given that these methods are frequently used in top-down EPF approaches we view this as the top-down element). This removes the need for manual calibration of the fundamental model allowing it to dynamically adapt to evolving market conditions.

By applying top-down techniques to a fundamental model and forecasting the behaviour of all market participants we are essentially pushing the fundamental model to its limits. What impact does this have on forecast accuracy? Can it still effectively capture market dynamics and reliably forecast price movements? These questions are at the core of the study.

The study presented several practical challenges. First, as detailed in chapter 3, acquiring, transforming, and generating the necessary datasets (i.e. both target and explanatory variables) was a substantial undertaking that had to be completed before any exploratory analysis or model design and evaluation could begin. Second, market participants have significant flexibility in how they structure their commercial orders necessitating bespoke order handling algorithms to standardise and transform raw order data into a format suitable for modelling. Finally, the scale of the analysis introduced computational and analytical challenges. With over 400 participants in the study period, we implemented a distinct *walk-forward validation* process for each one over a two year forecast horizon. This process was repeated multiple times to assess the impact of different modelling approaches.

While our research shares some high-level conceptual overlaps with the bottom-up approaches referenced in section 2.3.1, there are significant methodological differences as well as distinct research focuses and contributions. Our approach bears superficial similarities to reinforcement learning and agent-based simulation studies in that we also model individual participants (or agents) and examine market dynamics and outcomes. However, key differences emerge when considering the underlying modelling techniques. Additionally, our approach places greater emphasis on price outcomes rather than relative price movements. The motivation behind the hybrid bottom-up modelling approaches (i.e. [41, 51, 52]), to model and utilise what [51] refers to as the *'true source'* of the price formation process, aligns with our approach. However, differences arise in the underlying data used (i.e., participant order data versus anonymised demand and supply curves) and our additional focus on providing an explainability framework. While we acknowledge some conceptual overlap with the recently published preprint [48], this overlap is confined to the use of more granular modelling of market outcomes beyond aggregate bid and offer curves. The methodological approaches differ significantly; for example, [48] assumes inelastic demand, whereas our approach explicitly incorporates demand bidding, modelling demand at the participant level and allowing bids to vary over time to reflect evolving commercial behaviour rather than being treated as a price-insensitive quantity.

Finally, we are of the view that the applications of this research have the potential to extend beyond the specific market setting examined in the study. While the Irish electricity market setting is notably transparent in its publicly available data, similar forecasting methodologies can be applied in other jurisdictions with comparable DA market data transparency[21]. Additionally, many jurisdictions with balancing market mechanisms exhibit a high degree of commercial order data transparency, making them well-suited for similar modelling approaches. Beyond academic interest, a key practical application of our hybrid bottom-up methodology is its use as a calibration toolset for fundamental modelling approaches.

Chapter 5 This chapter presents two empirical analyses within the same short-term electricity market setting as explored in chapter 4. The first analysis examines the presence and behaviour of financial traders (i.e speculators). While identifying speculators may seem straightforward, market participants' commercial flexibility somewhat complicates this task (e.g. supplier units may engage in speculative trading etc.). The lack of clarity around the definition of and methodology for identifying speculators is something that has been noted in [60]. Once we have identified speculators we analyse their order and execution data over time, examining behavioural patterns, strategic shifts, and profitability at both aggregate and granular levels. We also assess the proportion of trading periods in which they are marginal. Although the influence of speculators on short-term electricity market spreads has been well studied in the USA (section 2.5), an empirical analysis of speculator prevalence and behavioural patterns in a European market setting is lacking.

In the second analysis, we investigate market price inertia and its evolution over time. Although electricity market prices are known for their volatility and susceptibility to sudden jumps, we introduce a simple and intuitive framework to measure how sensitive market prices are to small changes in demand and supply. To the best of our knowledge no literature in the wider electricity domain directly overlaps with our analysis/approach, although we acknowledge that [51] in the process of motivating their EPF model present a graphical illustration of a trading period example with conceptual similarities.

The motivations for the above analyses are twofold. First, prior industry ex-

[21] For instance, [8] highlights the *'ample publicly available data for the Iberian electricity market, which allow us to analyze firms' strategic behavior in sequential markets at high resolution'*.

perience suggests a growing presence of algorithmic and speculative traders, along with continued price spikes in short-term electricity markets. This observation motivated an investigation into whether these trends are indeed present in the study setting of choice. Second, findings from chapter 4, particularly the influence of speculators on forecast accuracy and the sensitivity of prices to small changes in demand and supply prompted a detailed follow-on empirical analysis. These analyses add to the existing literature by revealing subtle complexities in short-term electricity market dynamics and price formation, while also highlighting potential challenges for fundamental models[22].

Finally, as noted in chapter 5, the operation of the Day-Ahead (DA) market in the study setting (i.e. Ireland) changed during the study period following Britain's departure from the European Union. As will be shown in the empirical analysis there are notably distinct outcomes before and after this structural market change.

[22] i.e. the dynamic behaviour of speculators and occasional low price inertia levels may be difficult for such models to capture.

3 Data and Models

In this chapter, we provide an overview of some of the data and modelling techniques encountered in chapters 4 and 5. We begin with an overview of the Irish electricity market as it helps to contextualise the data section. We then describe the data pipeline used to source, persist, and subsequently transform the datasets into time series formats. We also provide a cursory description of the datasets. We conclude the chapter with a brief review of the main modelling techniques employed in the research.

3.1 Irish Electricity Market

The Irish electricity market, officially known as the *Integrated Single Electricity Market (I-SEM)*, underwent a major redesign in October 2018, the purpose of which was largely to ensure it aligned with the European target model discussed in section 1.3 [71]. An overview of market timeframes and routes to market is shown in figure 2, trading can range from financial contracts years or months ahead to physical markets closer to real-time. In the Irish electricity market participants seeking a physical position must trade in either the Day-Ahead, Intraday, Continuous or Balancing markets as self-dispatch is not permitted. In terms of relative importance, that is Day-Ahead vs Intraday vs Continuous quantities, as described below, we note that as of 2025 the Day-Ahead market accounts for approximately 84% of traded volumes while the Intraday market (specifically IDA1) contributes circa 13% [72]. For context on the overall size of the Irish electricity market we note that net electricity consumption in Ireland in 2023 was approximately 31TWh, the corresponding figure in Great Britain, one of Europe's more liquid and well-known electricity markets, was 294TWh [73].

The Irish electricity market offers the following short-term trading routes to market:

- Day-Ahead (DA) market: at 11am on day D participants submit orders to buy or sell electricity for hourly delivery periods in the [11pm D, 10pm D+1] interval. The market coupling algorithm, *EUPHEMIA* [74], takes these inputs, in conjunction with interconnector transmission capacities (and other factors) and determines hourly zonal prices and cross-border energy flows. When interconnector capacity is fully utilised, prices may diverge across zones reflecting differences in local supply vs. demand balances.

- *Intraday (ID) market*: following the clearance of the DA auction, three additional auctions, *IDA1*, *IDA2* and *IDA3* are held. These auctions are similar in structure to the DA auction but differ in two key respects: they occur closer to the delivery time and operate on 30-minute delivery intervals. There is overlap between the delivery horizons of the DA and ID auctions. The IDA1 auction occurs at 17:30 on day D and covers from 23:00 on D to 22:30 on D+1. The IDA2 auction occurs at 08:00 on D+1 and covers 11:00 to 22:30 on D+1, while the IDA3 auction takes place at 14:00 on D+1 and covers 17:00 to 22:30 on D+1.

- Continuous market: this is an important market in other European jurisdictions but in an Irish electricity market context it comprises less than half a percent of traded energy volumes over the study timeframes in chapters 4 and 5, hence it is ignored here.

- Balancing (B) market: one hour prior to delivery the ex-ante[23] trading opportunities cease and the *Transmission System Operator, TSO*, takes over. The TSO compares forecast demand versus forecast generation (including what volumes have been traded in the ex-ante markets) and dispatches on/off units to ensure that demand meets supply. The prices that result from these dispatch decisions are called *settlement imbalance* or balancing market prices.

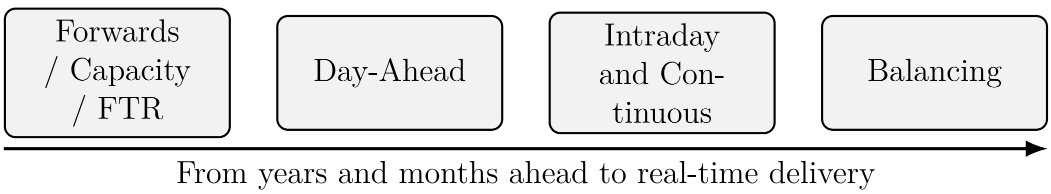

Figure 2: Routes to market available to participants in the Irish electricity market.

Structural Market Change For the timeframes analysed in chapters 4 and 5 the electricity markets in Ireland and Great Britain were connected via two interconnectors. Prior to Great Britain leaving the EU (i.e.*Brexit*), these markets were coupled in DA and IDA1 auctions; post 1st January 2021 the coupling now takes place at the IDA1 and IDA2 auctions [75, 76]. We note that in chapter 5, when presenting the empirical analyses we distinguish between time series data before and after this structural market change.

3.2 Data Pipeline

Figure 3 provides an overview of the data pipeline used to source, persist, and transform datasets. On the left of the flow chart are distinct data sources from

[23] In this context ex-ante means prior to delivery.

which raw data is obtained. Due to the variety of sources, datasets, publication frequencies, and unique characteristics of each (e.g. embedded business logic, datetime conventions etc), a robust pipeline was essential. The pipeline is built with a C# backend and a MySQL database. In the first step, raw data files are retrieved and stored in the database. Selected data is then transformed into time series formats (e.g. section 3.3 page 39). In a final step, which isn't reflected in the flowchart, the time series data is either combined with transformed participant commercial order data for model training and forecasting (chapter 4) or is used with raw data for empirical analysis (chapter 5). The pipeline runs weekly, gradually building a comprehensive historical dataset over time. The majority of the datasets are sourced from *SEMOpx*, which operates the Day-Ahead and Intraday markets in the Irish electricity market, and from *SEMO*, which is involved in market settlement in addition to other operational responsibilities.

From industry experience, it is common for multiple teams to be involved in the development and maintenance of data pipelines similar to, or more complex than, that shown in Figure 3, reflecting the increasing volume, heterogeneity, and operational complexity of datasets encountered in modern energy and commodity markets. When building the pipeline, we endeavoured to follow best practice and to deliver a robust solution capable of handling manifest errors; development was, however, constrained by the available resources (i.e. a single research student). During data acquisition, issues such as missing or corrupted source files were occasionally encountered and addressed through manual validation and no imputation of missing data was performed at any stage of the analysis. Sourcing and orchestrating the data was challenging and time-consuming, largely due to the complexity of certain datasets such as participant order and trade data source files. This is reflective of broader challenges in electricity market research where working with disparate datasets requires a combination of domain knowledge and technical software development skills before meaningful analysis and modelling can begin.

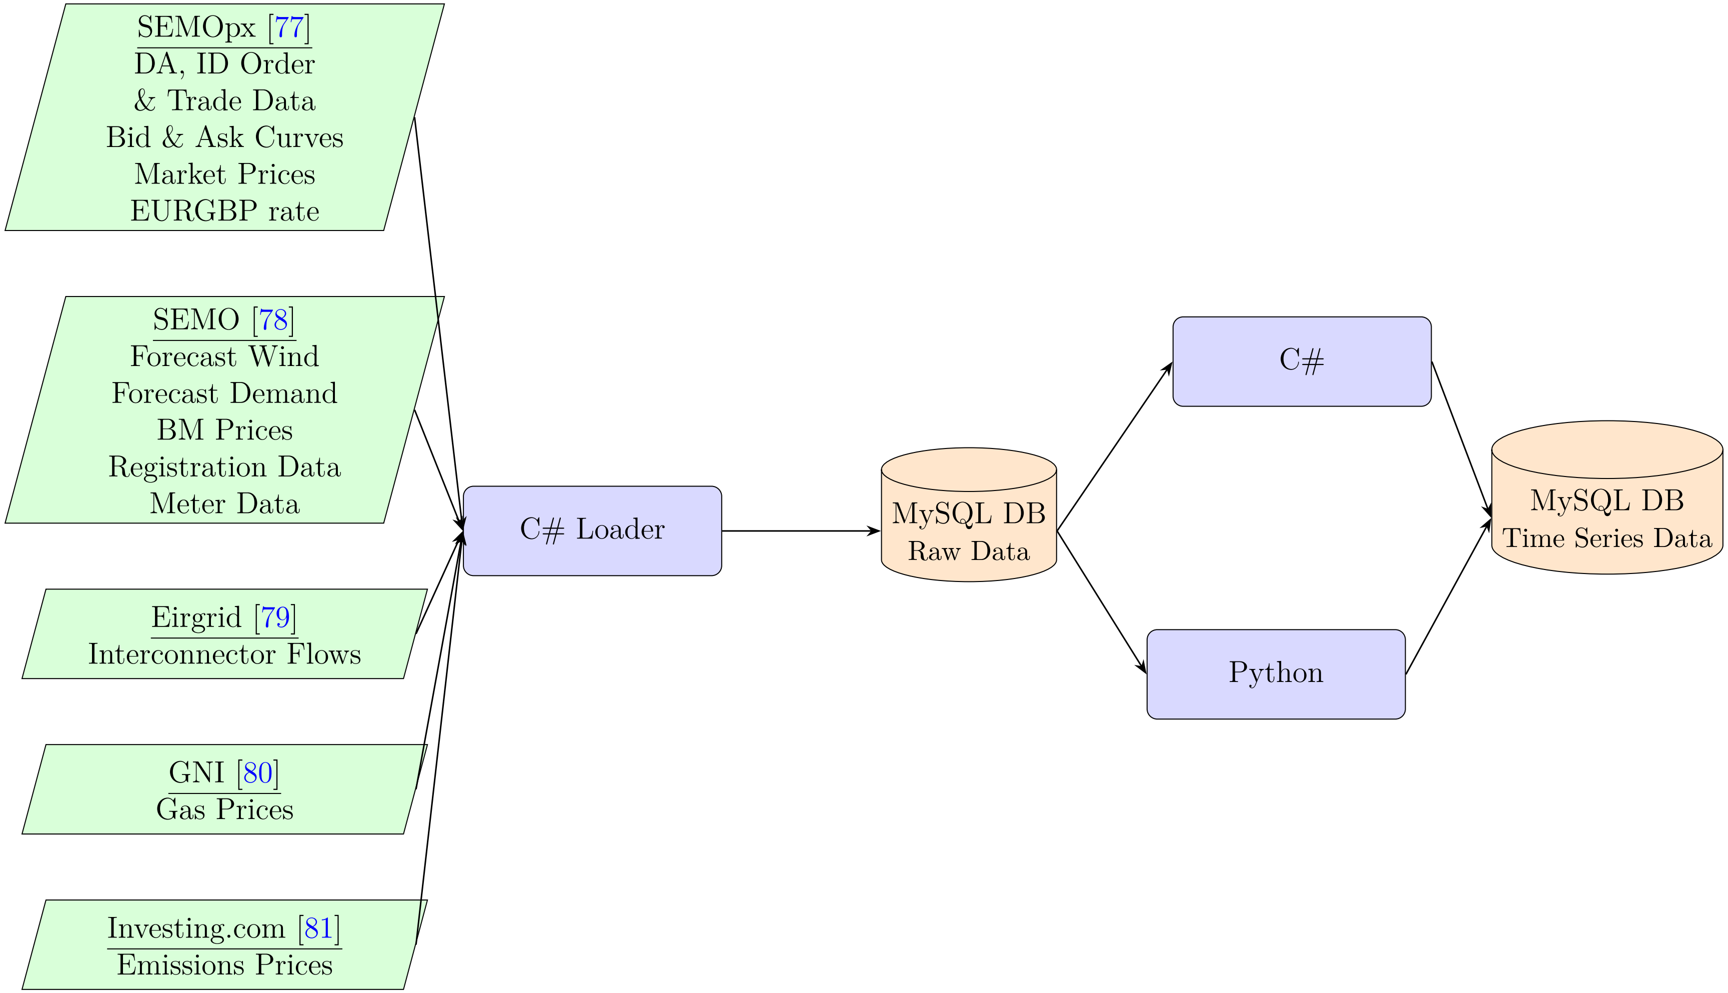

Figure 3: An overview of the data pipeline used in this thesis.

3.3 Datasets

Granular Participant Data SEMOpx publishes a distinct *ETS Bid File* for each of the DA, IDA1, IDA2, and IDA3 auctions enabling analysis of order and executed (trade) data at a participant and trading period level. Positive (negative) order quantities represent purchase (sell) orders, and the same convention applies to executed quantities (i.e. trades). The DA ETS Bid File is published on a D+1 basis relative to the trading day and typically contains over twenty thousand rows[24], with an average of over three hundred participants per auction. As the file includes data from market participants in the Republic of Ireland (ROI) and Northern Ireland (NI), monetary values are recorded in either € or £. Initially, the data is stored as is. In a subsequent processing step, £ amounts are converted to € using the relevant €/£ spot rate[25], the orders (both the converted orders and the original EUR orders) are then saved to distinct database tables. The file complexity can largely be attributed to the significant flexibility participants have in constructing their orders, as discussed in chapter 4.

Participant	Type	Value
AU_400003	Price	-500,-62,-62,-3,-3,27,27,57,57,82.8,82.8,84.8,84.8,89,89,119,119,164,164,3000
AU_400003	Quantity	110,110,98,98,68,68,48,48,38,38,-12,-12,-62,-62,-72,-72,-87,-87,-112,-112
GU_400530	Price	-500,55.98,55.98,56.98,56.98,57.08,57.08,57.94,57.94,76,76,3000
GU_400530	Quantity	0,0,-300,-300,-330,-330,-340,-340,-360,-360,-368,-368
GU_401930	Price	-500,-499.95,-499.95,3000
GU_401930	Quantity	0,0,-39.3,-39.3
SU_400071	Price	-500,435.63,435.63,522.76,522.76,609.89,609.89,697.01,697.01,784.14,784.14,871.27,871.27,3000
SU 400071	Quantity	1082.6,1082.6,1064.6,1064.6,1046.5,1046.5,1028.5,1028.5,1010.4,1010.4,992.4,992.4,974.3,974.3

Figure 4: Price and quantity values for the four market participants for trading period 6pm to 7pm 2nd February 2021.

For context, figure 4 displays the price and quantity data as extracted from a ETS Bid file for a specific trading period. Figure 5 visualises this data, following industry convention prices are on the vertical axis and quantities on the horizontal axis. While the rules governing market orders in the Irish electricity market are detailed in [82] and illustrated by example in chapter 4, it is sufficient to note here that a *Simple Order* submitted by a participant relates to a single trading period and must contain between 2 and 256 price-quantity pairs. These pairs form a stepwise function. For reference, the black dashed horizontal line in figure 5 indicates the market-clearing price for the trading period.

Bid Ask Curve Data SEMOpx publish a *BidAskCurve* file for each ex-ante auction containing a monotonic increasing (decreasing) and anonymised view of

[24] This large file size is in part due to the inclusion of both original orders and any updates that occur to the orders [76].

[25] A €/£ spot rate is associated with each auction and published by SEMOpx.

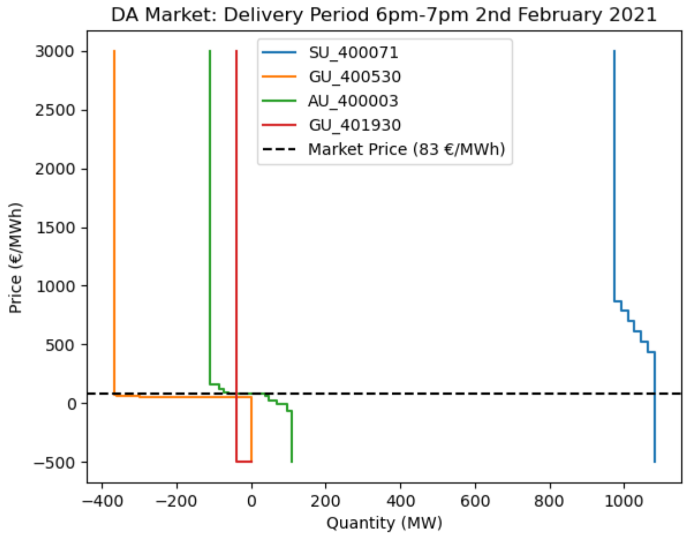

Figure 5: Order data for four market participants.

the sell (buy) orders for each trading period in the auction. Prior to July 2021 SEMOpx would publish two BidAskCurve files for each ex-ante auction, one containing information in £/MWh relating to the NI marketarea and the other containing information in €/MWh relating to the ROI marketarea. As of July 2021 SEMOpx have simplified matters by publishing a single BidAskCurve file for each ex-ante auction with the bid and ask information per trading period presented in €/MWh. The DA related BidAskCurve files are available on a D+1 basis relative to the trading day and in figure 6 we present an example of the ROI marketarea curve for trading period 6pm-7pm 2$^{\text{nd}}$ February 2021.

A general observation, also noted in section 2.5, is the lack of detailed public information regarding market algorithms and related processes. For example, no publication from SEMOpx or any other source we are aware of details how the published bid and ask curves are generated. This became relevant when we aggregated the order data from the ETSBidFile and compared it to the anonymised output in the BidAskCurve file (we did this as part of an independent data validation exercise). After considerable back-solving effort we managed to reconcile the BidAskCurve file with the ETSBidFile, as discussed in appendix C. Other questions however remained unanswered although they did not directly impact the research. For instance, we observed a step change in the nature of the demand and supply curve intersections in the BidAskCurve file (chapter 5). While it is possible to hypothesise about potential causes, the lack of documentation and limited visibility into the underlying algorithms and associated processes left the question unresolved.

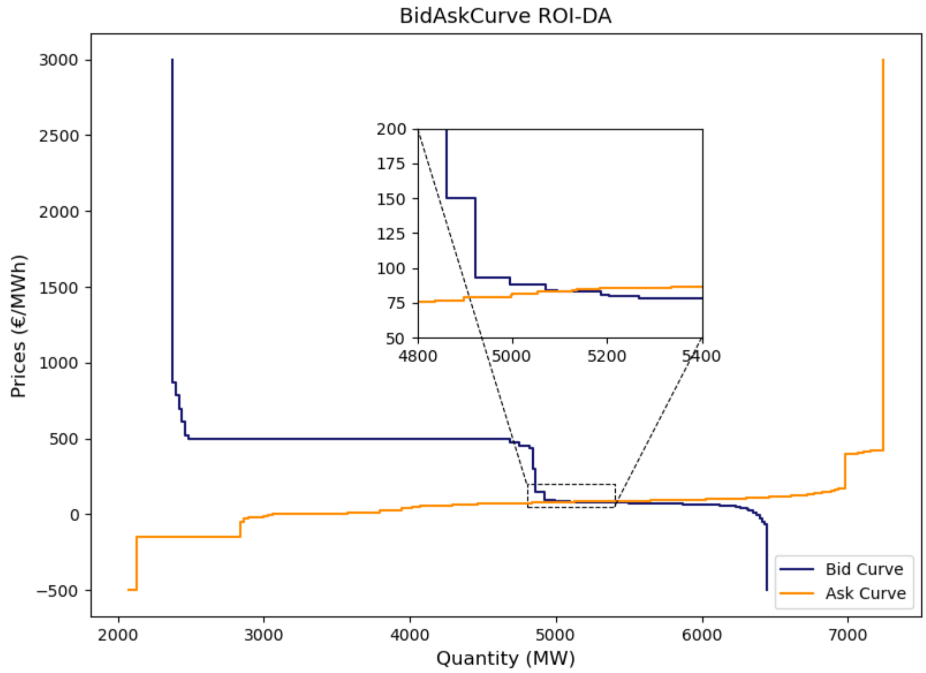

Figure 6: DA bid and ask (i.e. demand and supply) curve for ROI marketarea for 6pm-7pm 2$^{\text{nd}}$ February 2021.

Forecast System Wind and Demand As per [78] SEMO publishes system-wide forecast data for wind power generation and electricity demand. The *Four Day Aggregated Rolling Wind Unit Forecast Report* is produced four times daily, providing aggregated output (measured in megawatts, MW) of forecast wind generation across ROI and NI for the next four days. In addition the *Daily Load Forecast Summary Report* is published hourly, offering demand forecasts by jurisdiction for the same forecast horizon. For the purposes of the research in chapter 4 we are only interested in a subset of these files and specific forecast horizon intervals within each file. As explained in chapter 4, when forecasting participant DA commercial order data it is essential to use only information that would be available before the DA auction gate closure at 11am each day. Therefore, for demand, we use the file published around 10am each day, and for wind, the file published circa 6am each day. From both of these source files we use information relating to the [11pm D, 10pm D+1] forecast horizon[26]. In the flowchart shown in figure 3, the C# process on the right identifies the relevant file and forecast horizon data before storing it in database time series table. Figure 7 shows the forecast time series for a sample day (Sunday) in December 2022 showing both wind and demand forecasts in addition to actual (outturn) values. For reference, appendix A figure 42 shows forecast data for the full month of

[26]i.e. this assumes that we are generating the forecast before 11am on day D; the process is then repeated for each day in the forecast horizon.

December 2022. The plots illustrate that, in the Irish electricity market, forecast wind generation can represent a substantial share of forecast electricity demand. In addition, actual values can differ significantly from forecasts; this may be something that influences or impacts participants trading in short-term electricity markets.

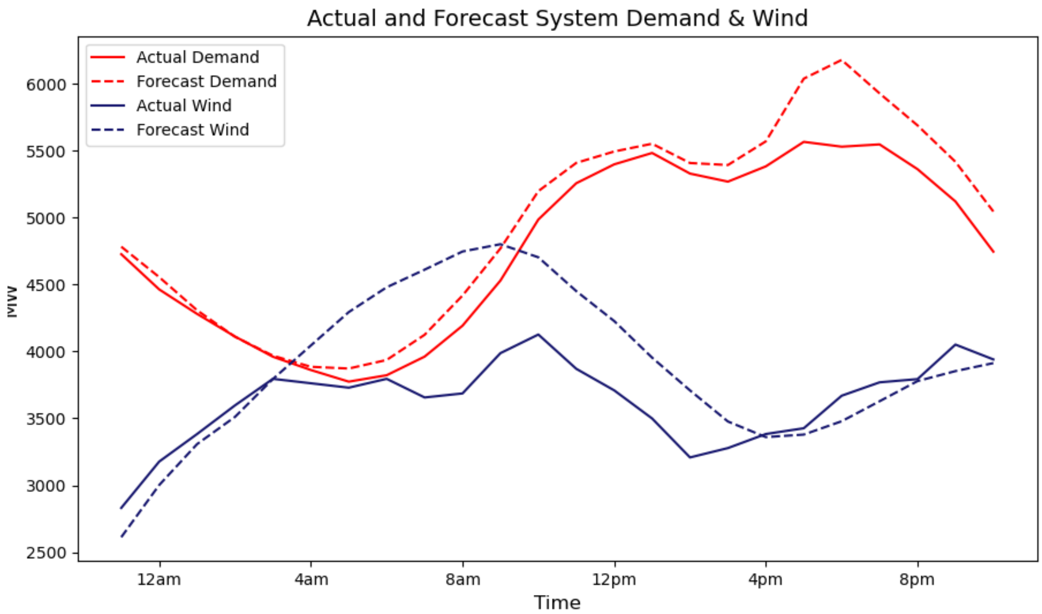

Figure 7: Forecast system wind and demand 18th December 2022.

DA, Intraday and Balancing market prices DA and ID market prices from the auctions are published by SEMOpx in the *ETS Market Results File* [77]. Although the file is published on a D+1 basis relative to the trading day, market participants can ascertain market prices prior to that via the market platform[27] which communicates their auction results and market prices after the auction concludes. Elsewhere, Balancing market prices are available in the *Imbalance Price Report (Imbalance Settlement Period)* and these thirty minute prices are published by SEMO after the end of each imbalance settlement period on the trading Day [78]. Figure 8 shows DA, IDA1, and Balancing Market prices from November 2018 to year-end 2022, while figure 9 presents the corresponding price series for three consecutive days in July 2021. In figure 8, periods of both reduced and elevated price volatility are apparent (notably during the COVID-19 pandemic and the subsequent European energy crisis). In addition, figure 9 highlights the strong correlation between DA and IDA1 prices and the generally higher volatility of Balancing Market prices.

[27] *ETS* platform which is owned and operated by *EPEX Spot*.

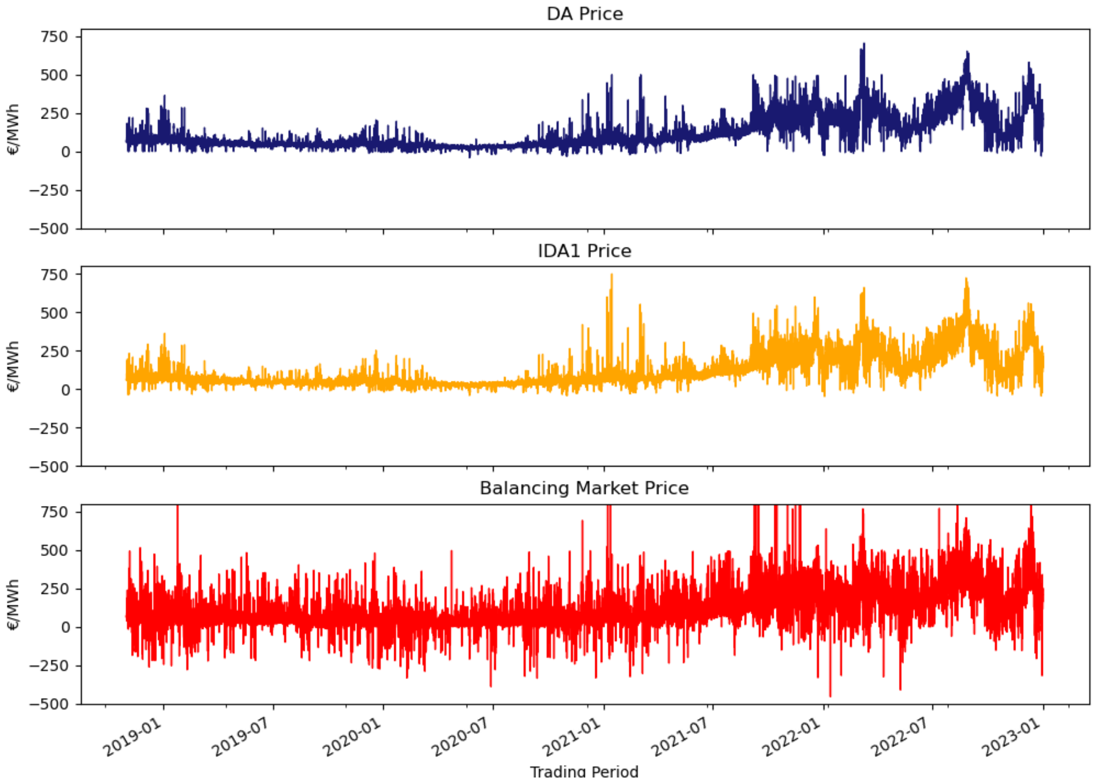

Figure 8: Time series of Irish DA, IDA1, and Balancing market prices (y-axis cropped to [-500, 800] €/MWh). From November 2018 and December 2022 the Balancing Market exceeded €800/MWh in 56 periods with a max price €4,800/MWh.

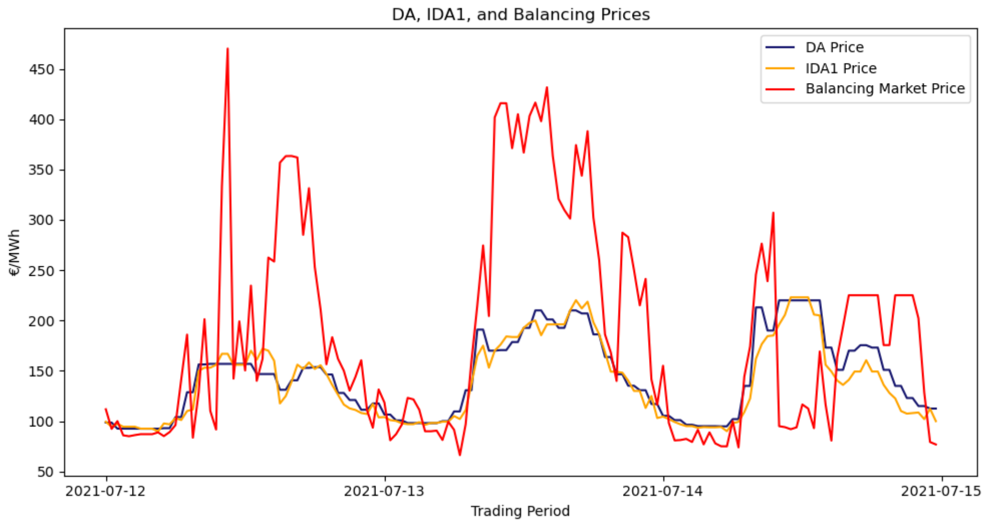

Figure 9: Time series data for DA, IDA1, and Balancing Market prices over three consecutive days in July 2021 (i.e. a zoomed-in view of the series shown in figure 8).

Gas and Carbon Prices Gas and carbon price time series are relevant to the analysis in chapter 4 when one considers that the thermal generation fleet in ROI and NI primarily comprises of generators that use gas as their primary fuel source [83]. Hence, for these generators, fuel input costs (i.e. gas and the costs of procuring carbon emission certificates) often influence participants' commercial order decisions. Although ROI has indigenous gas supplies, they are insufficient to meet demand necessitating gas imports from Great Britain (GB). As a result, the GB gas price effectively determines the wholesale gas market price in Ireland. For the purposes of modelling the day-ahead electricity market in chapter 4, the preferred gas price quote would logically be Day-Ahead Notional Balancing Point, NBP, price quotes (i.e. Day-Ahead GB gas market quotes) just before the DA Irish electricity market gate close at 11am. However, as we do not have access to this data we use the NBP System Average Price (NBP SAP) which is the average price of system balancing trades conducted by the gas system operator over the day[28]. Typically, the NBP DA gas price and NBP SAP are highly correlated. For carbon prices, we use the front-month carbon emissions future as displayed in [81]. Figure 10 provides a time series plot of gas prices (blue, left-hand axis) and the daily average DA electricity market prices (green, right-hand axis). The strong correlation between gas and electricity prices in the Irish market is evident and the underlying linkages become clearer in the context of the analysis presented in chapter 5.

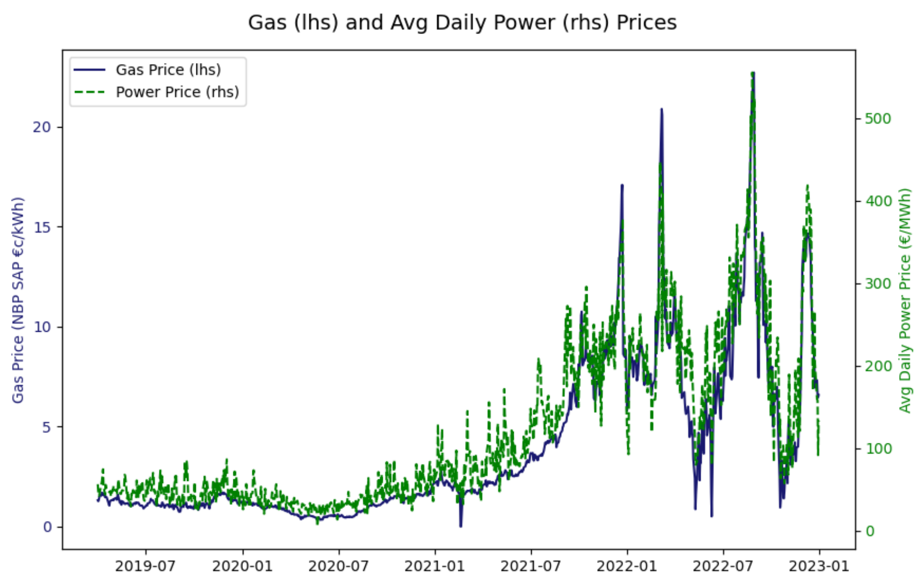

Figure 10: Time series of gas prices and daily average DA power prices.

[28]An alternative could be the Irish gas market equivalent, the IBP SAP, but its reliability is limited due to the high number of days with no balancing trades in the Irish gas market.

Proxy Aggregate Availability Data Combining sell order quantity data from the DA ETS Bid File with registered fuel type information from the *Registered Capacity Report*[29], we derive a time series of DA sell order quantities by fuel type, as shown in figure 11. Wind generator units are not included in this time series as they are already represented by a previously discussed dataset. The *None* category represents those market participants listed in the Registered Capacity Report that do not have an associated fuel type (this group is used to help identify speculators which are analysed in chapter 5). In chapter 4 we use a subset of this proxy availability data as explanatory variables. The rationale is that generation fleet availability, or expectations thereof, may influence the commercial behaviour of certain market participants. We note that this availability data approach was in part chosen for ease of implementation, alternative approaches and data sources are possible.

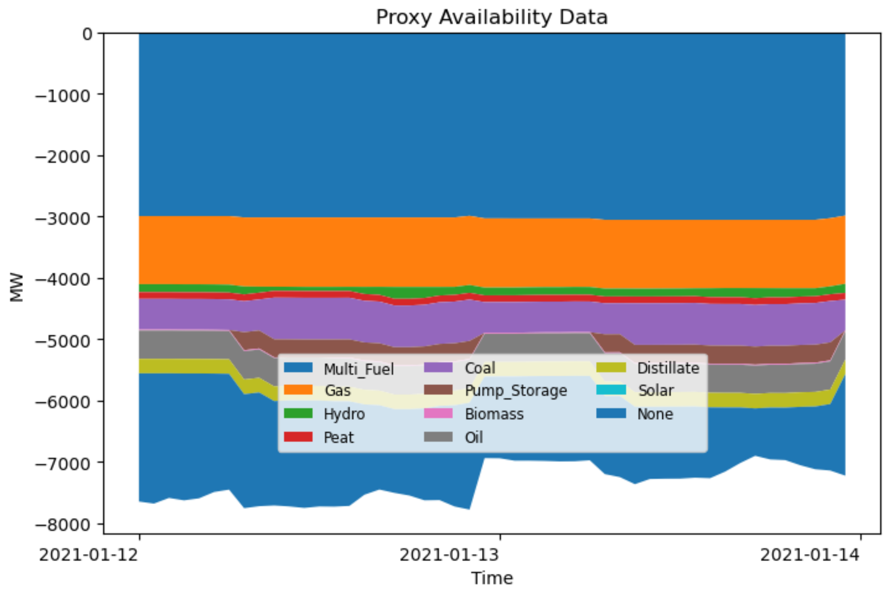

Figure 11: Proxy Availability Data over a two-day period showing submitted sell order volumes by fuel type in the DA market. The full time series is presented in appendix A figure 43.

Interconnector For the time horizons studied in chapters 4 and 5 the Moyle and East West Interconnector (EWIC) provided the physical links between the electricity markets in Ireland and Great Britain. Interconnector flows are available via [84] and figure 12 illustrates the flows (in MW) on both interconnectors in January and February 2021 for example. Positive values indicate flows from GB to Ireland, while negative values indicate flows from Ireland to GB. We note that

[29]This report provides registered capacity and fuel type details for all generator and demand-side units in the Irish electricity market [78].

the size of the interconnection, relative to the demand on the island of Ireland, is significant and hence their use as potential explanatory variables in chapter 4.

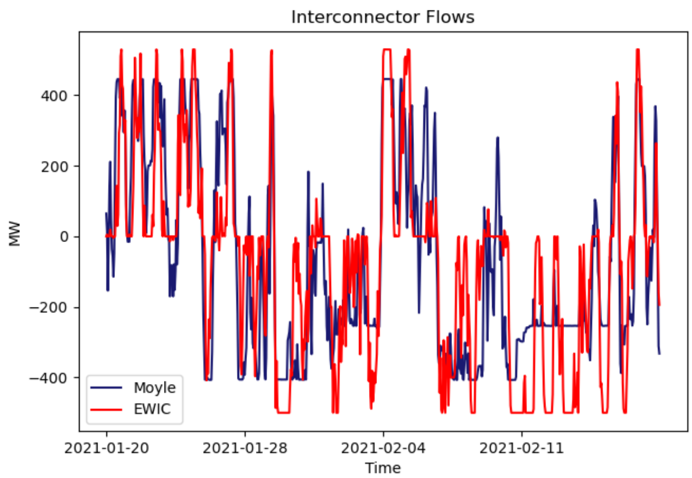

Figure 12: Interconnector flows between Ireland and GB over successive days in January and February 2021. Flows from Ireland to GB are negative.

Participant Meter Data The *Daily Meter Data Report*, published on a D+1 basis relative to the trading day by SEMO, provides metered generation (in MWh) for generator and interconnector units [78]. Figure 13 shows metered generation for a sample market participant over the course of November 2021. These datasets are of interest in chapter 4 when forecasting participant commercial order data. The underlying rationale is that participants may consider their operational status when constructing orders[30].

Additional Considerations As will be discussed in chapter 4, additional datasets may exist that are not incorporated into our analysis and may influence market participants' decisions. Specifically, datasets accessible through data provider subscription offerings or by trading in exchange or *Over the Counter (OTC)* markets such as:

- GB DA electricity market prices, GB DA gas prices.
- Forward prices.

Before the Irish DA electricity market gate closes at 11am, some participants may consider GB DA electricity prices from either the OTC market or the Nord Pool auction [85] when constructing their orders[31]. Separately, and as discussed

[30] For instance, if a generation unit is not running, the participant might include a fixed term in a *Complex Order* (discussed in chapter 4) to account for start-up costs.
[31] These prices may provide an indication of likely interconnector flow directions.

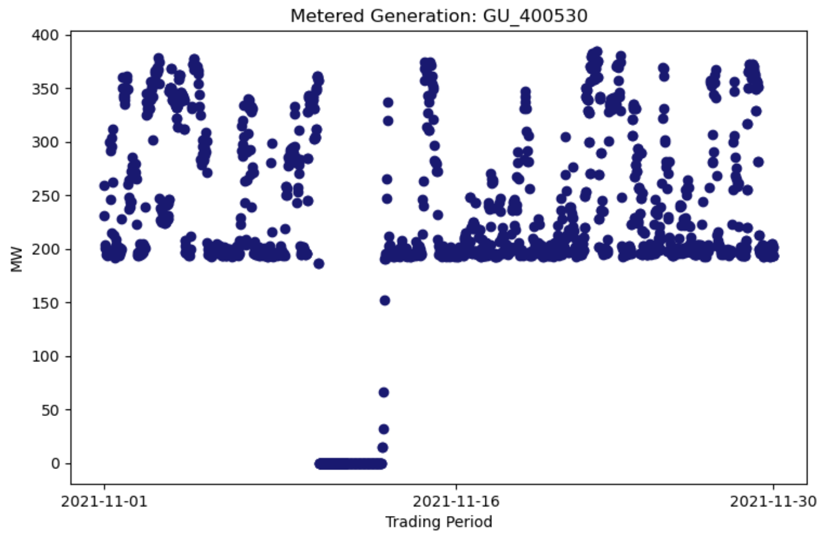

Figure 13: Metered generation for participant GU_400530.

earlier, OTC NBP DA gas prices (i.e. GB DA gas prices) which are typically available before the Irish electricity market 11 am gate closure might be a more accurate gas price indicator than the NBP SAP benchmark, albeit they are expected to be highly correlated. The final bullet point is simply a recognition that while most short-term EPF literature (discussed in chapter 2) does not include forward market prices as explanatory variables, they may still be relevant. The Irish electricity forward market has historically been relatively illiquid, making it reasonable to exclude related forward price data from our analysis. In contrast, the GB forward markets for gas and electricity are more liquid and incorporating these as explanatory variables may warrant further exploration. While such additions could potentially improve forecast accuracy, we maintain that their exclusion does not undermine the validity of the research presented in chapter 4.

3.4 Models

This section provides an overview of the main modelling techniques utilised in chapter 4. Given the extensive published literature on each model, our aim is not to present detailed derivations but rather to provide a synopsis. Our discussion presents the models in a single output context[32]. However, these models can be readily extended to a multi-output (or multitarget) framework where the output is a vector. It is in this context that they will be encountered in chapter 4. We begin with the naive model before proceeding to random forests, feedforward neural networks, linear regression and shrinkage techniques (*RIDGE*, *LASSO* and *ELASTIC NET*) and finally support vector machines. It is worth noting

[32]The output from the model is a scalar.

that while additional modelling techniques[33] were considered during the initial model exploratory analysis phase relating to chapter 4, these approaches were not incorporated into the final analysis and are therefore not discussed here.

Naive Model An example of the *'weekly persistent model, sometimes called naive model'* is given in [51]. Under this model, to generate a forecast of the DA market price for day d hour h for example, denoted by $\widehat{y}_{d,h}$, we can simply use the market price from the same hour seven days earlier (i.e. $y_{d-7,h}$) as the estimate. That is

$$\widehat{y}_{d,h} = y_{d-7,h} \tag{1}$$

Variants of this naive forecasting model are discussed in [35] and [57]. Under these alternative approaches, the forecast $\widehat{y}_{d,h}$ can be generated from

$$\widehat{y}_{d,h} = \begin{cases} y_{d-7,h}, & \text{if } d \in \{\text{Monday, Saturday, Sunday}\} \\ y_{d-1,h}, & \text{otherwise} \end{cases} \tag{2}$$

In chapter 4, we use a variant of the naive model to forecast participant commercial order data whereby the most recently available historical order for a participant is used as the forecast. Owing to a two-day publication lag in the underlying datasets, this corresponds to using a $d-2$ value on the right-hand side of equation 2. The naive model can serve as a benchmark for evaluating competing models, a role that becomes clearer when we discuss error metrics in section 3.5.

Random Forests The overview given in the following paragraphs closely follows standard treatments of tree-based methods, in particular those presented in [86] and [87]. For a discussion of Random Forests (RF) in a regression context (where the goal is to use explanatory variables to forecast a continuous target variable) it is first necessary to understand the concept of a *decision tree regressor*. The core idea is to recursively partition the dataset, starting from the full dataset (the parent node), into two subsets (the child nodes) by selecting an explanatory variable and a corresponding split point that minimises the squared error loss. That is, the algorithm searches for the split that leads to the greatest reduction in the sum of squared differences between the observed target values and their mean within each resulting subset. This splitting process continues recursively on each child node until a predefined stopping criterion is met, such as a minimum number of samples in a terminal node or when the maximum tree depth is reached. Once constructed, the decision tree regressor can be used for

[33] e.g., k-nearest neighbours, gradient boosting machines.

prediction: given a new input, the observation is passed down the tree according to its feature values until it reaches a terminal (leaf) node, where the forecast is given by the average of the target values in that node. To develop a deeper understanding of the approach, let us consider a set of N observations (x_i, y_i), where i indexes the observation pairs. Each x_i consists of n explanatory variables, represented as $(x_{i1}, x_{i2}, \ldots, x_{in})$, and y_i is the corresponding target value, that is the quantity we aim to predict. As explained above, starting with all N observations, the approach is to determine a splitting variable $k \in \{1, \ldots, n\}$ and a split point $t \in \mathbb{R}$ that partition the data into two distinct subsets, denoted $R_1(k,t)$ and $R_2(k,t)$, defined as

$$R_1(k,t) = \{x_i \mid x_{ik} \leq t\} \quad \text{and} \quad R_2(k,t) = \{x_i \mid x_{ik} > t\} \tag{3}$$

The variable k and split point t are chosen by solving

$$\min_{k,t} \left[\sum_{x_i \in R_1(k,t)} (y_i - \hat{c}_1)^2 + \sum_{x_i \in R_2(k,t)} (y_i - \hat{c}_2)^2 \right] \tag{4}$$

with

$$\hat{c}_1 = \frac{1}{|R_1(k,t)|} \sum_{x_i \in R_1(k,t)} y_i \quad \text{and} \quad \hat{c}_2 = \frac{1}{|R_2(k,t)|} \sum_{x_i \in R_2(k,t)} y_i \tag{5}$$

where $|R_j(k,t)|$ denotes the number of observations in subset R_j, for $j = 1, 2$. Here, min indicates that the values of k and t are chosen to minimise the total squared error across the two resulting regions. The process is then recursively applied to the resulting subsets, and the regression estimate is given by the mean of the y_i values in the terminal subset to which the observation belongs. Given this understanding, a RF is a collection of decision tree regressors, each trained on the same prediction task. In a RF the final regression estimate is given by the average of the predictions from the individual trees. It should be noted that when constructing each decision tree only a subset of the N observations is used (typically selected via sampling with replacement) and at each split point a randomly chosen subset of the n explanatory variables is considered when determining how to partition the data.

An interesting and attractive feature of RF is their ability to associate a feature importance score with each explanatory variable. One approach to calculating feature importance is described in [87] and is *'at each split in each tree, the improvement in the split-criterion is the importance measure attributed to the splitting variable, and is accumulated over all the trees in the forest separately for each variable'*. An alternative approach to calculating what [88] refers to as the *'permutation importance'* is also possible[34].

A subjective critique of RF is as follows. Advantages include their versatility for both regression and classification contexts and applicability across diverse problem domains. As noted earlier another advantage is that RF can provide feature importance values along with out-of-bag error estimates as useful by products. Disadvantages of RF include their *'black-box'* nature and susceptibility to overfitting which can lead to *'jagged'* regression estimates. The *'black-box'* characterisation is perhaps valid in that unlike a single decision tree a RF consists of multiple trees making it difficult to visualise individual decision making paths. RF models can underfit or overfit the training data (the latter often being associated with jaggedness), but careful hyperparameter tuning can help mitigate this.

Feedforward Neural Networks For an introduction to *artificial neural networks (ANN)* see [89]. Additional reference material include the chapter on neural networks in [87], the excellent and intuitively motivated overview by [90], and the books [86] and [91]. Both [86] and [91] provide practical guidance on implementing neural network models using open-source Python frameworks such as *TensorFlow* and *Keras*. The neural network overview given in the following paragraphs, along with its associated notation, closely follows the approach presented in *Deep Learning: An Introduction for Applied Mathematicians* [92].

A simplified description of a feedforward neural network (FFNN), and the process of training one, is as follows. Again, consider a dataset consisting of N observation pairs (x_i, y_i) where each $x_i = (x_{i1}, x_{i2}, \ldots, x_{in})$ represents a vector of n explanatory variables, and y_i is the corresponding target value, that is the quantity to be predicted. A neural network can be conceptualised as a flexible function approximator: a model designed to approximate the unknown relationship between inputs and outputs based on data. Denoting the function by f, our goal is to learn model parameters (or weights), denoted p, such that $\hat{y}_i = f(x_i; p) \approx y_i$. To train the model, we begin by passing an input observation x_i through the network to generate a prediction $\hat{y}_i$; this is known as the *forward propagation* step. We then adjust or update the parameters p in such a way as to to bring the prediction $\hat{y}_i$ closer in line to the true value y_i. The information needed to guide this update is computed efficiently through a process called *back-propagation* (described further below). These two steps, computing the network's

[34]For each *out of bag* sample, pass the observations down the associated tree (i.e. prediction step) and calculate the mean squared error. Next, taking one of the input variables, randomly permute its values in the out of bag sample, pass the observations down the tree, calculate the mean squared error. The process is repeated across the remaining trees. The average mean squared error difference across all the trees is the metric of interest. The process is repeated across the remaining input variables.

output and updating the weights, are repeated for each observation in the training set[35]. The training process continues until a stopping criterion is met or a predefined number of *epochs* (i.e. complete passes through the training data) has been completed.

For a deeper understanding let us consider a network with L layers, layer 1 being the input layer and layer L being the output layer. Each layer, l, contains n_l neurons. Every neuron in layer $l-1$ is connected, or acts as input, to each neuron in layer l; that is it is a fully connected network. A simple hypothetical example is shown in figure 14. We use the following notation: let $W^{[l]} \in \mathbb{R}^{n_l \times n_{l-1}}$

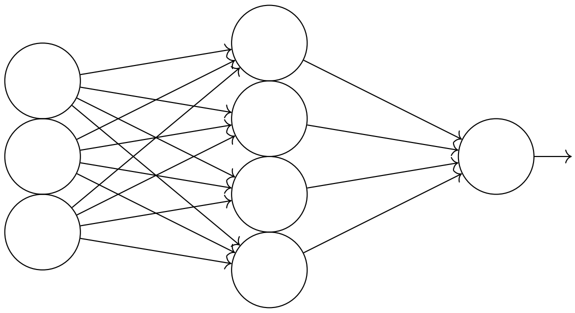

Figure 14: Feedforward neural network. The input layer consists of three nodes (representing the explanatory variables). The hidden layer consists of four nodes, fully connected to the input layer. The output layer consists of a single node.

be the weight matrix at layer l, where $w_{jk}^{[l]}$ represents the weight connecting node j in layer l to node k in layer $l-1$. Similarly, let $b^{[l]} \in \mathbb{R}^{n_l}$ be the bias vector, where $b_j^{[l]}$ is the bias for neuron j in layer l. For notational simplicity, we drop the observation index i and work with a representative single training example. We define $z^{[l]}$, $a^{[l]}$ and $\delta_j^{[l]}$ by

$$z^{[l]} = W^{[l]} a^{[l-1]} + b^{[l]} \in \mathbb{R}^{n_l} \; for \; l = 2, 3, ..., L \tag{6}$$

$$a^{[l]} = \sigma(z^{[l]}) \; for \; l = 2, 3, ..., L \tag{7}$$

$$\delta_j^{[l]} = \frac{\partial C}{\partial z_j^{[l]}} \; for \; 1 \leq j \leq n_l \; and \; 2 \leq l \leq L \tag{8}$$

In equation (7), σ is an activation function applied elementwise to the input vector $z^{[l]}$; the choice of activation function is discussed further below. Equations (6)

[35]This can be done individually for each observation (known as *stochastic gradient descent*) or, more commonly, in batches, referred to as *mini-batch gradient descent*.

and (7) describe the forward propagation step (whereas $\delta_j^{[l]}$ serves as an intermediate term in the derivation of the back propagation algorithm and as seen further below is subsequently used to compute network weight updates). Along similar lines to the example in [92], we define the cost function C, which measures the distance between the network output and the target, as[36]

$$C = \frac{1}{2}(y - a^{[L]})^2 \tag{9}$$

With these definitions in place, utilising the chain rule for differentiation, it can be shown that

$$\delta^{[L]} = \sigma'(z^{[L]}) \circ (a^L - y) \tag{10}$$

$$\delta^{[l]} = \sigma'(z^{[l]}) \circ (W^{[l+1]})^T \delta^{l+1} \; for \; 2 \leq l \leq L-1 \tag{11}$$

$$\frac{\partial C}{\partial b_j^{[l]}} = \delta_j^{[l]} \; for \; 2 \leq l \leq L \tag{12}$$

$$\frac{\partial C}{\partial w_{jk}^{[l]}} = \delta_j^{[l]} a_k^{[l-1]} \; for \; 2 \leq l \leq L \tag{13}$$

where $\circ$ denotes to element wise multiplication. Equations (10) and (11) describe the *back propagation* step, where computations start at the output layer and move backward to determine the δ values for earlier layers. These values are then used to compute the partial derivatives of the cost function with respect to the network parameters via equations (12) and (13), after which the network weights are updated using equations (14) and (15). The η parameter in equations (14) and (15) is referred to as the *learning rate*.

$$w_{jk}^{[l]} \rightarrow w_{jk}^{[l]} - \eta \frac{\partial C}{\partial w_{jk}^{[l]}} \quad \text{for } 2 \leq l \leq L \tag{14}$$

$$b_j^{[l]} \rightarrow b_j^{[l]} - \eta \frac{\partial C}{\partial b_j^{[l]}} \quad \text{for } 2 \leq l \leq L \tag{15}$$

As described in the high-level introduction, the process is repeated until some stopping criterion is met or a predefined number of epochs have completed.

When fitting a neural network, there are a number of considerations. For example, historically the activation function (i.e. the σ in equation (7)) of choice was the sigmoid function,

$$\sigma(x) = \frac{1}{1 + e^{-x}}, \tag{16}$$

which maps a scalar input x to the interval $[0, 1]$. Alternatives include the hyper-

[36] Alternative cost functions may result in a different equation (10).

bolic tangent (which outputs in $[-1, 1]$) and the leaky ReLU:

$$\text{leaky ReLU}(x) = \max(0.1x, x). \tag{17}$$

Leaky ReLU is computationally efficient and may help mitigate the vanishing or exploding gradients issue[37], where weight updates become excessively small or large, hindering training. The choice of learning rate η is also important; a large η may overshoot the minimum, while a small η may slow model training. Related to this is momentum optimisation, which incorporates past gradient updates. An intuitive discussion on this topic is found in [86]. Another consideration is *regularisation* and *dropout*, both used to mitigate overfitting. Regularisation, also discussed in the next subsection, modifies the cost function in equation (9) by adding a penalty term based on weight magnitudes whereas dropout randomly deactivates neurons during training. Other considerations include *parameter initialisation* and *batch normalisation* which are approaches to tackling the vanishing gradients problem. In [86], the author refers to a 2010 paper by Glorot and Bengio [93] in which they explain that the manner in which the activation function is initialised can result in vanishing gradients, in the same paper they propose an approach for initialising the sigmoid activation function which helps address/mitigate the issue[38]. Even using the aforementioned parameter initialisation technique, it was observed that the vanishing gradients problem could return during training. To address this, batch normalisation was proposed in 2015 and as explained in [86] in it involves *'adding an operation in the model just before the activation function of each layer, simply zero-centring and normalising the inputs, then scaling and shifting the result using two new parameters per layer....this operation lets the model learn the optimal scale and mean of the inputs for each layer'*. We also note that network topology (i.e. the number of layers and neurons per layer) has been considered in some of the EPF literature. As in other domains, there are no fixed rules for configuring a network; it is advised that some level of experimentation with different settings is performed based on the problem at hand. Other relevant topics, not covered here, include feature scaling[39] and early stopping[40]

What are the advantages and disadvantages of FFNN? While subjective, advantages (similar to those of RF) include their versatility in handling both classification and regression tasks. With proper tuning, FFNNs can be applied across diverse problem domains. Additionally, some EPF studies suggest that FFNN

[37] For an intuitive explanation of vanishing or exploding gradients, see chapter 5 in [90].

[38] Other researchers proposed initialisation methods for alternative activation functions.

[39] i.e. scaling of variables informed by the choice of the activation function.

[40] Aims to address overfitting by halting training when the validation error stops improving.

offer a superior, albeit slightly, forecasting performance. Potential disadvantages include that the models will typically take longer to train. Fitting a FFNN often requires a degree of experimentation, unless care is taken with configuring the network and hyperparameter settings then it is possible to under or over train the network.

Linear Regression and Shrinkage techniques We follow the exposition of linear regression and shrinkage methods presented in chapter 3 of [87] which offers a clear and insightful treatment of the topic. As in the preceding discussion, we are given a dataset of N observations (x_i, y_i), where each x_i is a vector consisting of n explanatory variables, $(x_{i1}, \ldots, x_{in})$, and y_i is the corresponding response. The standard linear regression model assumes the following form:

$$f(x_i) = \beta_0 + \sum_{j=1}^{n} x_{ij}\beta_j \tag{18}$$

where $\beta_0 \in \mathbb{R}$ is the intercept and $(\beta_1, \ldots, \beta_n) \in \mathbb{R}^n$ denotes the vector of coefficients associated with the explanatory variables. Given an input x_i, the function $f(x_i)$ maps the explanatory variables to a scalar output. To estimate the parameters of this model we determine the β values which minimise the *residual sum of squares (RSS)*, defined as:

$$RSS(\beta) = \sum_{i=1}^{N} \left(y_i - \beta_0 - \sum_{j=1}^{n} x_{ij}\beta_j \right)^2 . \tag{19}$$

Here, $\beta := (\beta_0, \beta_1, \ldots, \beta_n)$ denotes the vector of regression parameters, and the notation $RSS(\beta)$ indicates that the objective function depends on these parameters. Minimising the residual sum of squares in equation (19) yields the *ordinary least squares (OLS)* estimator. Once the model parameters have been estimated, the function $f(x_i)$ serves as a predictor for the response y_i, with the fitted value given by $\hat{y}_i = f(x_i)$. *RIDGE regression* extends the ordinary least squares framework by augmenting the RSS objective with an additional penalty term that discourages large coefficient values. As described in [87], *'the RIDGE coefficients minimize a penalized residual sum of squares'*:

$$\hat{\beta}^{\text{RIDGE}} = \underset{\beta}{\text{argmin}} \left\{ \sum_{i=1}^{N} \left(y_i - \beta_0 - \sum_{j=1}^{n} x_{ij}\beta_j \right)^2 + \lambda \sum_{j=1}^{n} \beta_j^2 \right\} . \tag{20}$$

The second term in equation (20) imposes a penalty on the size of the coefficients effectively shrinking them toward zero. The parameter λ is a tuning (or hyper-parameter) that controls the strength of the penalty: larger values of λ lead to greater shrinkage; argmin denotes the argument (i.e. the value of β) that minimises the expression inside the braces. In a similar manner the *least absolute shrinkage selector operator, LASSO*, estimate is given by

$$\hat{\beta}^{\text{LASSO}} = \underset{\beta}{\text{argmin}} \left\{ \sum_{i=1}^{N} (y_i - \beta_0 - \sum_{j=1}^{n} x_{ij}\beta_j)^2 + \lambda \sum_{j=1}^{n} |\beta_j| \right\} \tag{21}$$

LASSO regression, unlike RIDGE regression, can set some coefficient estimates exactly to zero, thereby performing variable selection and effectively removing the corresponding explanatory variables from the model. Unlike linear and RIDGE regression, there is no closed form solution for LASSO regression, instead the parameters can be estimated via coordinate descent[41]. Elsewhere, *ELASTIC NET*, is a combination of both the RIDGE and LASSO penalty terms and the parameter estimates are given by

$$\hat{\beta}^{\text{ELASTIC NET}} = \underset{\beta}{\text{argmin}} \left\{ \sum_{i=1}^{N} (y_i - \beta_0 - \sum_{j=1}^{n} x_{ij}\beta_j)^2 + \lambda \sum_{j=1}^{n} (\alpha\beta_j^2 + (1-\alpha)|\beta_j|) \right\} \tag{22}$$

It is worth noting that in addition to the λ hyperparameter in equation (22) there is an α hyperparameter which also needs to tuned.

A subjective view of the advantages of shrinkage methods is that given they are extension of linear regression they yield easily understood parameter estimates. Additionally, they are computationally efficient and quick to train. However, their effectiveness depends on the market and problem context. As evidenced by their inclusion in the EPF toolbox (section 2.3.2), they can offer strong forecasting performance in EPF settings. A disadvantage is that these models do not inherently capture nonlinear effects often requiring feature engineering in an effort to account for nonlinearities and interactions.

Support Vector Machines For a well-rounded introduction to Support Vector Machines (SVMs), particularly in the context of classification (Support Vector Classification, or SVC), see [86] and [87]. Given the research focus in chapter 4 we are primarily interested in Support Vector Regression (SVR), rather than SVC. References that provide an insightful overview of SVR include [95] and [96]. The following outline offers a summary of the key points from [95].

As before, we consider N observation pairs (x_i, y_i), with $x_i \in \mathbb{R}^n$ and $y_i \in \mathbb{R}$. As per [95], *'in ϵ-SV regression our goal is to find a function $f(x)$ that has at most ϵ deviation from the actually obtained targets y_i for all the training data and at the same time is as flat as possible'*. Starting with the case of a linear function

[41] The parameter estimates may also be derived via *Least Angle Regression, LARS*, [94].

f given by

$$f(x_i) = \langle w, x_i \rangle + b, \tag{23}$$

where $\langle w, x_i \rangle$ denotes the standard inner product between vectors w and x_i, and $w \in \mathbb{R}^n$ and $b \in \mathbb{R}$ are model parameters to be estimated. In this context, *flatness* refers to controlling the norm of the weight vector w, which limits the sensitivity of the function to variations in the input and thereby promotes good generalisation. As described in [95], the goal of achieving flatness while ensuring that predictions stay within a tolerance band of width 2ϵ around the true target values leads to the following optimisation problem:

$$\min \frac{1}{2}\|w\|^2 \quad \text{subject to} \quad \begin{cases} y_i - \langle w, x_i \rangle - b \leq \epsilon, \\ \langle w, x_i \rangle + b - y_i \leq \epsilon. \end{cases} \tag{24}$$

Here, min indicates that we aim to find model parameters that minimise the objective function $\frac{1}{2}\|w\|^2$, subject to the specified constraints. Similar to the concept of the *soft margin* in SVC, slack variables ξ_i, ξ_i^* are introduced to allow deviations from the ϵ tube (i.e. a tolerance zone where deviations within ϵ incur no penalty) in equation (24). Doing this yields

$$\min \frac{1}{2}\|w\|^2 + C\sum_{i=1}^{N}(\xi_i + \xi_i^*) \quad \text{subject to} \quad \begin{cases} y_i - \langle w, x_i \rangle - b \leq \epsilon + \xi_i, \\ \langle w, x_i \rangle + b - y_i \leq \epsilon + \xi_i^*, \\ \xi_i, \xi_i^* \geq 0, \quad i = 1, 2, ..., N. \end{cases} \tag{25}$$

The C parameter penalises instances outside the ϵ tube. Equation (25) represents the primal optimisation problem. To derive the corresponding dual optimisation problem (which enables the *kernel* trick and is briefly referenced further below), introducing Lagrangian Multipliers α_i, α_i^* for the margin constraints and η_i, η_i^* for the slack variable non negativity constraints, we obtain the Lagrangian:

$$\begin{aligned} L = \frac{1}{2}\|w\|^2 + C\sum_{i=1}^{N}(\xi_i + \xi_i^*) - \sum_{i=1}^{N}\alpha_i(\epsilon + \xi_i - y_i + \langle w, x_i \rangle + b) \\ - \sum_{i=1}^{N}\alpha_i^*(\epsilon + \xi_i^* + y_i - \langle w, x_i \rangle - b) - \sum_{i=1}^{N}\eta_i\xi_i - \sum_{i=1}^{N}\eta_i^*\xi_i^*. \end{aligned} \tag{26}$$

Taking partial derivatives of the Lagrangian with respect to the model parameters, and applying standard simplifications, yields the following dual optimisation problem:

$$\max_{\alpha_i,\alpha_i^*} \sum_{i=1}^{N} y_i(\alpha_i - \alpha_i^*) - \epsilon \sum_{i=1}^{N}(\alpha_i + \alpha_i^*) \\ - \frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{N}(\alpha_i - \alpha_i^*)(\alpha_j - \alpha_j^*)\langle x_i, x_j\rangle \\ \text{subject to} \quad \begin{cases} \sum_{i=1}^{N}(\alpha_i - \alpha_i^*) = 0, \\ 0 \le \alpha_i, \alpha_i^* \le C, \quad i = 1, \ldots, N. \end{cases} \tag{27}$$

The dual optimisation problem is a quadratic programming (QP) problem and can be solved using standard QP solvers. Using the primal form in equation (23) and intermediate steps from the derivation of the dual, it can be shown that

$$f(x) = \sum_{i=1}^{N}(\alpha_i - \alpha_i^*)\langle x_i, x\rangle + b. \tag{28}$$

Therefore, to compute $f(x)$, only those training inputs x_i for which $\beta_i = \alpha_i - \alpha_i^* \neq 0$ are required, these are the *support vectors*. The above treatment handles the linear regression case. For non-linear regression, the previously mentioned kernel trick is utilised. Here we purposefully gloss over the associated mathematical complexities but an excellent conceptual summary is given by [86]: *'a kernel is a function capable of computing the dot product $\langle\phi(a), \phi(b)\rangle$ based only on the original vectors a and b without having to compute (or even to know about) the transformation ϕ'.* Armed with that intuitive explanation, replacing $\langle x_i, x_s\rangle$ in equations (27) and (28) with the kernel function yields[42]

$$\max \sum_{i=1}^{N} y_i(\alpha_i - \alpha_i^*) - \epsilon \sum_{i=1}^{N}(\alpha_i + \alpha_i^*) - \frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{N}(\alpha_i - \alpha_i^*)(\alpha_j - \alpha_j^*)k(x_i, x_j) \\ \text{subject to} \quad \begin{cases} \sum_{i=1}^{N}(\alpha_i - \alpha_i^*) = 0, \\ 0 \le \alpha_i, \alpha_i^* \le C, \quad i = 1, 2, \ldots, N. \end{cases} \tag{31}$$

$$f(x) = \sum_{i=1}^{N}(\alpha_i - \alpha_i^*)k(x_i, x) + b. \tag{32}$$

[42] As per [95], following the standard derivation of Support Vector Regression, the bias term b in equation (28) is given by

$$b = \frac{1}{|S|}\sum_{s\in S}\left(y_s - \sum_{i=1}^{N}\beta_i\langle x_i, x_s\rangle\right). \tag{29}$$

where S is the set of support vectors; hence b in equation (31) is given by

$$b = \frac{1}{|S|}\sum_{s\in S}\left(y_s - \sum_{i=1}^{N}\beta_i k(x_i, x_s)\right) \tag{30}$$

With the choice of an appropriate kernel function (the details on the conditions the kernel function need to satisfy are outside of the scope of this thesis), the dual optimisation problem given by equation (31) is solved, this then allows estimates of $f(x)$ to be computed via equation (32) (and (30)). As per [95] and [86], a selection of kernel functions include:

$$k(x_i, x_j) = \langle x_i, x_j \rangle = x_i^T x_j \tag{33}$$

$$k(x_i, x_j) = \left(x_i^T x_j + c\right)^p \tag{34}$$

$$k(x_i, x_j) = \exp\left(-\frac{\|x_i - x_j\|^2}{2\sigma^2}\right) \tag{35}$$

The kernel in equation (33) corresponds to the earlier linear regression case. Equations (34) and (35) are referred to as the polynomial and Gaussian (or Radial Basis Function) kernels, respectively. These kernels allow SVMs to model non-linear relationships by implicitly mapping the data into higher-dimensional feature spaces, without requiring explicit computation of those mappings. Parameters such as the degree p and offset c in the polynomial kernel, or the bandwidth σ in the Gaussian kernel, influence the form of the kernel function and are typically selected through model tuning.

A subjective view of the strengths of SVMs is again their broad applicability in both regression and classification contexts. They have the capability to incorporate nonlinearities in data through kernel methods and they have featured in prior EPF literature. However, SVMs present challenges in terms of interpretability and require careful hyperparameter tuning. Additionally, their theoretical framework is arguably more complex than that of some other modelling approaches discussed in this chapter, making it more conceptually challenging for users to grasp the underlying mechanics.

3.5 Error Metrics

For completeness, we briefly reference common error metrics used in the EPF literature. Given the focus in chapter 4 we describe point forecast error metrics[43], a fuller discussion of these metrics can be found in [38]. For clarity, the point forecast error metric formulas below are presented in the context of the out of sample data (i.e. the test set). We also note that in chapters 4 and 5 we develop custom metrics to explain forecast outcomes and gauge market price inertia.

Given test set target observations $\{y_i\}_{i=1}^N$ and corresponding point forecast estimates $\{\hat{y}_i\}_{i=1}^N$, the error metrics referenced in section 2.4, namely the root

[43] For probabilistic forecast evaluation metrics see [35].

mean squared error (RMSE), mean absolute error (MAE), mean absolute percentage error (MAPE) and symmetric mean absolute percentage error (sMAPE) are given by

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2} \tag{36}$$

$$\mathrm{MAE} = \frac{1}{N}\sum_{i=1}^{N}|y_i - \hat{y}_i| \tag{37}$$

$$\mathrm{MAPE} = \frac{1}{N}\sum_{i=1}^{N}\frac{|y_i - \hat{y}_i|}{|y_i|} \tag{38}$$

$$\mathrm{sMAPE} = \frac{1}{N}\sum_{i=1}^{N}\frac{2|y_i - \hat{y}_i|}{|y_i| + |\hat{y}_i|} \tag{39}$$

While some or all of the above metrics are reported in the EPF literature, our anecdotal industry experience is that the MAE or the RMSE are more commonly referenced while the MAPE is less frequently used (likely this is a result of it being undefined in periods where the price y_i is 0 which appears to be an increasingly common occurrence in short-term European electricity market settings [97, 98]). Another useful and intuitive metric is the *relative mean absolute error (rMAE)* which is calculated by dividing a model's MAE by the MAE of a corresponding naive model. It is given by

$$\mathrm{rMAE} = \frac{\frac{1}{N}\sum_{i=1}^{N}|y_i - \hat{y}_i|}{\frac{1}{N}\sum_{i=1}^{N}|y_i - \tilde{y}_i^{\mathrm{naive}}|} \tag{40}$$

with $\tilde{y}_i^{\mathrm{naive}}$ denoting the point forecast from the naive model for trading period i. While chapter 4 reports MAE values, the presence of different price regimes in the study (see figure 8) makes it useful to also report rMAE alongside the custom metrics developed in our research.

4 Day-Ahead Electricity Market Modelling

4.1 Preamble

This chapter presents a slightly expanded version of the paper referenced in section 1.5, titled *A New Approach to Fundamental Electricity Market Modelling: Exploring Market Dynamics and Speculator Influence*. While the introduction in section 4.2 overlaps somewhat with sections 2.3 and 2.6, the concise overview provides helpful context and preserves narrative flow.

4.2 Introduction

Electricity markets have undergone significant change in recent decades. Factors influencing the change include electricity market liberalisation and the build out of renewable generation to meet climate change targets. While electricity market liberalisation has resulted in a diverse set of market participants [99], the ramp-up in renewable generation has been associated with an increasing volatility in short-term electricity market prices [100, 101, 102]. The increasing complexity and volatility in electricity market prices is an important consideration for market participants as they try to optimise their hedging and operational strategies. Given this context it is unsurprising that the electricity price forecasting (EPF) domain has received increasing attention in the literature in recent decades [11]. When conceptualising the EPF domain it helps to differentiate between the forecast horizon (i.e. short/medium/longer term) and the modelling approach. In this chapter our interest is in short-term forecast horizons with a particular focus on the Day-Ahead (DA) market as it represents one of the primary short-term electricity market trading settings. In the DA market participants submit orders on day D to buy or sell electricity for each hourly delivery period in day D+1.

In this setting a suite of modelling approaches frequently encountered in both the literature and industry, for example [46, 101, 103], is fundamental or cost of production type models . We colloquially refer to this class of models as bottom-up modelling approaches. As the name suggests these models typically aim to forecast the underlying supply (and demand) components. An attractive feature of these bottom-up modelling approaches is that they have the capacity to provide insights into both market and price dynamics, however a key challenge lies in the data availability and the assumptions embedded within the simulations [11]. Both [49] and [50] report that forecast prices from fundamental models tend to be flatter and less volatile than market prices outcomes, a pattern that may stem from the omission or absence of strategic bidding and trading behaviours. Furthermore, [50] notes that *'backtesting results for this type of model class are rather scarce in* *the literature'* and appropriately cautions that *'when assessing policy measures, we hence advise to be aware of the drawbacks of consulting untested or insufficiently backtested models'.* Alternatively, what we refer to as top-down modelling approaches use endogenous and exogenous explanatory variables, which are often at the system-level, in conjunction with statistical/machine-learning/neural network techniques to generate DA price forecasts. For instance, [38] presents 'state of the art' models (one being a *LASSO Estimated Autoregressive, LEAR,* model and the other a *Deep Neural Network, DNN*[44]) and datasets against which DA EPF models can be benchmarked. While these models demonstrate strong price forecasting capabilities, they typically do not address short-term electricity market dynamics.

Elsewhere, papers such as [41] and [51] can be viewed as hybrid approaches combining bottom-up and top-down features to generate DA price forecasts. In [41] the authors utilise a functional data approach to implement a fundamental, or merit order, model which *'assumes that the spot prices at electricity exchanges are based on the marginal generation costs of the last power plant that is required to cover the demand'* whereas in [51] time series techniques are used to model DA sale and purchase curves. While both approaches have the ability to forecast electricity market dynamics, absent is an explainability framework to systematically evaluate differences between forecast and actual outcomes. There is also a wider body of short-term electricity market literature that considers market dynamics, behaviours and price outcomes. Papers such as [42], [44] and [104] use agent-based reinforcement learning to examine how regulatory, structural, or market participant behavioural changes have the potential to influence market outcomes. However, as noted in [11] the application of reinforcement learning approaches in these settings *'generally focus on qualitative issues rather than quantitative results'*.

A further tangential strand of literature relates to the role of financial traders, hereafter referred to as *speculators*, in energy and commodity markets. While not the primary focus of this chapter, this is of interest because part of our analysis examines the contribution of different participant cohorts to forecast performance, and the speculator cohort, whose commercial behaviours are intuitively expected to be dynamic, may have a non-trivial influence on outcomes. Relevant works include [60], [62], [63], and [64]. In [60] the authors review research on the role of financial traders, noting that in some of the literature there is ambiguity both in how such traders are defined and in the metrics used to assess their impact. We address the first source of ambiguity by adopting a transparent classifica-

[44]The DNN implementation corresponds to a feed forward neural network architecture described in section 3.4.

tion framework. Conceptually we define speculators as market participants that neither consume nor generate electricity but instead earn or pay a cash amount based on the spread at which they buy and sell across market timeframes. It is worth noting that while participants with physical assets (e.g. pumped storage, thermal, wind or solar generation) or those representing end-customer demand (e.g. supplier units) may display behaviours that resemble speculation, we do not classify them as speculators. Our speculator label is reserved strictly for units without physical characteristics, whose trading activity is not consistently aligned with either generation or consumption roles[45]. We also recognise that larger firms may comprise portfolios with several participants, some of which may fall into our speculator cohort, even though speculative behaviours can also occur outside this group. We address the second source of ambiguity by employing clear and consistent forecast error metrics, ensuring that the impact of speculators and other participant cohorts is measured in a transparent and comparable way. Beyond these definitional considerations, related empirical work in short-term US electricity markets, including [62], [63], and [64], examines the influence of speculators on outcomes such as market price spreads. Although these papers do not focus on forecasting electricity market prices or dynamics, they are relevant given the growing presence of speculators in short-term electricity markets. The discussion in [62] regarding the complexities of market pricing algorithms is also noteworthy.

Research Questions, Approach, and Contributions In this chapter, to generate forecasts of the DA market we implement a fundamental electricity market modelling approach. In this approach we use statistical, machine learning, and neural network models to forecast individual participants' commercial order data. This forecast data is subsequently aggregated to produce demand (bid) and supply (ask) curves for each trading period. Our primary research questions are:

- Can such an approach capture the characteristic behaviours of the DA market at aggregate, cohort, and individual levels?
- Can we provide a level of explainability into differences or similarities between forecast and actual demand and supply at these levels?
- How does price forecast accuracy compare with other DA EPF benchmark models?

[45] Further details on the cohort classification process are provided in Section 4.4.4.

To address these questions, we present an innovative fundamental short-term electricity market model that overcomes one of the main shortcomings of such approaches, namely strong assumptions embedded within the model. We pair this with an explainability framework, filling a gap in the EPF literature by enabling systematic evaluation of differences between forecast and realised market outcomes. When applied together, our modelling framework reveals that speculators have a significant influence on overall price forecast accuracy, an aspect not directly addressed in existing work. These results, alongside prior research, highlight a trade-off inherent to fundamental approaches: excluding trading behaviour can simplify forecasting but reduces realism, yielding flatter and less volatile price forecasts; including it improves realism but introduces behavioural complexity that can reduce price forecast accuracy (relative to top-down models). These insights are relevant given the continued use of fundamental models in industry, academia, and regulatory settings. By integrating elements from multiple strands of literature, our approach emphasises both price forecast accuracy and the exploration of market behaviours and dynamics, offering deeper insights into the price formation process.

Chapter Structure The remainder of the chapter is organised as follows. Section 4.3 outlines the structure of participant DA order data, the transformation into a model ready format (including a robustness-check variant), and the associated information loss analysis. Section 4.4 summarises the forecasting process and implementation for individual, aggregate, and cohort analyses; section 4.5 presents results; section 4.6 interprets them in light of the research questions and section 4.7 concludes. Supplementary details are provided in the appendices.

4.3 Participant Order Data

In section 4.3.1 we describe the participant DA order data (to familiarise the reader with the market information accessible in practice) while in section 4.3.2, we provide the motivation as to why it is a transformed version of this data, rather than the raw participant order data itself, that we use in model training and forecasting. The transformation approaches are then detailed in section 4.3.3, while section 4.3.4 presents metrics that quantify the information loss associated with these transformations.

4.3.1 Order data example

As explained in section 3.3 and as detailed in [82] a Simple Order relates to a single trading period and a specified participant. It must contain at least two and not more than 256 price quantity, PQ, pairs which represent the price and quantity of electricity for sale or purchase in that trading period. The prices must fall within the minimum (-€500/MWh) and maximum (€3000/MWh[46]) price range and each Simple Order must contain at least one price at the minimum price point and one price at the maximum price point. In the Irish electricity market Simple Orders are stepwise functions with either an incremental or decremental quantity at each price point, the prices associated with sell (buy) orders are monotonically increasing (decreasing). The market convention is that negative (positive) quantities represent sell (buy) orders. With this definition in mind, consider the following hypothetical DA Simple Order for a participant, with ϕ_P and ϕ_Q denoting its stepwise price and stepwise quantity vectors respectively:

$$\phi_P = \text{Stepwise Prices (€/MWh)} = [-500, p_1, p_1, p_2, p_2, p_3, p_3, 3000] \qquad (41)$$

$$\phi_Q = \text{Stepwise Quantities (MW)} = [0, 0, -q_1, -q_1, -q_2, -q_2, -q_3, -q_3] \qquad (42)$$

Reading the vectors from left to right it can be seen that the participant is not willing to sell any quantity at the minimum price. It is willing to sell $|q_1|$ at a price of p_1, an additional $|q_2 - q_1|$ at a price of p_2 and finally another $|q_3 - q_2|$ at a price of p_3. It does not offer any supplemental quantity at the maximum price point. It is clear that the salient information in this hypothetical Simple Order are at the points where there is a change in quantity. In this example, without loss of information, a more compact representation of the stepwise order is given by the corresponding *piecewise* vectors; namely the piecewise ask price vector, ω_P^a, given by $[p_1, p_2, p_3]$ and the piecewise ask quantity vector, ω_Q^a, given by $[-q_1, -q_2, -q_3]$, where the a superscript denotes that it is the ask component. In this hypothetical example the participant is only selling and we use a convention that its corresponding piecewise bid price vector, ω_P^b, and piecewise bid quantity vector, ω_Q^b, are empty. In figure 15 we plot this hypothetical example.

In the Irish electricity market participants can also utilise *Complex Order* types (as of January 2023 Complex Orders were replaced by *Scalable Complex Orders*). While, as described earlier, a Simple Order specifies a price and quantity pair for a single trading period and is independently cleared, a Complex Order links one or more simple sell orders with additional conditions or constraints (e.g. minimum revenue or execution conditions across multiple periods). Complex Order types are only available for sell orders. In this chapter we restrict ourselves to focusing on the Simple Order component of such order types. On average 92% (8%) of the participant sell orders are Simple (Complex) order types. For the

[46] Subsequently amended to €4000/MWh in May 2022 as per [105] and to €5000/MWh in August 2022 [106].

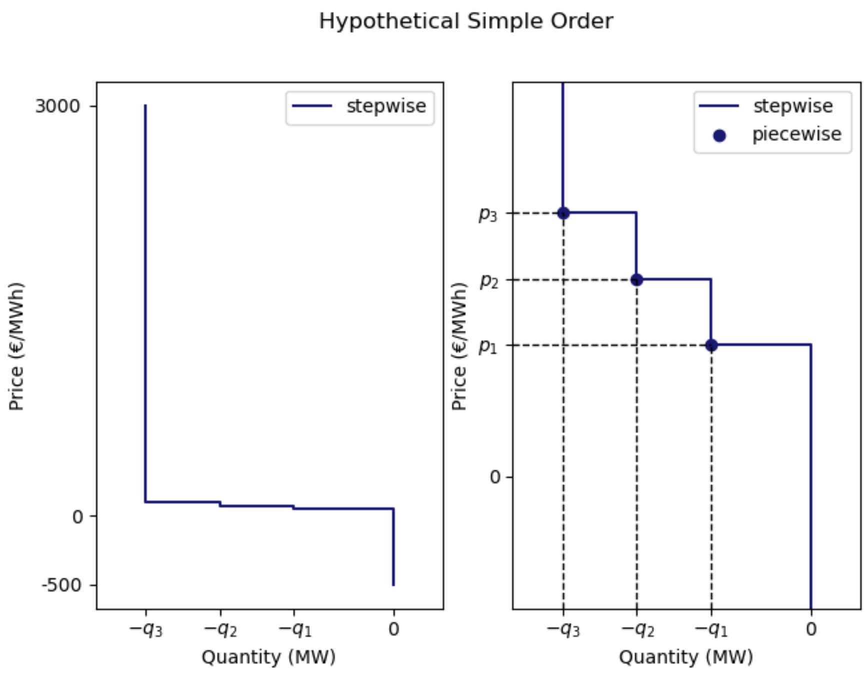

Figure 15: Plot on lhs is a line plot of the stepwise Simple Order data from eqns (41) and (42); plot on rhs presents similar information except we zoom in on the price axis, we also include a (scatter) plot of the corresponding piecewise ask data.

quantities sold, that is sell orders that are matched, circa 65% (35%) are Simple (Complex) order types.

4.3.2 Transformation motivation

From section 4.3.1 it is clear that participants are afforded a significant degree of flexibility in how they construct their order data. This is also evident if we consider how over the January 2021 to December 2022 timeframe 12% of market participants change the dimension, or number of points in their Simple Order data, 150 or more times[47]. A number of speculators change the dimensionality at even greater frequencies. An additional challenge is clear when we consider the structure, or granularity, of the Irish electricity market Simple Order data. While on average 88% of the stepwise order data consists of between 2 to 6 points, circa 8% of the stepwise order data is comprised of between 6 to 16 points and an average of approximately 3% (sometimes spiking to 7%) of the stepwise order data consists of greater than 16 points[48].

Thus, when forecasting participant order data we face the twin challenge of first trying to predict the dimensionality of the order and second the granular

[47] 23% of market participants show 10 or more changes in the dimensionality of their order data over the corresponding timeframe.

[48] For reference, on average 1% of the stepwise order data contains 30 or more points

order details (i.e. prices and quantities). With the goal of reducing the modelling complexity our approach is to instead take the participant order data and transform it into vectors of a fixed dimension. While the transformed data is more amenable to analysis/modelling, it may lead to a loss of information at the aggregate level (which can be quantified).

4.3.3 Transformation Approaches

In this section, we introduce two approaches for transforming participant stepwise order data into a model-ready, fixed-dimension format and present the underlying algorithms for completeness. Our primary focus in this chapter is on Approach 1, which preserves as much granularity from the original order data as possible[49]. To test robustness, we also designed Approach 2, which focuses on quantities around the market price[50].

Before outlining Approach 1, we restate the key notation and introduce a few additional symbols. For readability, we omit the subscripts i (participants) and j (trading periods), though they will appear explicitly in the algorithms later in this chapter. Let ϕ_P and ϕ_Q denote the participant's stepwise price and quantity vectors, which can be decomposed into piecewise bid vectors ω_P^b and ω_Q^b and piecewise ask vectors ω_P^a and ω_Q^a, with lengths u and v respectively (i.e. u and v are the dimensions of the bid and ask vectors). We introduce n as the target transformation dimension to which the piecewise order data will be mapped. The resulting transformed vectors are τ_P^b, τ_Q^b, τ_P^a, and τ_Q^a, each lying in $\mathbb{R}^n$.

Approach 1

1. Split the participant's stepwise order data into separate piecewise bid and ask vectors.

2. Select n, the common dimension for all transformed vectors.

3. For each side (bid/ask):

   - If $n \geq$ the length of the corresponding piecewise vector (u or v), append zeros to the piecewise price and quantity vectors until dimension n is reached.

[49] In some cases, Approach 1 produces surplus elements containing no information (e.g. τ_P^b and τ_Q^b may be entirely zero for certain participants); the impact of this is examined in a supplemental modelling scenario in table 3. Variants of Approach 1, such as omitting the piecewise vector bid/ask split, are also possible but are left for future research.

[50] Scenarios using Approach 1 transformations are reported in section 4.4.1 (table 2), while additional scenarios with Approach 2 appear in table 3.

- If $n <$ the length of the corresponding piecewise vector, divide the quantity range into $n-1$ equal intervals and linearly interpolate the corresponding prices at these quantity points using the piecewise data.

The output is a $4n$-dimensional vector $[\tau_P^b, \tau_Q^b, \tau_P^a, \tau_Q^a]$ for each participant and period. To illustrate the approach, let us consider the hypothetical example from section 4.3.1. With $\omega_P^a = [p_1, p_2, p_3]$, $\omega_Q^a = [-q_1, -q_2, -q_3]$, $\omega_P^b = \omega_Q^b = []$ and a choice of $n = 5$, then:

$$\tau_P^b = \tau_Q^b = [0, 0, 0, 0, 0], \quad \tau_P^a = [p_1, p_2, p_3, 0, 0], \quad \tau_Q^a = [-q_1, -q_2, -q_3, 0, 0].$$

If instead the dimensionality of a participant's piecewise vector exceeds n (e.g. $v > n$), interpolation is applied. Figure 16 provides an example of an Approach 1 transformation with $v > n$, where some information is lost[51].

[51]Stepwise vectors in figure 16 are:

$$\begin{aligned}
\phi_P &= [-500, 10.92, 10.92, 54.83, 54.83, 64.91, 64.91, 76.77, 76.77, 101.81, 101.81, 3000] \\
\phi_Q &= [0, 0, -110.5, -110.5, -211, -211, -321.6, -321.6, -381.9, -381.9, -409, -409].
\end{aligned} \tag{43}$$

This results in piecewise vectors: $\omega_P^b = \omega_Q^b = []$, $\omega_P^a = [10.92, 54.83, 64.91, 76.77, 101.81]$ and $\omega_Q^a = [-110.5, -211, -321.6, -381.9, -409]$. Applying Approach 1 with $n = 4$ gives: $\tau_P^b = \tau_Q^b = [0, 0, 0, 0]$, $\tau_P^a = [10.92, 54.39, 63.81, 101.81]$ and $\tau_Q^a = [-110.5, -210.0, -309.5, -409]$.

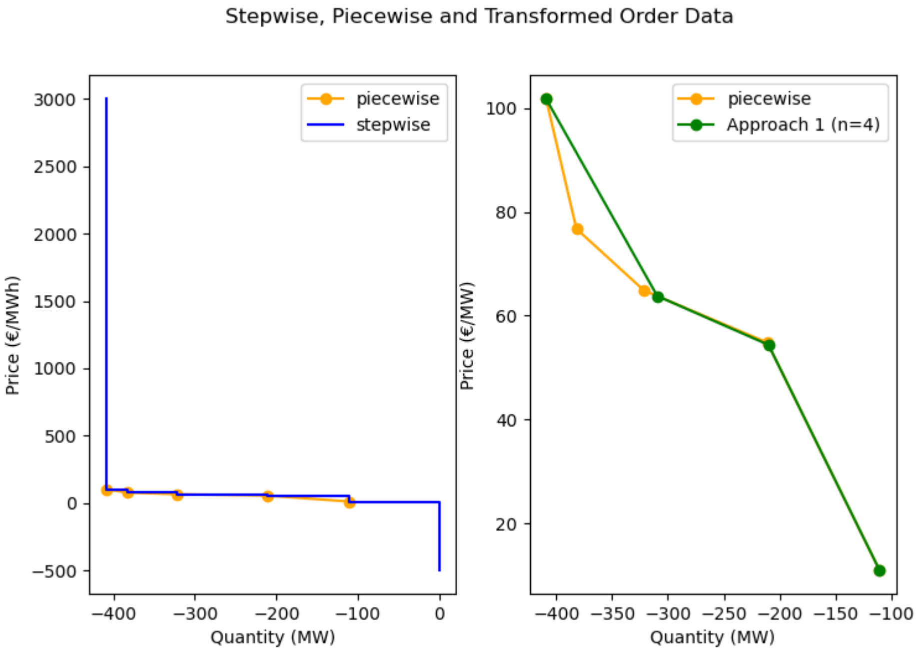

Figure 16: Plot on the lhs shows a lineplot of the stepwise and corresponding piecewise sell order data for a thermal generator participant in a single trading period in February 2021. Plot on the rhs displays lineplots of the same piecewise sell order data and the corresponding transformed piecewise sell order data where the transformation is done under Approach 1 with $n = 4$. The piecewise sell vectors are of dimension 5 and hence under the transformation there is a loss of information given the choice of $n = 4$. The piecewise buy vectors are empty and are ignored for display purposes.

Before outlining Approach 2, we introduce additional notation (for readability, we again omit the participant i and trading period j subscripts here). Let $\chi_P \in \mathbb{R}^n$ denote the *price quantile* vector, containing n market price quantile values derived from recent I-SEM DA market prices[52]. This vector is identical across all participants for a given trading period. Let $\chi_Q \in \mathbb{R}^n$ denote the *quantity quantile* vector for a given participant, obtained by mapping each price in χ_P to the corresponding quantity in the participant's stepwise order data. The vectors $\chi_P, \chi_Q \in \mathbb{R}^n$ constitutes the transformed order data for that participant under Approach 2[53].

Approach 2

1. Select n, the target vector dimension (e.g. $n = 5$).

2. Using hourly I-SEM DA market price observations for that delivery hour and its neighbouring hours over the past X days, compute n quantiles to form χ_P. For example, if $n = 5$, one might choose the $0^{\text{th}}, 25^{\text{th}}, 50^{\text{th}}, 75^{\text{th}}$, and 100^{th} percentiles of observed DA prices.

3. For each price $p \in \chi_P$, locate the interval in the participant's stepwise price vector ϕ_P that contains p, and take the corresponding quantity from ϕ_Q to form χ_Q.

Using the hypothetical example from Section 4.3.1, and a choice of

$$\chi_P = [P^{\text{DA},0}, P^{\text{DA},25}, P^{\text{DA},50}, P^{\text{DA},75}, P^{\text{DA},100}],$$

where $P^{\text{DA},k}$ denotes the kth price quantile of the DA market, the approach is illustrated in figure 17. When these price quantiles are mapped to the participant's stepwise order data, the first three quantiles correspond to a quantity of 0, while the last two correspond to $-q_1$ and $-q_3$, giving

$$\chi_Q = [0, 0, 0, -q_1, -q_3].$$

The impact of different choices of n are considered in section 4.5.1[54].

[52] The procedure for computing these quantiles is described in section 4.5.1.

[53] Since χ_P is identical across participants for a given trading period, the object of interest in our modelling and forecasting is effectively χ_Q.

[54] Extensions of Approach 2, similar to those in [51], include: (a) using a bucketing approach to determine the price quantile vector, and (b) further subdividing each bucket once the participant's quantity per bucket has been determined. These are left for future research.

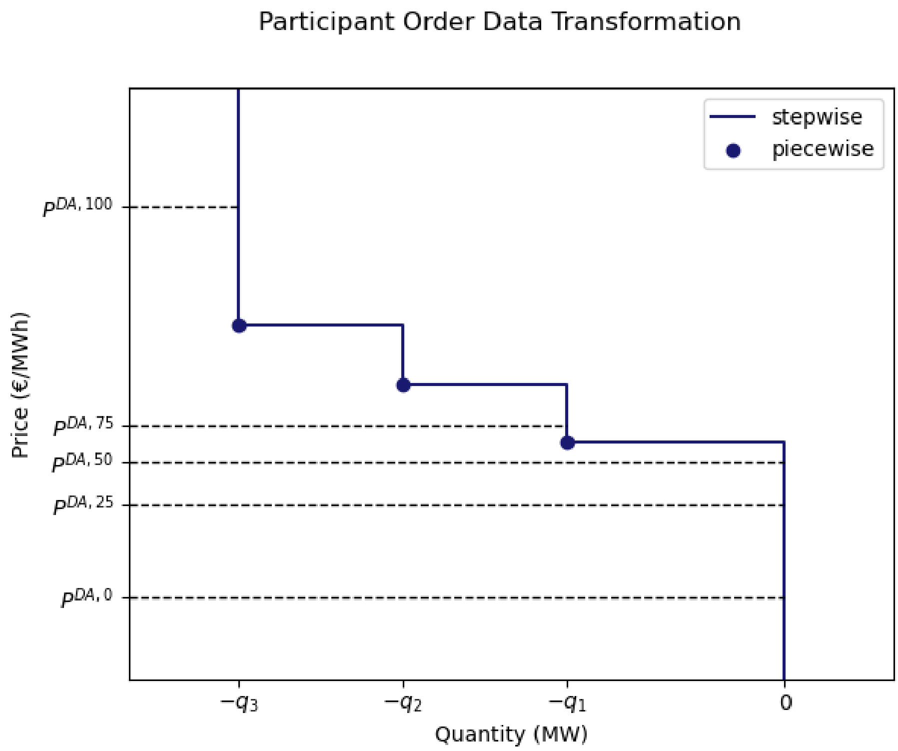

Figure 17: Lineplot (scatterplot) of the hypothetical stepwise (piecewise) Simple Order data in which we zoom in on the price axis. As per the Approach 2 example, we draw horizontal dashed lines at the price quantile values $P^{\mathrm{DA},0}$, $P^{\mathrm{DA},25}$, $P^{\mathrm{DA},50}$, $P^{\mathrm{DA},75}$, and $P^{\mathrm{DA},100}$ to derive the participant's corresponding quantity quantile vector $\chi_Q = [0, 0, 0, -q_1, -q_3]$.

Algorithms For completeness, algorithms 1, 2 and 3 present the Approach 1 implementation details (algorithms 2 and 3 handle the transformation under a choice of $n = 1$). It is worth noting that in Algorithm 1, when a participant's piecewise order data vectors have a dimension greater than n, the transformed piecewise order quantity vectors contain redundant quantity information. Specifically, only the maximum and minimum quantities are necessary as the remaining transformed quantities are positioned at equally spaced intervals of width h. While the algorithm could be modified to retain only the participant's minimum and maximum quantities in the transformed piecewise order data, we do not adopt this approach as it would lead to a loss of information for participants whose piecewise order data vectors have a dimension less than n and whose piecewise quantities are not positioned at equally spaced intervals. Algorithm 4 presents the Approach 2 implementation details. Finally, we note that the if statement in the algorithms is a convention we apply to handle situations in which (a) participants only submit orders for particular portions of the trading day or (b) scenarios in which a participant does not submit any order in the trading day.

Algorithm 1: Populating $\tau^b_{Pij}, \tau^b_{Qij}, \tau^a_{Pij}, \tau^a_{Qij}$ for $n > 1$

Data: Participant i, trading period j.

Input : Stepwise vectors, ϕ_{Pij} and ϕ_{Qij}.

Output: Transformed piecewise vectors, τ^b_{Pij}, τ^b_{Qij}, τ^a_{Pij}, τ^a_{Qij} $\in \mathrm{I\!R}^{\mathrm{n}}$

if $\phi_{Pij} = \phi_{Qij} = [\,]$ **then**

  if *Participant submits an order for* $\geq$ *1 periods in the trading day* **then**

    $\tau^b_{Pij} = \tau^b_{Qij} = \tau^a_{Pij} = \tau^a_{Qij} = [0_1, \ldots, 0_n]$

  else

    $\tau^b_{Pij} = \tau^b_{Qij} = \tau^a_{Pij} = \tau^a_{Qij} = []$

  end

else

  1. Convert stepwise order into corresponding piecewise vectors

  $\phi_{Pij}, \phi_{Qij} \rightarrow \omega^b_{Pij}, \omega^b_{Qij}, \omega^a_{Pij}, \omega^a_{Qij}$

  where $\omega^b_{Pij}, \omega^b_{Qij} \in \mathrm{I\!R}^{\mathrm{u}}$ and $\omega^a_{Pij}, \omega^a_{Qij} \in \mathrm{I\!R}^{\mathrm{v}}$.

  2. Handling $\omega^b_{Pij}, \omega^b_{Qij}$ for $u \leq n$:

  If u = 0 **then** $\tau^b_{Pij} = \tau^b_{Qij} = [0_1, \ldots, 0_n]$

  If u = n **then** $\tau^b_{Pij} = \omega^b_{Pij}$ and $\tau^b_{Qij} = \omega^b_{Qij}$

  If $0 < u < n$ **then** $\tau^b_{Pij} = [\omega^b_{Pij}, 0_{u+1}, \ldots, 0_n]$ and $\tau^b_{Qij} = [\omega^b_{Qij}, 0_{u+1}, \ldots, 0_n]$

  3. Handling $\omega^b_{Pij}, \omega^b_{Qij}$ for $u > n$:

  Calculate

$$h = \frac{\omega^b_{Qiju} - \omega^b_{Qij1}}{n-1}$$

  Then

  $\tau^b_{Qij} = [\omega^b_{Qij1}, \omega^b_{Qij1} + h, \ldots, \omega^b_{Qij1} + (n-1) * h]$.

  Next, for each quantity in τ^b_{Qij} linearly interpolate the corresponding price using the ω^b_{Pij} and ω^b_{Qij} vectors.

  5. Repeat steps 3 and 4 using $\omega^a_{Pij}, \omega^a_{Qij}$ and v.

end

Algorithm 2: Populating $\tau^b_{Pij}, \tau^b_{Qij}, \tau^a_{Pij}, \tau^a_{Qij}$ for $n = 1$

Data: Participant i, trading period j.

Input : Stepwise vectors, ϕ_{Pij} and ϕ_{Qij}.

Output: Transformed piecewise vectors, τ^b_{Pij}, τ^b_{Qij}, τ^a_{Pij}, $\tau^a_{Qij} \in \mathbb{R}^1$.

if $\phi_{Pij} = \phi_{Qij} = [\,]$ **then**

  if *Participant submits an order for* $\geq$ *1 periods in the trading day* **then**

    $\tau^b_{Pij} = \tau^b_{Qij} = \tau^a_{Pij} = \tau^a_{Qij} = [0]$

  else

    $\tau^b_{Pij} = \tau^b_{Qij} = \tau^a_{Pij} = \tau^a_{Qij} = []$

  end

else

  1. Convert stepwise order into piecewise bid and ask orders
  $\phi_{Pij}, \phi_{Qij} \rightarrow \omega^b_{Pij}, \omega^b_{Qij}, \omega^a_{Pij}, \omega^a_{Qij}$
  where $\omega^b_{Pij}, \omega^b_{Qij} \in \mathbb{R}^{\mathrm{u}}$ and $\omega^a_{Pij}, \omega^a_{Qij} \in \mathbb{R}^{\mathrm{v}}$.

  2. Handling $\omega^b_{Pij}, \omega^b_{Qij}$:

  If u $= 0$ **then** $\tau^b_{Pij} = \tau^b_{Qij} = [0]$

  If u $= 1$ **then** $\tau^b_{Pij} = \omega^b_{Pij}$ and $\tau^b_{Qij} = \omega^b_{Qij}$

  If $u > 1$ **then** $\tau^b_{Pij}, \tau^b_{Qij}$ = Algorithm3($\omega^b_{Pij}, \omega^b_{Qij}$)

  3. Repeat step 2 using $\omega^a_{Pij}, \omega^a_{Qij}$ and v.

end

Algorithm 3: Converting $[p_1, \ldots, p_m]$ and $[q_1, \ldots, q_m]$ into price quantity pair $[p^*]$ and $[q^*]$.

Input : Price and quantity vectors, $[p_1, \ldots, p_m]$ and $[q_1, \ldots, q_m]$.

Output: $[p^*]$ and $[q^*]$ with both vectors $\in \mathbb{R}^1$.

Step 1: $[p^*]$ is given by

$$p^* = \frac{p_1 q_1 + \sum_{i=2}^{m}(q_i - q_{i-1})p_i}{q_m}$$

Step 2: If p^* from Step 1 equals any of the $[p_1, \ldots, p_m]$ then q^* equals the corresponding element from $[q_1, \ldots, q_m]$. Otherwise, determine the interval $[p_j, p_{j+1}]$, which contains p^*. Then $[q^*]$ is given by

$$q^* = q_j + \frac{(p^* - p_j)(q_{j+1} - q_j)}{p_{j+1} - p_j}$$

Algorithm 4: Populating χ_{Qij} given $\chi_{Pj} \in \mathbb{R}^{\mathrm{n}}$

Data: Participant i, trading period j.

Input : Stepwise vectors ϕ_{Pij}, ϕ_{Qij}; price quantile vector χ_{Pj} with $\chi_{Pj} \in \mathbb{R}^{\mathrm{n}}$.

Output: Quantity quantile vector $\chi_{Qij} \in \mathbb{R}^{\mathrm{n}}$

if $\phi_{Pij} = \phi_{Qij} = [\,]$ **then**

    if *Participant submits an order for* $\geq$ *1 periods in the trading day* **then**

        $\chi_{Qij} = [0_1, \ldots, 0_n]$

    else

        $\chi_{Qij} = []$

    end

else

    1. Set $\chi_{Qij} = []$.

    2. For each price p in χ_{Pj}:

    Determine the index (indices) of ϕ_{Pij} which contain p.

    q is the value at the corresponding index in ϕ_{Qij}.

    Append q to χ_{Qij}.

end

4.3.4 Transformation error metrics

The error metrics described in this section are based on the following high-level approach; we take the transformed participant order data and aggregate it to construct what we call transformed bid and ask curves, we then compare these transformed bid and ask curves to the published bid and ask curves. We calculate and compare price and quantity intersection points for the transformed and published bid and ask curves, we also consider area under the curve metrics (in later sections we will use similar area under the curve metrics to derive insights into good/bad forecast performance).

As explained in section 3.3 for each DA auction the market operator publishes a BidAskCurve file containing a stepwise monotonic increasing (decreasing) and anonymised view of the sell (buy) orders for each trading period in the auction. From this file, using the bid and ask curves for a particular trading period, j, we calculate the price and quantity values, denoted by P_j^{DA} and Q_j^{DA} respectively, at which the curves intersect. For the same trading period, taking all participants orders and transforming them using one of the approaches outlined in section 4.3.3, we construct the corresponding (stepwise) transformed bid and ask curve. Using these transformed curves we determine the price and quantity values at which they intersect, henceforth denoted as P_j^{TDA} and Q_j^{TDA}. Subsequently, we calculate mean absolute error, MAE, metrics in the usual way:

$$\text{Price MAE} = \frac{1}{N}\sum_j |P_j^{\text{DA}} - P_j^{\text{TDA}}| \tag{44}$$

$$\text{Quantity MAE} = \frac{1}{N}\sum_j |Q_j^{\text{DA}} - Q_j^{\text{TDA}}| \tag{45}$$

Here the summation is over all trading periods in the timeframe of interest and N represents the number of trading periods in that time range. We choose the MAE metric as it is one of the most commonly used metrics in the EPF literature, alternative choices are possible. For area metrics, let us consider the scenario in which we want to calculate the area under the ask curve in trading period j, denoted as $A_j^{a,\text{DA}}$, on some price interval $[lower, upper]$. We first partition the price axis into equal spaced intervals, for example a partition of width €1/MWh is defined by

$$\left\{[x_0, x_1], ..., [x_{n-1}, x_n]\right\} = \left\{[lower, lower + 1], ..., [upper - 1, upper]\right\} \tag{46}$$

Using this partition, we use the Riemann sum to approximate the area under the ask curve:

$$A_j^{a,DA} \approx \sum_{i=0}^{n-1} f(x_i)\Delta x_i = \sum_{i=0}^{n-1} f(x_i) \tag{47}$$

The value $f(x_i)$ is the quantity at price point x_i which is trivial to determine from the published ask curve. In a similar manner, we can take the transformed stepwise ask curve for trading period j, and calculate the area under the transformed ask curve, $A_j^{a,\text{TDA}}$. Subsequently, we can calculate an ask curve MAE related metric via

$$\text{Ask Curve MAE} = \frac{1}{N}\sum_j |A_j^{a,\text{DA}} - A_j^{a,\text{TDA}}| \tag{48}$$

The process can be repeated for the bid curve, that is calculate $A_j^{b,\text{DA}}$ and $A_j^{b,\text{TDA}}$ and in turn the corresponding Bid Curve MAE.

4.4 Modelling

An overview of the end-to-end process is given in figure 18. In the first step we specify the modelling approach that applies to the participants. In step 2 we loop through each participant and generate forecasts of their DA order data over the horizon of interest. In step 3 we combine the individual participant forecasts to get an aggregate view of the forecast bid and ask curves. In step 4 we examine cohorts of market participants as it helps provide additional insights into the forecast performance. We repeat steps 1 through 4 several times to investigate the impact of different modelling approaches on the overall forecast performance. In the next subsection, we describe step 1 in detail setting out the modelling approach (problem formulation, target data, explanatory variables, and models) and then explaining how variations in aspects such as calibration length and explanatory variable subsets give rise to the different modelling scenarios summarised in table 2.

Step 1: Specify the modelling approach

↓

Step 2: Generate individual participant forecasts

↓

Step 3: Combine participant forecasts

↓

Step 4: Cohort forecast analysis

Figure 18: The steps in the bottom-up modelling approach.

4.4.1 Specify the modelling approach

Problem Before 11am on day D participant i submits a distinct order to the Irish electricity market DA auction for each hourly trading period in the interval [11pm D, . . . , 10pm D+1] . Mimicking real world conditions (i.e. ensuring to use only information that is available on/before 11am on day D) we want to forecast participant i's order data, for each trading period $j \in$ [11pm D, . . . , 10pm D+1].

Target Data For the reasons outlined in section 4.3.2 it is the transformed participant piecewise order data that we forecast. We therefore have the choice as to the transformation approach we apply in addition to the value of n (the dimension of the transformed vectors). If for example we choose Approach 1, then as previously explained for participant i in trading period j its stepwise order data, ϕ_{Pij} and ϕ_{Qij}, is converted into transformed piecewise bid and ask vectors τ^b_{Pij}, τ^b_{Qij}, τ^a_{Pij},$\tau^a_{Qij} \in \mathbb{R}^n$. The vector given by $[\tau^b_{Pij}, \tau^b_{Qij}, \tau^a_{Pij}, \tau^a_{Qij}] \in \mathbb{R}^{4n}$ constitutes the target data[55].

[55] In subsequent sections we also examine the impact of subdividing the transformed vectors obtained under Approach 1 into $[\tau^b_{Pij}, \tau^b_{Qij}]$ and $[\tau^a_{Pij}, \tau^a_{Qij}]$ and then forecasting them separately.

Explanatory Data To ensure the forecasting process is valid, we use only explanatory variables that would have been available at the time each forecast was made, and all datasets used in the analysis are publicly available. The explanatory variables considered here follow the main categories commonly used in the EPF literature (e.g. demand and renewable forecasts, lagged prices, fuel and carbon costs, and calendar effects; for examples see [56, 57, 58, 39]), while also incorporating participant-level information. The selected variables were further shaped by industry knowledge/experience and exploratory analysis[56]. The explanatory data includes both forward and backward looking variables. Some relate to the wider system or environment, while others are participant specific. Details include:

- (Forward looking, system) W_j, D_j, G_j, C_j: with j representing the trading period, W and D denote the I-SEM wind and demand system-level forecasts in MW. The variables G and C denote the spot market gas and carbon prices respectively, with G measured in euro cents per therm (€c/th) and C measured in euros per metric tonne (€/t)[57] (and are lagged by two and one day to ensure alignment with the principle of using only information available at the forecast time).

- (Backward looking, system) P^{DA}_{j-lag}, P^{B}_{j-lag},W_{j-lag}, D_{j-lag}, G_{j-lag}, C_{j-lag}, A_{j-lag}, I_{j-lag}: with the subscript $j-lag$ again indicating the trading period, P^{DA} and P^{B} denote the Irish electricity market DA and balancing market prices respectively. Variable A, measured in MW, is the thermal generation fleet availability[58]. I, also in MW, represents the electricity flow across the interconnectors, Moyle and EWIC, joining Great Britain and Ireland. Of the eight variables, the first four use lags of $\{24, 48, 72, 168\}$ hours, the latter four only use lags of $\{24, 168\}$ hours; however we note that P^{B}_{j-24} and I_{j-24} are not observable at the time the forecast is made for the entirety of the forecast horizon and hence are excluded as explanatory variables.

- (Backward looking, participant) $E^{\text{auction}}_{i,j-lag}$, $M_{i,j-lag}$, $\mathrm{T}_{i,j-lag}$: subscript $i, j-lag$ denotes that the information relates to participant i, trading period $j - lag$ with $lag \in \{48, 168\}$. We note that 48 hours is the minimum lag

[56]The exploratory analysis was a mixture of both filter (i.e. correlation) and embedded (e.g. feature importance and *L1* regularisation) methods.

[57]Gas and carbon prices are at a daily granularity and hence the time series values are the same for each $j \in$ [11pm D-1, . . . , 10pm D]. We use the *NBP SAP* price as a proxy for the Day-Ahead gas price and the nearest/open EUA December contract as an estimate of the carbon spot price.

[58]We sum all Irish electricity market DA sell orders in trading period $j - lag$ where the participant is a generator with a registered fuel type of either multi-fuel, gas or coal. Alternative source files, and hence approaches, can be used to estimate availability.

given the data publication schedule. E denotes the executed quantity, that is the amount bought or sold, by the participant at auction with auction $\in \{\text{DA}, \text{IDA1}, \text{IDA2}, \text{IDA3}\}$. M represents the metered data. T denotes the target data. We apply a convention whereby if the data does not exist for a participant (e.g. some participants will not have any M values), we use a value of 0.

- (Categorical) Hour_j, Weekday_j: given trading period j, we apply a one-hot encoding to the hour and weekday values.

We emphasise that although the choice of explanatory variables is influential, it is not the main focus of the paper. We also note that for certain modelling approaches we only utilise a subset of the explanatory variables, as discussed further below.

Modelling Taxonomy Given the nature of participant i's target data it is clear we are dealing with a multitarget or multioutput regression (i.e. simultaneous prediction of more than one variable or output). As discussed in Section 2.3.2, participant i forecasts can be generated using either separate models for each hour (i.e. 24 models) or a single model for all hours. We adopt the latter, implementing a univariate approach as defined in [59][59].

Models We performed an initial investigation on a wider cross section of models and subsequently focused on the subset listed below. The rationale for this selection, along with the choice of model parameters, calibration periods, and explanatory variable subsets is discussed further below. In the interests of brevity we refer the reader to the model overviews presented in section 3.4; additional details include:

- Naive: We use a lag of 48 hours given this is the most recent publicly available information relating to participant i but other options are possible; for example a lag of 168 hours for demand participants may be preferable.

- Least Absolute Shrinkage and Selection Operator, LASSO: In python we use the *Scikit-learn* library (the *linear_model.Lasso* class).

- Random Forests, RF: We use the *RandomForestRegressor* class from Scikit-learn's *ensemble* module.

[59]That is, we train a single model for participant i and use it to generate 24 distinct forecasts of participant i's transformed order data using unseen test data, one for each $j \in [\text{11pm D}, \ldots, \text{10pm D+1}]$.

- Support Vector Machines, SVM: In python, we use the *SVR* class (from *sklearn.svm*) in conjunction with the *MultiOutputRegressor* class (from *sklearn.multioutput*).

- Feedforward Neural Network, FFNN: In Python we use the Tensorflow and Keras libraries in the implementation, we use the rectified linear unit (ReLU) as the activation function, in addition we apply early stopping using the most recent seven days worth of training data as the validation dataset.

Derivation of Modelling Scenarios The goal of this step was to define a manageable set of modelling scenarios to be used in the main experiment. To construct these scenarios, we examined how model choice, calibration window length, and explanatory variable selection affect forecast accuracy. We did so using an out-of-sample model selection process (where models are evaluated on data not used for estimation) with walk-forward validation (where models are repeatedly re-estimated on a rolling historical window and evaluated on the subsequent period). Candidate models were evaluated on the same set of sampled participants with held-out 2023 data, and the best on-average performing configurations were identified. These configurations form the basis of the modelling scenarios summarised in table 2. The process was as follows:

- Select a transformation approach for generating participant target data. From section 4.5.1 table 4 we observe that under Approach 1 values of $n \leq 3$ lead to non-trivial information loss; we therefore chose a value of $n = 4$ for Approach 1.

- Identify the subset of modelling approaches by:
  - Selecting one model at a time.
  - Sampling one or more participants from each cohort (the process for cohort identification is outlined in section 4.4.4); this same set of sampled participants is used for all modelling approaches when determining the chosen settings.
  - Using 2023 data (i.e. ensuring no overlap with the study horizon referenced in section 4.4.2):
    1. Perform a grid search to select model hyperparameters, number of calibration days and explanatory variable subsets.
    2. Run a walk-forward validation process for each sampled participant using the selected settings.

- Determine the configuration that performs best on average across sampled participants, while aiming to avoid overfitting. This "best on average" configuration is reflected in table 2.

As supplementary material appendix D details model hyperparameter settings. The process leading to table 2 was also shaped by considerations outlined in the next section, namely (a) the large number of participants which discourages an exhaustive search for a tailored configuration for each individual participant and (b) the forecast horizon which motivates a purposeful limitation on the number of modelling approaches analysed.

Table 2: Table summarising the modelling approaches used. The target data for each approach is the participant's transformed piecewise order data from Approach 1 with $n = 4$. The *Explanatory* row shows whether all or only some explanatory variables are used in the multitarget regression. *Categorical* indicates if one-hot encoded Hour and Weekday variables are included. *No. days* refers to the intended number of days of data used for model training; in practice, the full amount may not always be available, in which case the model is trained on the data that are available. If fewer than 14 days of training data are available, a naive model is used instead.

No.	0	1	2	3	4	5	6	7
Model	Naive	RF	RF	LASSO	LASSO	SVM	SVM	FFNN
Explanatory		All	All	Subset	Subset	Subset	Subset	Subset
Categorical		True	True	True	False	False	False	False
No. days		75	30	120	30	75	120	120

Finally, to assess the impact of alternative transformations and other factors on forecast accuracy, and to further ensure robustness, we analysed the additional modelling approaches presented in table 3. These supplemental scenarios were informed by the main experiments in table 2, which form the core of our analysis.

Table 3: Building on a selection of baseline modelling approaches, namely LASSO (No.4), FFNN (No.7), SVM (No.6), and RF (No.1) from table 2, this table explores alternative modelling approaches. These include: (a) different transformation methods; (b) splitting the target data under Approach 1 and modelling bid and ask components separately; (c) modifying a single hyperparameter; (d) adjusting model settings; (e) forecasting speculators as a group rather than individually and (f) changing explanatory variables. Unlike the derivation of table 2, no grid search was performed to optimise combinations of explanatory variables, hyperparameters, or other settings.

No.	Model	Target	Comment
8	LASSO	Approach 1, $n = 5$	(a) Increasing the dimensionality
9	LASSO	Approach 2, $n = 25$	(a) Alternative transformation
10	LASSO	Approach 2, $n = 15$	(a) Alternative transformation
11	FFNN	Approach 2, $n = 15$	(a) Alternative transformation
12	FFNN	Approach 1, $n = 4$, bid Approach 1, $n = 4$, ask	(b) $[\tau^b_{Pij}, \tau^b_{Qij}]$ and $[\tau^a_{Pij}, \tau^a_{Qij}]$ modeled separately
13	LASSO	Approach 1, $n = 4$, bid Approach 1, $n = 4$, ask	(b) $[\tau^b_{Pij}, \tau^b_{Qij}]$ and $[\tau^a_{Pij}, \tau^a_{Qij}]$ modeled separately
14	FFNN	Approach 1, $n = 4$	(d) Using drop_out = 0.2
15	SVM	Approach 1, $n = 4$	(c) Setting C = 1.0
16	RF	Approach 1, $n = 4$	(d) Setting No. days = 120, Explanatory = Subset, Categorical = False
17	LASSO	Approach 2, $n = 25$	(a)(e) Approach No. 9 with speculators forecast as a combined group
18	FFNN	Approach 2, $n = 15$	(a)(e) Approach No. 11 with speculators forecast as a combined group
19	LASSO	Approach 1, $n = 4$	(f), removing G and C as explanatory variables

4.4.2 Generate Individual Participant Forecasts

The forecast horizon is from 11pm on the 31st December 2020, the *test_start* date, to 11pm on 31st December 2022, the *test_end* date. Within this forecast horizon we ignore trading periods that coincide with days when the clock change occurs or periods with missing public data (this equates to circa 1.8% of trading periods in the forecast horizon). Over this time horizon there are 411 market participants. For each modelling approach, for each participant, we perform a distinct walk forward validation. In our python walk forward validation implementation, if the participant has less than two weeks of historic data on which to train the model (e.g. participant is a new market entrant) then we use a naive model to generate the forecast (see *No. days* description in table 2). For the participant forecast error metrics we restrict our attention to the MAE and the rMAE. Given that participant's target data vectors have indices with different units and scales, the MAE and rMAE metrics are calculated for each index separately.

4.4.3 Combine Participant Forecasts

For a particular modelling approach, restricting our attention to Simple Order types, we then use the forecast participant output from the previous section to construct an aggregate view of the forecast bid and ask curves in each trading period over the horizon of interest. Using similar notation to section 4.3.4, P_j^{FDA} and Q_j^{FDA} denote the price and quantity values at which the forecast curves intersect in trading period$j$with $A_j^{b,\text{FDA}}$ and $A_j^{a,\text{FDA}}$ denoting the area under the bid and ask curves respectively. We then mimic equations (44), (45) and (48) (i.e. the TDA superscript is replaced by FDA) to calculate price, quantity and area under the curve error metrics. It should be noted that if we examine the forecast participant transformed piecewise order data generated in section 4.4.2 it is seen that a percentage (it will be dependent on the modelling approach but it ranges from 7% to 22%) of the forecasts are not in accordance with the market rules i.e. requirement that bid forecasts are monotonic decreasing and ask forecasts are monotonic increasing. Hence, prior to constructing the forecast bid and ask curves we programmatically adjust non-compliant participant forecasts[60]. A potential research extension would be to investigate whether the number of such occurrences could be reduced by including a related penalty in the cost function.

For comparison purposes, we also apply top-down models to forecast the DA market price directly, using the open source toolbox described in [38]. In this setup, the response variable is the I-SEM DA market price, P_j^{DA}, and the explanatory variables include lagged values of the target variable as well as the forward looking DA wind and demand forecasts (W_j and D_j; see section 4.4.1). We evaluate two LEAR models and five DNN models. The LEAR models vary in the number of calibration days used for training, while the DNN models vary in (a) the number of layers and (b) the number of calibration days; for the DNN configurations, hyperparameters were optimised following the toolbox methodology. All models were evaluated using walk-forward validation over the forecast horizon and results are reported for the best performing top-down model.

[60] e.g a hypothetical forecast transformed piecewise quantity vector of $[200, 199, 230, 250]$ would be adjusted to $[200, 200 + 0.01, 230, 250]$.

4.4.4 Cohort Forecast Analysis

To analyse aggregate forecast performance, we group participants into cohorts rather than examine individual forecasts, which is impractical given the number of participants. Classification proceeds as follows: participants with published physical characteristics (e.g. fuel type) are assigned to the corresponding cohort. If such information is unavailable (e.g. demand participants), classification is based on commercial behaviour, such as whether the participant is consistently buying or selling. Remaining cases are assessed for more complex configurations, such as demand and generation mixes (e.g. wind generation $<$ 10MW in capacity can be registered against a supplier unit). Any remaining unclassified participants are placed in the speculator category with a brief review of their commercial order data to identify any clear misclassifications. Speculators are almost exclusively identified by an *AU* prefix, the official market identifier denoting an *Assetless Unit*, and may belong to members with none, one, or several other registered participants (recall from section 4.2 that we consider *speculators* to be the participants that are financial traders and it does not include speculative actions that other participants or cohorts may engage in). While alternative classification methods could have been applied, the combination of publicly available participant data with manual inspection provides a sufficiently robust basis. Any residual misclassification is unlikely to materially affect the results. Full details are provided in appendix B.

Then, with each participant assigned to a cohort, for a particular modelling approach, for a given trading period we then perform the following distinct pieces of analysis

- For each cohort calculate and compare the areas under the actual[61] and forecast bid and ask curves.
- Take the stepwise order data for all participants. For a particular cohort replace its stepwise order data by its forecast order data and reconstruct the bid and ask curves, note the intersection price and compare it to the market price. Repeat the process for the remaining cohorts.

For the first bullet point, if we calculate the area metrics on an interval in the environs of the market price then it will provide information as to each cohorts contribution to the overall forecast error. The second bullet point can be viewed as a coarse estimate of the cohort impact given that it does not take into account potentially offsetting or reinforcing impacts of other cohorts. The above analysis

[61] For clarity bid and ask curves are not published at the cohort level; we derive them using participant stepwise order data.

is then repeated for each trading period over the forecast horizon from which we derive high-level insights for each cohort for a given modelling approach.

Finally, we note that in sections 4.4.3 and 4.4.4 it could be argued that a more reflective comparison would be between the transformed curves and the corresponding forecast curves (since it is the transformed participant piecewise order data that we forecast). However, our aim is to evaluate real-world performance, and for this reason we compare the actual and forecast curves.

4.5 Results

From sections 4.3 and 4.4 it is clear that results are available at both the micro and macro-levels. This occurs in both the time (i.e. individual trading period vs entire forecast horizon) and participant dimensions (i.e. aggregate vs cohort vs individual). Given the breadth of the analysis we are faced with the challenge of presenting a concise overview of the results. For example, taking a single modelling approach and configuration setting from table 2, executing steps 1 to 2 in figure 18 produces over 5 million forecast observations where each forecast is a multitarget output (i.e., a vector). Hence, our approach is to start each subsection by presenting a sample trading period, this helps reinforce some of the concepts introduced in sections 4.3 and 4.4. We then move onto highlighting a selection of noteworthy macro-level results.

4.5.1 Transformation approaches

This subsection evaluates the impact of different transformation approaches applied to stepwise participant order data. These transformations, introduced in section 4.3.3, are assessed using the transformation error metrics defined in section 4.3.4. The aim is to understand how the choice of transformation affects the aggregate accuracy (we do this prior to undertaking any of the forecasting steps outlined in figure 18). In figure 19 we plot the published and transformed bid and ask curves for a particular trading period. From the plots it can be seen that for this trading period the published curves (and intersection point) in blue largely overlap with the corresponding red curves which were derived using transformation Approach 1 with $n = 4$. In contrast the green curves, derived using Approach 1 with $n = 1$, appear to provide a poor approximation. In table 4 we summarise the results of various transformation approaches, the results cover the November 2018 to September 2023 timeframe. The columns correspond to different transformations, the rows correspond to transformation error metrics. In detail:

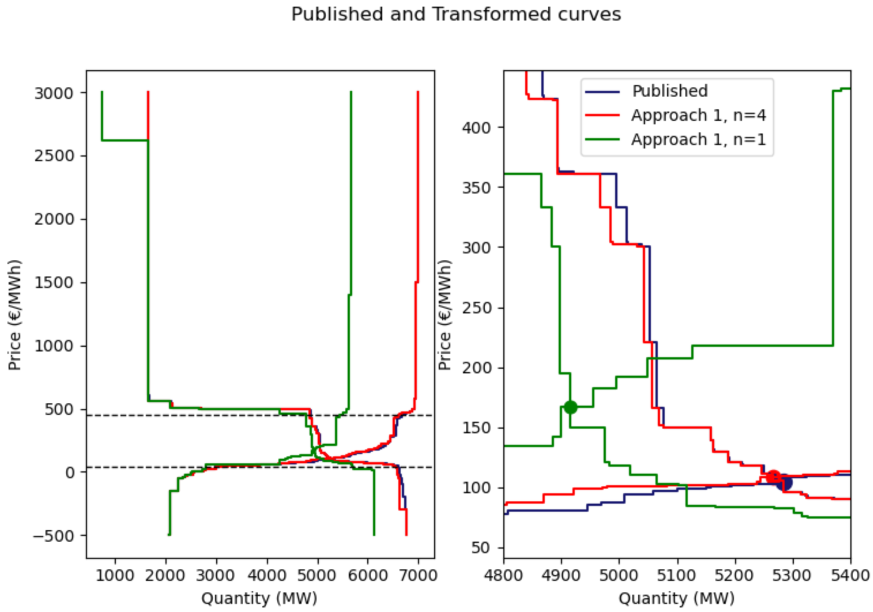

Figure 19: Published and transformed bid and ask curves for 4pm 9$^{\text{th}}$ January 2021. The left plot shows the full scale with horizontal lines at the 0$^{\text{th}}$ and 100$^{\text{th}}$ price quantiles. The right plot zooms in on the curve intersection regions.

- dim denotes the dimension of the $[\tau^b_{Pij}, \tau^b_{Qij}, \tau^a_{Pij}, \tau^a_{Qij}]$ vector under Approach 1 and the dimension of the χ_Q vector under Approach 2.
- Approach 1: we show different choices of n. Given bid and ask, price and quantity vectors, dim $= n * 4$.
- Approach 2: the market price quantile for trading period j, $\chi_{P,j}$, is calculated using DA market prices from hours $[j-1, j, j+1]$ over the previous 20 days (i.e. similar time periods from approximately the most recent three weeks). We then vary the number of quantile points, n, to assess the impact of different resolution levels; here dim $= n$.
- Note: the number of quantiles in Approach 2 is chosen to balance dimensionality with forecast feasibility; increasing the number of quantiles substantially would raise the dimensionality of the target data, with a potentially offsetting or negative impact on forecast performance. In addition, as Approaches 1 and 2 are alternative ways of generating the target data, we keep their dimensionalities similar to allow a fair comparison of forecast outputs in subsequent sections.
- Metric: price, quantity and curve MAE's are calculated for each month using eqns (44), (45) and (48) respectively. We normalise the MAE values by dividing by the average price, quantity and area value for each month which are derived using the published bid and ask curves. The resulting

monthly percentage values are then averaged to produce the values shown in the table. Note: as per eqn (46), we need to specify a $[lower, upper]$ price interval when calculating the area . We let the 0^{th} and 100^{th} quantile observations from $\chi_{P,j}$ define the interval[62].

Table 4: Normalised transformation error metrics for different transformation approaches, values rounded to the nearest .1%. The corresponding non-normalised transformation error metrics are available in appendix D table 12.

Metric	**Approach 1**						**Approach 2**	
	$n=7$ $(dim=28)$	$n=5$ $(dim=20)$	$n=4$ $(dim=16)$	$n=3$ $(dim=12)$	$n=2$ $(dim=8)$	$n=1$ $(dim=4)$	χ_P $(dim=25)$	χ_P $(dim=15)$
Price MAE	0.4%	1.3%	2.8%	4.8%	12.4%	39.2%	2.1%	4.4%
Quantity MAE	0.1%	0.1%	0.3%	0.8%	2.1%	4.6%	0.3%	0.7%
Bid Curve MAE	0.1%	0.1%	0.2%	0.7%	1.7%	4.3%	0.9%	1.4%
Ask Curve MAE	0.1%	0.3%	0.8%	1.6%	4.2%	9.8%	2.4%	3.8%

For Approach 1, from table 4 we see that increasing values of n leads to improved transformation error metrics. In addition the quantity and area related metrics are small relative to the corresponding price metrics and the bid curve MAE is noticeably smaller than the related ask curve MAE (an apparent exception is simply a rounding/display issue). For Approach 2, we again observe that increasing n (i.e. dim) results in improved transformation error metrics.

We interpret these results as follows. Dependent on the choice of n, the relatively small quantity and area related error metrics indicate that the transformations effectively capture the shape and positions of the bid and ask curves, albeit the price accuracy is not as high. In subsequent sections we will need to balance a desire to improve price accuracy (i.e. increasing n) without a potentially offsetting impact on the forecast performance (driven by an increase in dimensionality of the target vectors). For roughly similar levels of dimensionality, the outperformance of Approach 1 relative to Approach 2 becomes clear when we consider that Approach 1 captures granular participant level prices and quantities versus Approach 2 which yields distinct quantities but fewer, and shared, price buckets. The latter results in fewer steps in the aggregated curves; Approach 2 research extensions suggested in section 4.3.3 footnote 54 may help address this. Finally, taking transformations under Approach 1 with $n = 4$ as an example, appendix D table 13 provides an indication of the cohorts (e.g speculator, multi-fuel gas and coal) that are impacted by a loss of information.

[62] Alternative intervals have been examined but for succinctness we do not report them here.

4.5.2 Aggregate Forecasts

In this subsection we present aggregate forecast results for the modelling approaches defined in section 4.4.1. We report forecast error metrics analogous to those in section 4.3.4 (i.e. price, quantity and area); additionally, we benchmark our participant level approach against a publicly available top-down forecasting model. In figure 20 we plot the published bid and ask curve for a sample trading period and compare it to the corresponding forecasts from Naive (No. 0) and LASSO (No. 4) modelling approaches from table 2 for example. From the plot on the lhs it is clear that relative to the naive modelling approach the LASSO model reflects the general shape of the published bid and ask curve; however from the plot on the rhs we observe that the naive curve intersection price of €154.76/MWh is closer to the market price[63] of €151.78/MWh than the corresponding LASSO intersection price of €126.34/MWh.

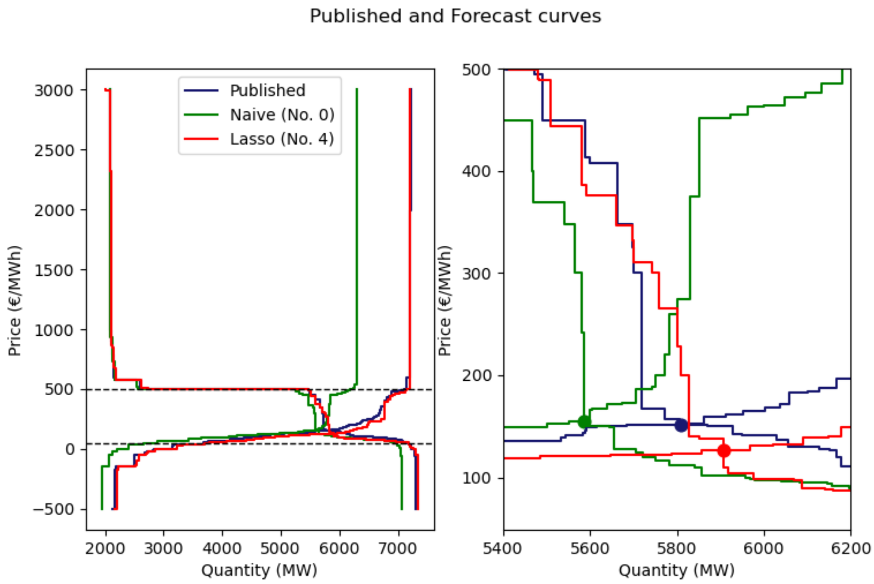

Figure 20: Published and forecast bid and ask curves for LASSO and Naive modelling approaches from table 2 for 5pm 8$^{\text{th}}$ March 2021. The plot on the left shows the full scale with horizontal lines at the 0$^{\text{th}}$ and 100$^{\text{th}}$ price quantile; the plot on the right zooms in on the curve intersection regions with points indicating the intersections. The remaining modelling approaches from table 2 are shown in figure 44 in appendix D.

Next, as is shown in fig 21, we calculate the average of the P_j^{DA} values for each day and compare it to the average of the P_j^{FDA} values for each day for a selection of modelling approaches. The Naive (No. 0) and LASSO (No. 4) modelling approaches are as described in table 2. The top-down forecast, described

[63]For reference, the transformed bid and ask curve, where the transformation is done under Approach 1 with $n = 4$, has an intersection price of €165/MWh. This is an example of a trading period where the difference between P^{DA} and P^{TDA} is non-trivial.

in section 4.4.3 and derived using the open-source toolbox provided by [38] is included for comparison purposes. From this plot it can be seen that both the top-down and LASSO (No. 4) modelling approaches appear to match the trend in the market prices, albeit there are some notable exceptions such as December 2022 for the LASSO (No. 4) model. In fig 22 we present a kernel density estimate, that is a non-parametric estimate of the empirical distribution, and time series plot of the forecast errors (i.e. $P_j^{DA} - P_j^{\mathrm{FDA}}$) for the top-down and LASSO (No. 4) modelling approaches from which it is seen that the forecast errors for LASSO (No. 4) have heavier tails. From the underlying data we observe that in circa 58% of trading periods the top-down modelling approach has smaller forecast errors than the LASSO (No. 4) modelling approach; however we also observe days in which the LASSO (No. 4) modelling approach produces a superior forecast relative to the Top-Down modelling approach (the opposite is also true). For instance, figure 23 presents actual and forecast prices over two consecutive days in March 2021. On the first day, the top-down forecast is more accurate, while on the second day, the bottom-up approach has a better forecast.

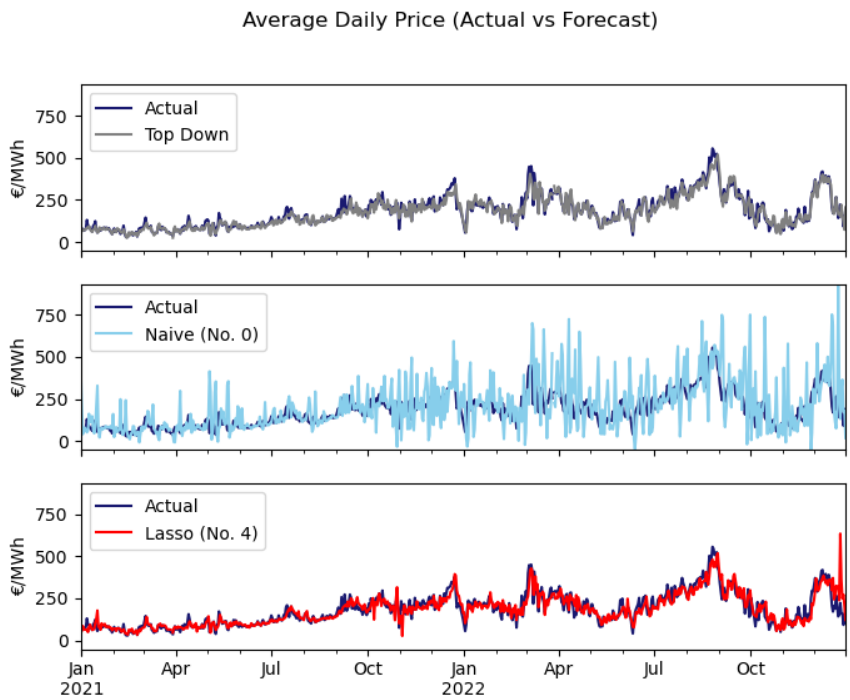

Figure 21: Plot of the average daily price, actual (i.e. published) versus forecast for three modelling approaches.

Next, table 5 presents normalised forecast error metrics for the modelling approaches from table 2 over the forecast horizon. To arrive at the table we replace the TDA with FDA superscript in equations (44), (45) and (48) and subsequently follow similar steps as to those used to derive table 4 in section 4.5.1. For the

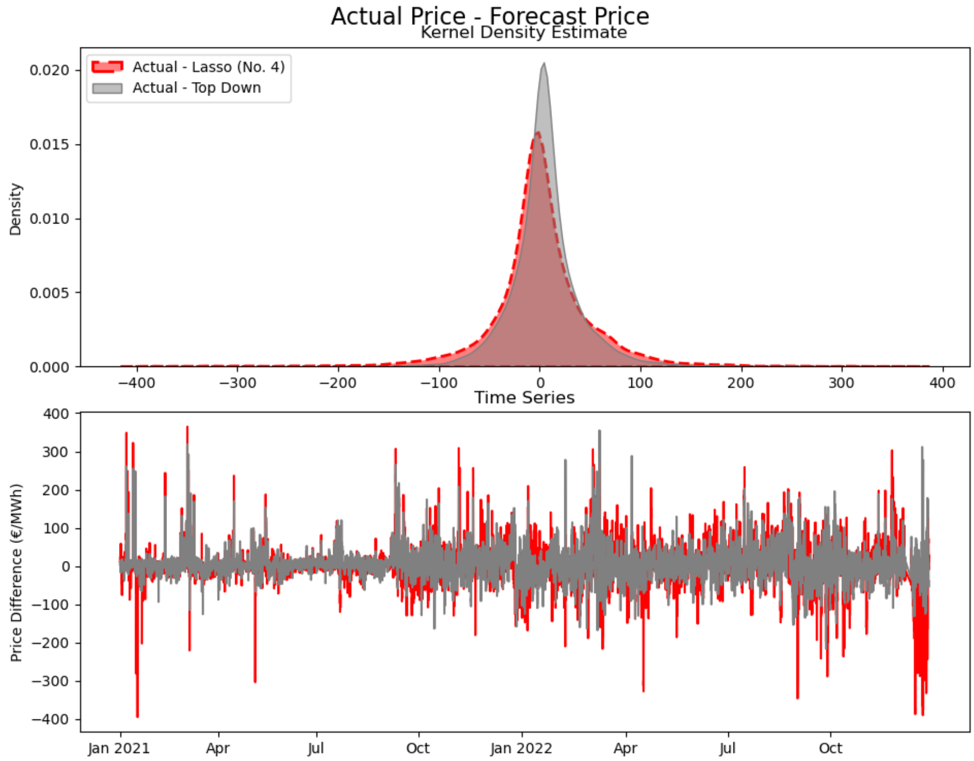

Figure 22: Kernel density and time series plots of $P_j^{DA} - P_j^{\mathrm{FDA}}$. Note: 20 forecast errors from the LASSO (No.4) model exceed 400 in magnitude and are excluded from the display (total observations $> 17,000$).

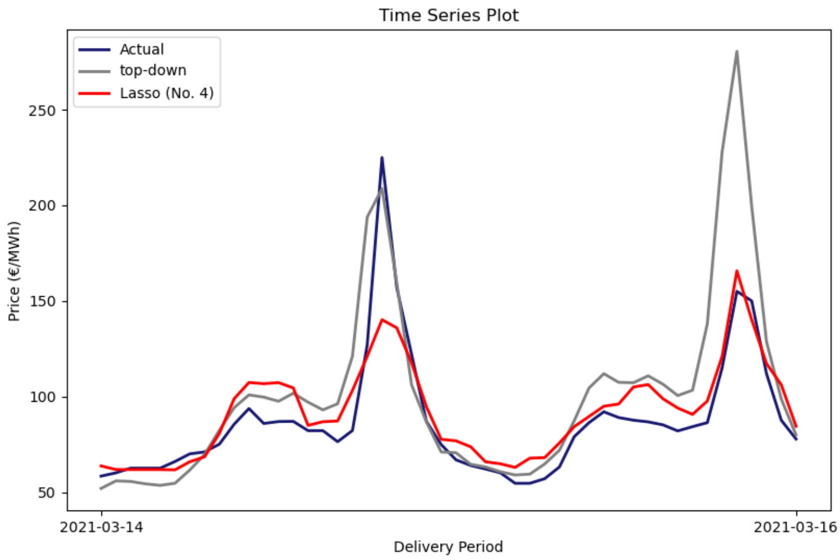

Figure 23: P_j^{DA} and P_j^{FDA} values over two successive days in March 2021. The forecasts are taken from the top-down benchmark model and LASSO (No. 4) described in table 2.

top-down modelling approach we are restricted to reporting price related metrics given the nature of the model output. From table 5 it is clear that (a) the top-down modelling approach has the lowest normalised Price MAE (b) LASSO models 3 and 4, SVM model 6 and FFNN model 7 appear to be the best bottom-up approaches (c) the normalised Bid Curve MAEs are significantly below the corresponding Ask Curve MAEs[64]. For completeness, table 6 presents the normalised forecast error metrics for the modelling approaches in table 3 over the forecast horizon (i.e. supplemental modelling approaches). Factors such as the choice of transformation approach[65], whether the target data generated via transformation Approach 1 is split and modelled separately, or whether speculators are modelled as a single entity, do not appear to significantly impact accuracy (i.e. the forecast error metrics remain in the same range). Additionally, figure 24, which visualises target and forecast data for the LASSO (No. 9) model from table 3 for example, suggests that at least visually the bottom-up modelling approach effectively captures trends in the transformed bid and ask curve data.

Table 5: Normalised forecast error metrics for different modelling approaches described in table 2 (the non-normalised transformation error metrics are available in appendix D table 14). Values rounded to the nearest 0.1%.

Metric	**Naive** (No. 0)	**RF** (No. 1)	**RF** (No. 2)	**LASSO** (No. 3)	**LASSO** (No. 4)	**SVM** (No. 5)	**SVM** (No. 6)	**FFNN** (No. 7)	**Top-Down** (DNN)
Price MAE	58.7%	29.1%	30.4%	19.2%	18.6%	27.1%	19.6%	18.3%	13.7%
Quantity MAE	8.4%	4.1%	4.2%	3.6%	3.4%	5.9%	3.7%	3.3%	
Bid Curve MAE	9.1%	4.7%	4.7%	4.9%	4.5%	7.5%	4.7%	4.3%	
Ask Curve MAE	25.3%	14.4%	14.8%	8.1%	7.9%	12.5%	8.8%	7.7%	

[64]We also note that if we derive the Bid Curve MAE and Ask Curve MAE values in table 5 using price intervals of $[P_j^{\mathrm{DA}} - 1, P_j^{\mathrm{DA}} + 1]$ in lieu of the 0^{th} and 100^{th} price quantiles (i.e. we are more interested in area differences in the immediate vicinity of the DA price) then the normalised area error metrics are by and large of the same

[65]i.e assuming similar dimensions, the choice of using transformation Approach 1 versus transformation Approach 2

Table 6: Normalised forecast error metrics for the different modelling approaches in table 3 (i.e. supplemental modelling approaches). Values rounded to the nearest .1%. The table was derived, and can be interpreted in a manner similar to that of table 5.

Metric	**LASSO** (No. 8)	**LASSO** (No. 9)	**LASSO** (No. 10)	**FFNN** (No. 11)	**FFNN** (No. 12)	**LASSO** (No. 13)
Price MAE	18.5%	18.7%	19.5%	20.5%	18.5%	18.6%
Quantity MAE	3.4%	3.2%	3.4%	3.2%	3.3%	3.4%
Bid Curve MAE	4.5%	4.8%	4.9%	4.9%	4.3%	4.5%
Ask Curve MAE	8.0%	8.2%	8.5%	8.2%	7.9%	7.9%

Metric	**FFNN** (No. 14)	**SVM** (No. 15)	**RF** (No. 16)	**LASSO** (No. 17)	**FFNN** (No. 18)	**LASSO** (No. 19)
Price MAE	22.3%	25.0%	23.7%	18.5%	20.5%	19.4%
Quantity MAE	3.8%	5.4%	3.6%	3.1%	3.0%	3.4%
Bid Curve MAE	4.9%	6.9%	4.4%	4.7%	4.7%	4.5%
Ask Curve MAE	10.6%	11.4%	11.3%	8.1%	8.1%	8.3%

We interpret the results as follows. The Naive (No. 0) forecast in figure 20 shows that while it is possible to have a reasonably poor overall fit to the bid and ask curve it may still lead to a good price forecast. This helps illustrate that with a bottom-up modelling approach a small price MAE is not a sufficient indicator of a good forecasting performance. While it may seem from figure 21 that the performance of a bottom-up modelling approach is in-line with that of a top-down modelling approach, the difference in the normalised price MAEs in table 5 stands out. We also highlight figure 25 in which we plot the monthly price MAEs for the top-down modelling approach against a selection of bottom-up modelling approaches; while there are months in which the relative performances are close, there are no months in which the top-down modelling approach is clearly outperformed.

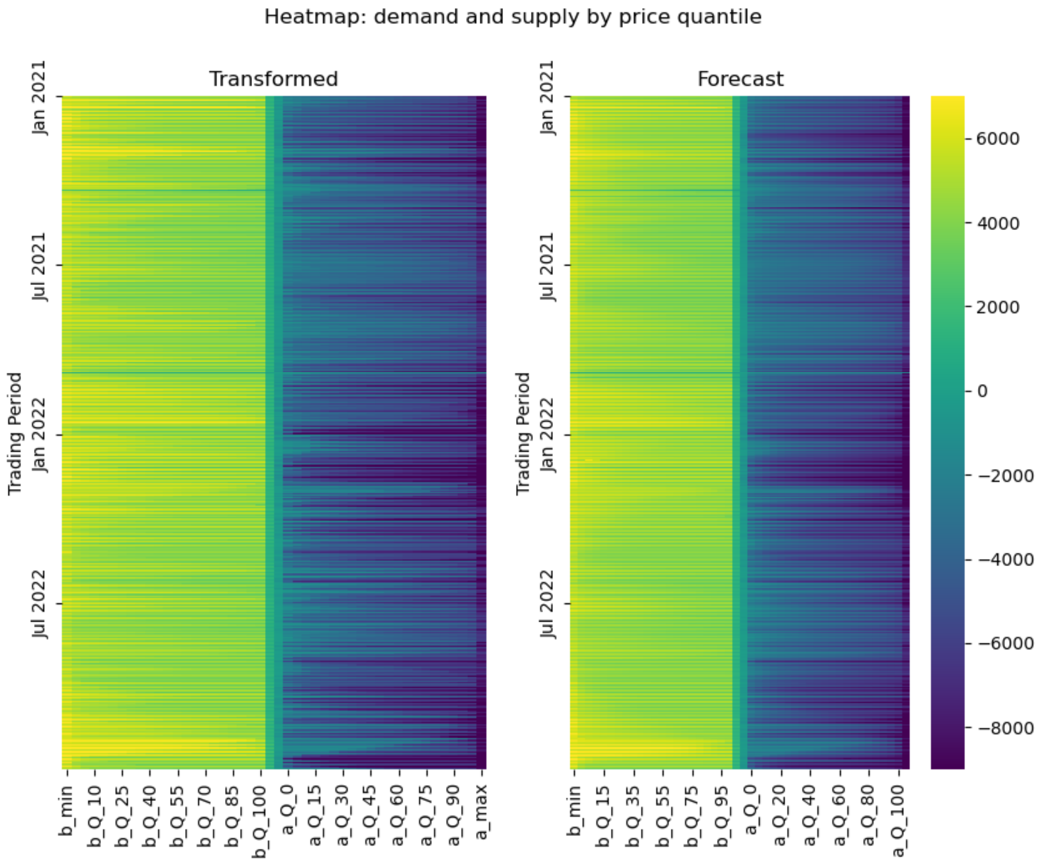

Figure 24: Heatmaps of bid and ask quantity data (Approach 2, dim = 25), aggregated across all participants. Left: transformed target data. Right: corresponding forecast data from the LASSO (No.9) model in table 3. Rows represent trading periods; the 25 bid and 25 ask columns correspond to aggregate quantities at each price quantile.

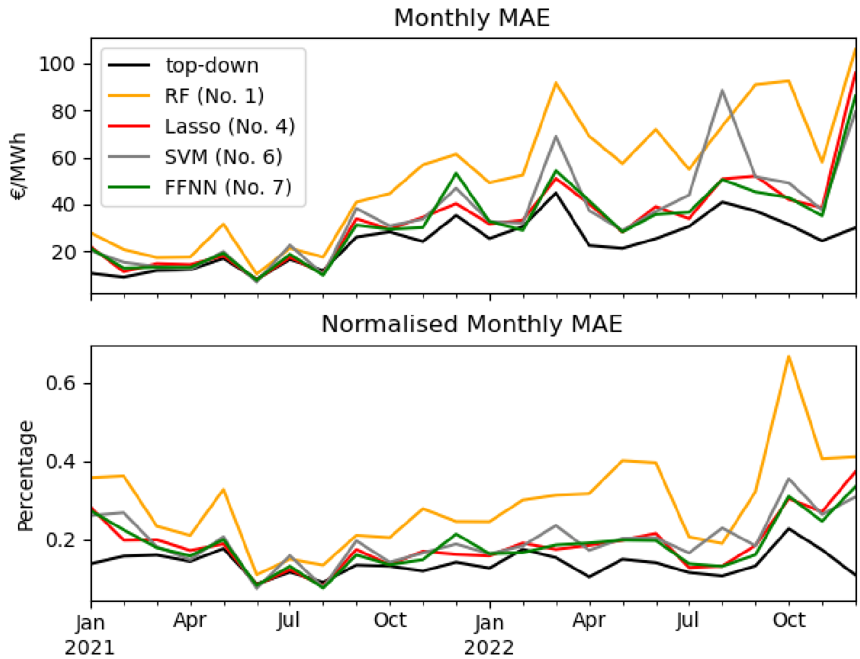

Figure 25: MAE and normalised MAE for each month in the forecast horizon for a selection of modelling approaches from table 2.

4.5.3 Cohort Forecasts

In this subsection we present results that highlight each cohort's contribution to overall forecast performance. We do so by building upon the area-based explainability framework to present metrics that compare forecast and actual values at the cohort level. Additionally, by applying techniques similar to those in section 4.4.4 (e.g. substituting a cohort's forecast values with its target values) we gain an intuitive understanding of how each cohort influences overall price and quantity forecast accuracy. We conclude with a discussion of cohort commercial behaviours, providing economic insights in the process.

As per the first bullet point in section 4.4.4 to understand the difference between the published and LASSO (No. 4) bid and ask curves in figure 20, for each cohort we construct the forecast and actual, bid and ask curves; then on a price interval of the DA price plus or minus €1/MWh, a fixed and deliberately simple interval chosen to allow consistent comparison across trading periods, we calculate and compare areas under the curve (this concept is illustrated visually in figure 26, right-hand panel). For this particular trading period a selection of cohorts with the largest area differences are shown in table 7. From both figure 20 and table 7 it is clear that the poor LASSO (No. 4) forecast is largely due to the ask curve[66] and this in turn is primarily due to the speculator cohort, albeit

[66]This is also evident from figure 20 if we consider the price resulting from (a) the ask curve

the pump storage cohort has a partially offsetting effect. In figure 26 we plot the corresponding actual and forecast ask curve for the speculator cohort; it is seen that the forecast does a reasonable job in predicting the overall shape but underestimates the prices that speculator participants are willing to sell in the market price region. On calculating the area difference between the actual and forecast speculator order data at the individual participant level, we identified 3 out of 26 active speculators in that trading period that were largely responsible for the deviation. The forecasts overestimated their sale quantities and underestimated their sale prices, with the latter being higher and more volatile than usual (see also appendix D figure 45 where we present the transformed and forecast order data for one of the three speculator participants around this time). Repeating the process for the Naive (No. 0) forecast in figure 20, we found that the differences in the forecast and actual, bid and ask curve areas were similar. The under-forecast in the Naive (No. 0) ask curve, mainly due to the wind cohort, was almost equivalent to the under-forecast in the Naive (No. 0) bid curve, primarily due to the buy cohort, effectively shifting the market price intersection point horizontally to the left (i.e. despite these incorrect forecasts, the offsetting effects meant the Naive intersection price was closer to the observed market price).

Table 7: The table shows the cohorts with the largest area differences between the forecast (LASSO No. 4 from table 2) and actual bid and ask curves for the 5pm trading period on 8^{th} March 2021. Areas are calculated on the price interval $P_j^{\text{DA}} \pm 1$.

Metric	buy	speculator	multi-fuel gas	wind	coal	pump storage
Bid Curve: Area Diff	-162	151	0	0	0	0
Ask Curve: Area Diff	0	1,441	65	45	-39	-490

for LASSO (No. 4) intersecting the published bid curve (b) the bid curve for LASSO (No. 4) intersecting the published ask curve.

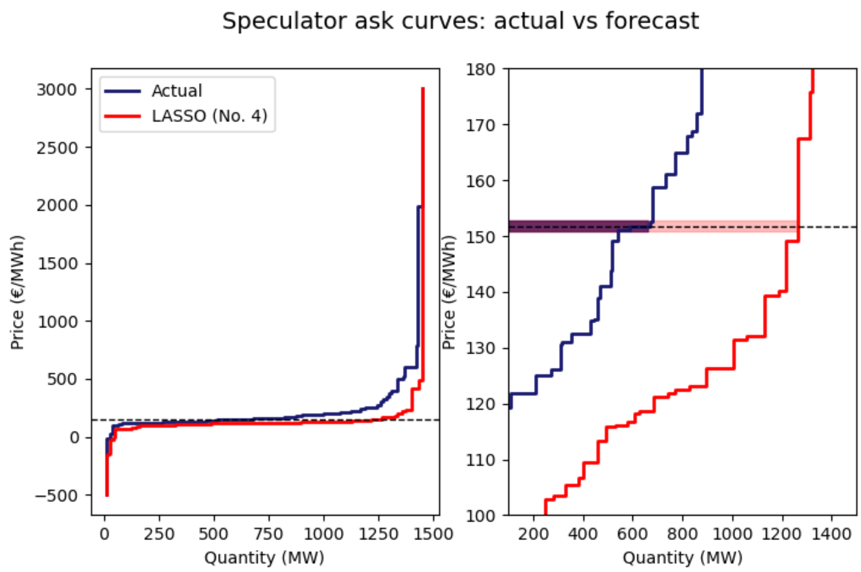

Figure 26: Speculator ask curves for 5pm on 8th March 2021. Left: full scale; right: zoomed view with shaded rectangular areas representing the area under the curve with respect to the price axis. The horizontal dashed line indicates the market price. The actual curve uses stepwise order data, the forecast curve is based on forecast output from LASSO No.4 table 2.

Table 8: The table shows cohorts with the largest average absolute area differences between forecast and actual curves calculated on price intervals $P_j^{\mathrm{DA}} \pm 1$ from January 2021 to December 2022 for the modelling approaches in table 2. Cohorts are groupings of participants and the process of assigning a participant to a cohort is described in section 4.4.4.

Cohort	**Curve**	**Naive** (No. 0)	**RF** (No. 1)	**RF** (No. 2)	**LASSO** (No. 3)	**LASSO** (No. 4)	**SVM** (No. 5)	**SVM** (No. 6)	**FFNN** (No. 7)
buy	bid	681	280	283	323	305	619	324	267
speculator	bid	433	308	314	296	304	344	305	309
pump storage	bid	48	52	51	56	55	61	51	55
speculator	ask	858	618	630	598	599	869	672	657
multi-fuel gas	ask	295	276	277	286	291	321	283	309
wind	ask	1,736	943	984	397	397	598	456	359
coal	ask	89	93	92	108	103	85	86	100
pump storage	ask	50	62	60	77	74	112	64	61
sell	ask	139	80	83	54	56	62	56	61
wind demand	ask	153	84	89	44	43	65	47	43

Generalising the approach in table 7 we arrive at table 8. With the columns representing the different modelling approaches and the rows representing a selection of some of the more noteworthy cohorts, the values represent the average of the absolute area differences between the forecast and actual curves (the areas being calculated on a price interval of the DA price plus or minus €1/MWh). If we consider the values on a relative basis (a) the speculator cohort exhibits large differences across both the bid and ask curves (b) both RF modelling approaches show large differences for the wind cohort (c) modelling approach SVM (No. 5) has large values for the speculator, wind and buy cohorts. For additional context, in table 17 appendix D we indicate what percentage of DA bought and DA sold quantities each cohort contribute.

For a coarse estimate of the impact that the cohort forecasts have on the price and quantity intersections we follow the second bullet point from section 4.4.4. That is taking stepwise order data for all participants, for the cohort of interest we replace its stepwise order data with its forecast order data, we then reconstruct the bid and ask curves and compare them to the published versions. Following this approach we arrive at table 9 which shows the normalised Price and Quantity MAE values that result when we replace the stepwise order data with the LASSO (No. 4) forecast data for example for the respective cohorts. The price metric for the speculator cohort, and to a lesser extent the buy/multi-fuel gas/wind cohorts, stand out. An alternative approach to assessing a cohorts impact is to replace their forecast order data with their actual stepwise order data (while continuing to use the forecast data for the remaining cohorts). If we follow this approach then the LASSO (No. 4) and FFNN (No. 7) normalised Price MAEs from table 5 for example of 18.6% and 18.3% both reduce to circa 8.5% when we use speculator stepwise order data in lieu of speculator forecast order data. Repeating the process, but for the wind, and then separately the buy cohort, the resulting normalised Price MAEs are in the 16.6% to 16.9% range (i.e. a much smaller reduction).

Table 9: Impact on the published bid and ask curves when, for a given cohort, its stepwise order data is replaced with the corresponding LASSO (No. 4) forecast data, while all other participant data remain unchanged. The table reports normalised price and quantity MAE values rounded to the nearest 0.1%. The remaining cohorts are shown in appendix D table 15.

Metric	**Buy**	**Speculator**	**Multi-Fuel Gas**	**Wind**	**Coal**	**Pump Storage**
Price MAE	3.9%	12.7%	3.0%	4.7%	1.1%	1.3%
Quantity MAE	2.1%	1.6%	0.8%	1.1%	0.3%	0.4%

We interpret the results as follows. The single trading period example in table 7 and figure 26 shows that comparing actual versus forecast areas under the curve at the cohort level, and subsequently applying it to individual participants, provides insights or explainability into the aggregate forecast performance. Presenting an average of the absolute area differences for the different cohorts and modelling approaches, as is done in table 8, is perhaps a simplification in that it does not consider potentially offsetting cohort impacts, but it signposts likely root causes in the differences in the overall forecast performance seen in table 5. Despite the smaller percentage of DA bought and DA sold quantities that the speculator cohort contributes, table 8 indicates that the speculator area differences are significant. This aligns with the results in table 9, and the alternative analysis in which we replaced cohort forecast data with cohort actual data (whilst using forecast data for the remaining cohorts), which indicate that the speculator cohort have a significant, almost outsize, negative impact on the overall forecast performance. While the bottom-up modelling approach appears to perform reasonably well, at least visually, in forecasting speculator cohort behaviours (for example the heatmap in figure 27 which compares transformed and forecast speculator data over the full forecast horizon), individual trading period data often reveals that the speculator cohort is the largest contributor to aggregate actual and forecast area differences. The speculator cohort is also a contributor to, or responsible for, a number of spikes in the bottom-up forecast error time series plot in figure 22.

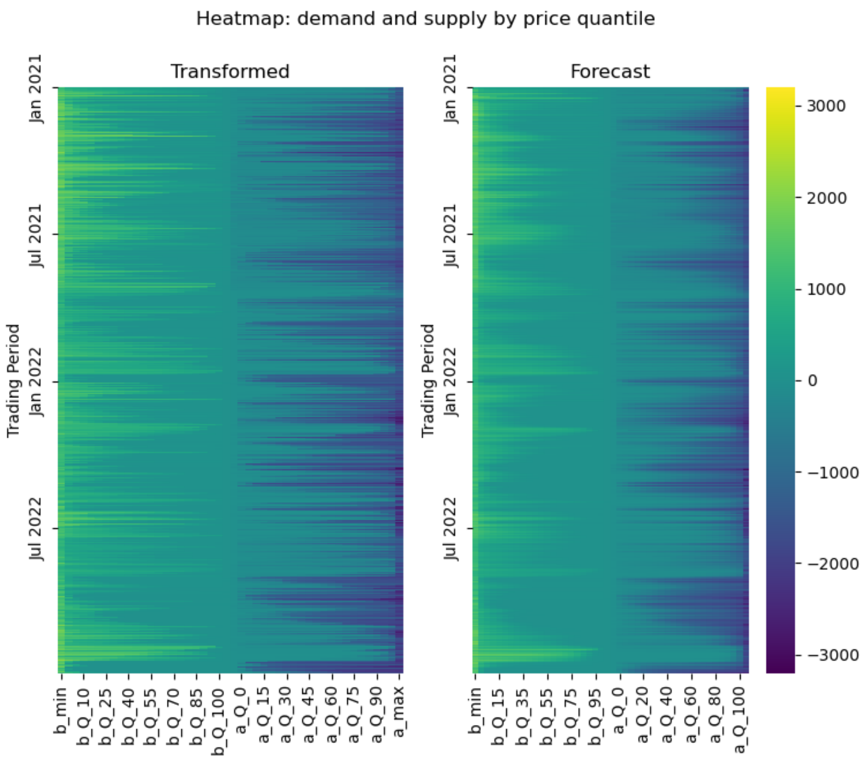

Figure 27: Similar to figure 24 but aggregated over speculator participants only. Columns represent quantity at each price quantile (25 bid, 25 ask), and rows are trading periods. Left: transformed (i.e. target) data; right: forecast data from the LASSO (No.9) model in table 3. The forecast heatmap appears less peaky than the transformed heatmap.

At this juncture, we present a small number of high-level discussion points relating to actual and forecast cohort behaviours in the Irish DA electricity market.

These points are intended to highlight selected aspects of a broader, complex, and evolving market environment, rather than to provide an exhaustive analysis.

- The inelastic nature of demand in wholesale electricity markets is well documented (e.g. [107]) and similar patterns are observed in both actual and forecast behaviours for the buy cohort. For this cohort approx 30% of submitted order quantities are placed at the market cap (i.e. the maximum allowable market price) thereby acting as price takers and roughly 90% are priced above the 100$^{\text{th}}$ market price quantile. Most buy cohort participants effectively submit single price-quantity pairs per order. We find that several models in table 2 forecast the cohort behaviour well, understandably so given the relatively stable and predictable quantity patterns that the cohort exhibits. Notably, SVM No.5 performs poorly for a few participants with quantity forecasts missing demand peaks and troughs; in contrast, SVM No.6, which applies a different regularisation parameter, offers improved fit for these cases.

- Wind generation accounted for approximately 25%[67] of DA sold quantities during the study period making the wind cohort the primary source of variable renewable generation supply. The impact of variable renewables on DA prices is well established (e.g. [108]), so here we focus briefly on cohort behaviours. Like the buy cohort, wind participants tend to submit an effective single price-quantity pair per order. Around 17% of volumes are priced at €0/MWh with the remaining predominantly at negative prices (24% of wind cohort quantities are offered at the market minimum price). As is the case in other electricity markets, negative offer prices for wind generators are often incentivised via renewable support schemes. Similar to the buy cohort, forecasting wind participant DA orders is largely a quantity prediction task. While most models in table 2 perform reasonably in this regard, random forests (RF No.1 and RF No.2) underperform primarily due to errors around peaks and troughs. Incorporating local weather forecasts and further model tuning (e.g. calibration length[68] or hyperparameter adjustments) could improve performance.

- The presence of oil and distillate generators in the DA market (i.e. when their offers are partially or fully matched) is notable as these thermal units

[67] The 25% figure likely understates total wind generation as units under 10MW are not required to participate directly in the DA market and therefore are not captured in the wind cohort. Their absence does not affect the order-level behaviour or price formation analysis in this thesis.

[68] RF No.16, tables 3 and 6, uses a larger calibration window and shows improved forecast performance.

are typically less efficient and more expensive than gas or coal-fired plants. Their offer prices generally exceed prevailing market levels placing them *'out of the money'* in the merit order. When they have matched (i.e. sold) volumes in the DA market it may signal market stress such as generation capacity shortages or spikes in input fuel prices for other generation cohorts. Over the study period the cohort equalled circa 1% of DA sold volumes and was marginal in a similar proportion of trading periods. While order formats vary across participants, commercial behaviour for this cohort is relatively stable. Most models in table 2 predict the behaviour well at both aggregate and participant levels.

- The multi-fuel gas cohort is a key player in the DA market accounting for 42% of sold volumes and setting the marginal price in circa 20% of trading periods. Commercial behaviour across this group is both diverse and dynamic. For example, participants often submit lower quantities for overnight periods and also submit offer prices that diverge from short-term fuel cost trends which is perhaps indicative of strategic bidding. Such behaviour is likely to be reflective of intra-cohort competition for merit order positioning, as well as broader considerations such as system wind and demand levels in addition to individual participants' wider portfolio considerations. While aggregate forecasts in table 2 generally capture the cohort's aggregate behaviour well, they sometimes incorrectly forecast participant level strategies.

- Speculators accounted for close to 5% of DA bought quantities and 12% of DA sold quantities during the study period, yet were marginal in circa 60% of trading intervals. Intuitively, this helps explain why forecast errors for this cohort, whether at the participant or aggregate level, can materially affect forecast market prices. This would also align with the literature referenced in the introduction which outlines the potential for speculators to impact/influence market prices. Examining the underlying speculator order data we observe dynamic and heterogeneous behaviours with frequent shifts in positioning both day-on-day and intraday.

4.5.4 Individual Participant Forecasts

This subsection provides a high-level overview of participant level forecast outputs (i.e. the outputs generated via step 2 in fig 18). In light of the cohort and participant level behavioural discussions presented at the end of section 4.5.3, and for the reasons outlined further below, we purposefully keep the discussion concise.

In figure 28 we take a single trading period and plot the target and forecast output for two modelling approaches for a small, illustrative set of participants spanning demand, thermal generation, speculator, and pumped storage categories. These participants were selected for clarity and continuity, as some were previously used in chapter 3, rather than being chosen at random[69]. The differing price and quantity scales in each of the four plots are indicative of distinct commercial strategies. The demand participant has target and forecast prices significantly in excess of the market price and hence in this instance qualitatively we are more interested in the quantities. For the thermal generator participant the target and forecast quantity data are largely the same, it is the prices that differ. In table 10, taking the LASSO (No. 4) modelling approach as an example, we present rMAE metrics for these four participants.

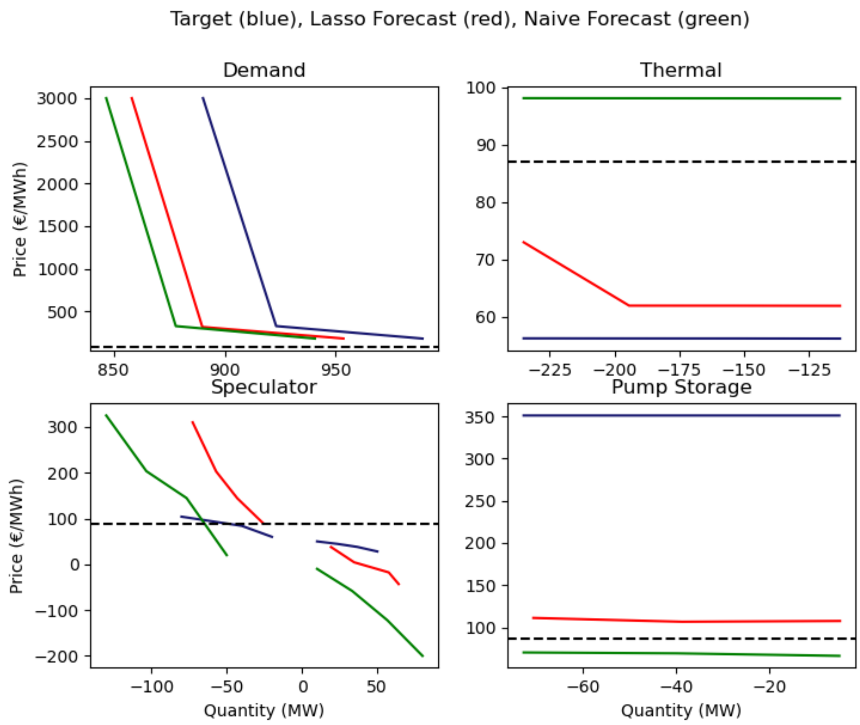

Figure 28: Transformed order data (target) and forecast data from Naive (No.0) and RF (No.4) modelling approaches for four participants at 7pm on 19th February 2021. Each plot shows price (Y-axis) versus quantity (X-axis), the horizontal dashed line indicates the market price.

When assessing the performance of a modelling approach, it may seem intuitive to examine participant level rMAE (or MAE) metrics across the entire population, analogous to the summary presented in table 10. However, considered in isolation, such metrics can be misleading as they abstract from important con-

[69] Identifiers for the demand, thermal generator, pump storage and speculator are SU_400071, GU_500130, GU_400363 and AU_400006 respectively.

Table 10: Relative MAE (rMAE) metrics for the LASSO (No.4) model from table 2, calculated over the two-year forecast horizon for the participants in figure 28. The target vectors τ_P^b, τ_Q^b, τ_P^a, and τ_Q^a are in $\mathbb{R}^4$, so the rMAE is computed for each component. For some indices, the rMAE values are undefined and are therefore left empty/blank.

Participant	rMAE τ_Q^b	rMAE τ_P^b	rMAE τ_Q^s	rMAE τ_P^s
Demand	$[0.6, 0.6, 0.6, 0.6]$	$[0.6, 1.2, 1.2, 1.4]$		
Thermal			$[0.5, 0.4, 0.4, 0.4]$	$[1.0, 1.0, 1.0, 1.0]$
Speculator	$[0.8, 0.8, 0.8, 0.8]$	$[0.8, 0.7, 0.8, 0.9]$	$[0.8, 0.8, 0.8, 0.9]$	$[0.8, 0.8, 0.8, 0.8]$
Pump Storage	$[2.0, , ,]$	$[1.4, , ,]$	$[1.8, 1.9, 1.9,]$	$[1.4, 1.5, 1.5,]$

textual factors such as the scale and number of observations and the behavioural drivers underlying forecast errors. In practice, we find that greater insight is obtained either by visualising participant level target and forecast series (as in figure 28) or by examining participant contributions to aggregate deviations via area based diagnostics as per section 4.5.3. Given the number of participants, trading periods, and modelling configurations, any concise summary necessarily masks substantial heterogeneity and isolated exceptions are unavoidable. For this reason, we restrict the discussion here to high-level observations (see appendix D.2 figures 46 and 47 for rMAE distributional examples for a selection of cohorts): most participants in the buy cohort have quantity rMAEs < 1; for buy participants with rMAEs > 1, the corresponding MAE is small relative to their overall demand levels and such cases typically involve far fewer observation points. Similar to the multi-fuel gas participant in fig 28, for this cohort we focus more on prices than on target and forecast quantities (as quantities are largely constrained by plant availability and operational considerations while bidding behaviour varies mainly through price). While most multi-fuel gas participants have price rMAEs < 1, this does not necessarily indicate good forecast performance. As previously explained, the modelling approaches tend to do well in capturing discernible trends (e.g. multi-fuel gas participants that two-shift[70]), but for some participants there are relatively frequent changes in the target data which don't seem commensurate with changes in fuel input costs, such instances appear challenging for the models to capture. For speculator participants, we are interested in both price and quantity forecasts. While some individual speculator forecasts capture commercial intent and are an improvement on naive models, the overall results are less clear cut and require individual consideration. For pump storage participants, we are interested in the price they submit and whether they are likely to buy, sell, or do nothing in a trading period. On examining the output,

[70]Running during the day when prices are often higher and off overnight when prices tend to be lower.

the modelling approaches in table 2 perform reasonably in predicting the commercial behaviours of individual pump storage participants; however, instances with rMAE > 1 indicate underperformance relative to the naïve benchmark.

4.6 Discussion

Section 4.5.1 showed that both transformation approaches reflected published bid and ask curves well, with performance depending on dimensionality. In addition to quantifying information loss, the approaches proved capable of handling diverse and complex participant order behaviours. For comparable dimensionalities, Approach 2 error metrics were slightly worse than Approach 1, although the extensions referenced in footnote 54 page 67 could likely mitigate this. For the Approach 1 transformation with $n = 4$, most information loss stemmed from speculator, multi-fuel gas, and coal cohorts.

Before considering forecast results, two limitations should be noted. First, as also true of some prior EPF studies, the construction of aggregate forecast bid and ask curves ignores the full complexity of the EUPHEMIA algorithm (e.g. treatment of Complex Orders). Ideally, forecast order data would be submitted directly to EUPHEMIA. In its absence, tools such as *PyPSA* [47] could approximate its functionality, allowing incorporation of Complex Orders alongside the Simple Orders modelled here[71]. While outside the present scope, this would represent a valuable extension. Second, some forecasting difficulties are structural: a new participant's first two days of activity are invisible due to data publication lags; participants occasionally submit manifest trading errors (e.g. infeasible generator offers); and unexpected outages periodically occur. These factors are noted for completeness but are regarded as second-order effects in the broader analysis.

In section 4.5.2, the benchmark top-down model outperformed the bottom-up approaches on price forecast error, with lower errors in about 58% of trading periods. Nonetheless, several bottom-up approaches performed reasonably well in matching overall price shape, and offered custom metrics that provided additional insight into forecast performance. In addition, forecast error magnitudes were broadly similar regardless of whether target data were generated with Approach 1 or Approach 2 transformations.

Section 4.5.3, Cohort Forecasts, highlighted the contributions of different participant groups, with speculators standing out. Despite representing a smaller

[71] Forecasting Complex Orders is feasible in principle as they involve a limited number of additional parameters that are often less volatile than Simple Order components; however, their inclusion may marginally reduce forecast accuracy due to the increased number of forecast components. Assessing this effect requires testing and is left for future work.

share of DA traded volumes, the speculator cohort exhibited large area differences and strongly influenced overall forecast accuracy. Replacing speculator forecasts with actual order data sharply reduced normalised forecast price and quantity MAE, an effect not seen to the same extent for other cohorts. The examples demonstrated both how forecast errors from a small number of speculator participants can significantly impact overall price forecast accuracy and more generally the dynamism in speculator commercial behaviours (e.g. figure 27). On examining the underlying data it was the case that speculators were a primary contributor to spikes in forecast error in a number of trading periods.

Section 4.5.4, Individual Participant Forecasts, provided high-level observations on population-level forecasts. Our aim was to apply a reasonable, robust process rather than optimise at the participant level. In practice, some participants (particularly speculators) were not well captured by the approaches in table 2, even when rMAE values remained below 1. Several factors may explain this:

- Unidentified explanatory variables: while the variables in section 4.4.1 align with literature standards, some participants may rely on additional factors not captured in our dataset.

- Modelling approach: alternative approaches such as optimising model configurations at the participant level, could improve results. Separately, a custom loss function emphasising errors near market price levels may also be beneficial, particularly for speculators, and could help reduce non-compliant forecasts (section 4.4.3).

- Market volatility: the forecast horizon coincided with extreme fluctuations in European gas and power prices, likely leading participants to frequently adjust their strategies, complicating model training.

- Human intervention: for some participants there is human activity or involvement in the auction process that may result in ad-hoc changes to the participant order data which is hard to anticipate.

Finally, we believe the methods developed here have broad applicability wherever granular participant level order data are available, including balancing markets with high order visibility, where anticipating order evolution may be strategically valuable. As noted in section 4.2, our approach could also support calibration or generation of input assumptions for cost of production (i.e. fundamental) models, or serve as an environment model in reinforcement learning settings. While some extensions may not necessarily improve price forecast accuracy, the contributions of this research include (i) explainability of results, (ii) practical methods for order data transformation, and (iii) explicit quantification of information loss, all of which add value to the literature.

4.7 Conclusion

In this chapter, in a day-ahead electricity market setting using statistical, machine-learning and neural network methods we implement an innovative data-driven bottom-up modelling approach and compare it to a top-down benchmark model. To implement the bottom-up modelling approach we first develop and apply participant order data handling algorithms, we subsequently perform a walk forward validation exercise across all market participants. Next, restricting our attention to Simple Order types, we aggregate the forecast participant output and compare it to publicly available bid and ask curve data. Using explainability related custom metrics we identify key factors influencing our bottom-up forecast performance finding that forecast errors associated with the speculator cohort significantly and negatively impact overall price forecast accuracy. The analysis highlights the complex nature of the price formation process (i.e. as per the single trading period example in section 4.5.3, incorrect forecasts across a small number of participants can have a non-trivial impact on price) in addition to the dynamism in individual participant, and cohort, commercial behaviours. We conclude by noting potential research extensions in addition to suggesting wider use cases for the data-driven bottom-up modelling.

5 Speculators and Price Inertia

5.1 Preamble

This chapter presents a slightly elongated version of the paper referenced in section 1.5 titled *Speculators and Price Inertia in a Day-Ahead Electricity Market: An Irish Case Study*. We note an overlap between the introduction in section 5.2 and the non-EPF literature review in section 2.5. Both are concise and to maintain narrative flow the structure and content of section 5.2 remain unchanged.

5.2 Introduction

Electricity markets have become a prominent area of academic research in recent decades driven by transformative developments such as market deregulation, the rapid expansion of renewable generation and the growing availability of short-term trading opportunities. These changes have spurred research across diverse areas including market power and competition [7, 8], price volatility [9], electricity price forecasting [10, 11, 57, 109], portfolio and asset optimisation [110, 111], and market dynamics [44]. In many jurisdictions the Day-Ahead (DA) market is one of the main routes to market through which electricity participants hedge or rebalance their positions ahead of physical delivery. DA market prices serve as reference prices for a range of financial derivative contracts and also influence retail electricity pricing. Given its central role, further insights into the functioning of DA markets remain important. This chapter contributes to that effort by examining two distinct but interrelated topics within a DA market setting: the activity of financial traders and price inertia, the latter defined as the extent to which market prices resist small shifts in demand or supply. A deeper understanding of these aspects is relevant for regulators evaluating market design and liquidity, as well as for participants active in short-term electricity markets.

As [60] observes, while the role of financial traders in commodity and energy markets has been widely studied, some of the literature exhibits ambiguity in defining financial traders and in specifying the metrics used. In a short-term electricity market context, we refer to financial traders as *speculators*, defined as participants that neither generate nor consume electricity but instead earn profits or incur losses based on the spread between their buy and sell prices either within the same market or across different market timeframes. While other market participants (e.g. thermal generators, batteries, renewable generation, demand etc) may engage in speculative behaviour, we reserve the term speculator for financial traders with no clear consumption or generation role. In several U.S. electricity markets, speculators participate through *Virtual Bidding*, which allows them to profit from the spread between DA and Real-Time (RT) prices[72]. In related studies [62] presents instances which would narrow the DA-RT spread with limited system benefits, whereas [63] find a reduction in both the spread and its volatility in California, [64] links lower transaction costs in MISO[73] to greater speculator participation and improved welfare, while [65] highlight reductions in forward premiums constrained by capital and regulatory limits[74]. While this U.S. focused literature offers valuable insights into price spreads and market efficiency, our study addresses a different and underexplored question: the presence and commercial behaviour of speculators in a European DA market. In contrast to the definitional ambiguity found in some of the broader commodity and energy market literature noted by [60], our analysis offers a clear and transparent framework for identifying speculators and quantifying their presence in a short-term European electricity market setting.

While prior research has examined short-term electricity market price volatility and the occurrence of price spikes, relatively little attention has been given to the concept of price inertia. As introduced earlier, we use this term to describe local price sensitivity to marginal shocks in demand or supply, which is conceptually distinct from volatility: volatility captures variability across periods, whereas inertia reflects price stability in response to incremental changes (see appendix E.1 figure 48 for a conceptual illustration distinguishing price inertia from price volatility). We do not address broader notions of inertia that concern long-run market stability, such as how structural reforms or system integration affect overall price dynamics at the market level. Studies addressing volatility include, among others, [69, 70, 9, 66]. Separately, systematic patterns such as calendar effects have been shown to impact electricity prices [114]. While these reflect broader shifts in demand and supply, our analysis concentrates on marginal, auction level changes. In terms of inertia in DA markets, studies such as [115] and [116] touch on the concept indirectly, as their methodologies examine the effect of larger supply shifts on prices. In a DA forecasting context, [51] illustrate how steep supply or demand slopes can generate high price sensitivity to shocks, a similar intuition underlying our measure of price inertia. One rare

[72]While Virtual Bidding does not exist in European electricity markets, participants may engage in similar speculative strategies through other mechanisms. For example, in the British market, *Net Imbalance Volume (NIV) chasing* allows traders to profit from spreads between ex-ante prices and balancing market prices [112, 113]. Similarly, [8] document speculative behaviour by wind generators in the Iberian market, who exploit price differences across sequential trading periods.

[73]MISO (Midcontinent Independent System Operator) operates the wholesale electricity market and transmission system across much of the Midwestern United States.

[74][65] note: "*When financial participants' ability to arbitrage the premium is limited, we find evidence suggesting that for large actors it may be more profitable to use their bids to increase the value of related instruments instead of arbitraging the premium.*"

study to consider market stability explicitly is [68], which finds that increased integration and competition in the Nordic market led to more stable and less volatile prices, where stability refers to longer run market behaviour rather than the local auction level price inertia examined in this chapter. While conceptually related, their approach differs methodologically and relies on daily averages, whereas we focus on hourly price data. In this chapter, we propose an intuitive and easily understood method for quantifying price inertia in DA markets. Our approach is analogous to liquidity and depth measures familiar from continuous trading, but designed specifically for an auction based setting.

Research Questions and Contributions This chapter addresses two distinct but related research questions:

- What is the extent of speculator participation in the market, and how has it evolved over time?
- Can we define a clear and intuitive approach to measuring price inertia, and how has it changed over time?

To address the first question, we use granular participant-level order and trade data to quantify the share of activity attributable to speculators and to measure their influence on price setting, specifically how often their bids or offers equal the market clearing price (i.e. marginality[75]). To address the second, we introduce a sensitivity based method for quantifying DA price inertia using publicly available demand and supply curves. Our findings show that speculators, while a small share of total activity, are frequently marginal in the DA auction (i.e. their bids or offers equal the market-clearing price). We also observe changes in price inertia after a structural market shift.

The chapter contributes in two ways. First, it provides empirical evidence on speculators in a European auction-based DA market, an area largely unexplored in contrast to the extensive U.S. literature. Second, it develops an intuitive and transparent methodology for quantifying DA price inertia, offering a practical lens for assessing how sensitive prices are to small shifts in demand or supply. This metric complements existing volatility measures and has direct relevance for ongoing market design and reform debates[76].

While our methodology is deliberately straightforward and transparent, it is informed by the complexity of the market environment, making an observational, data-driven approach both appropriate and necessary. Key simplifying assumptions are outlined in subsequent sections.

[75] *Marginal pricing* refers to the mechanism by which short-term electricity market prices are determined [6, 13, 117, 31].

[76] e.g. Great Britain's *Review of Electricity Market Arrangements* [118, 119]

Chapter Structure The structure of the remainder of the chapter is as follows: in section 5.3 we present a brief overview of the (short-term) market structure and pricing algorithms encountered in the Irish electricity market, we also provide a hypothetical example of a speculator in this setting. In section 5.4 we detail the data and methodology. Section 5.5 presents the empirical analyses whilst in section 5.6 we discuss the main findings. In section 5.7 we conclude.

5.3 Market Structure

Market The Irish electricity market is the domain in which the research takes place. Noteworthy points from the market overview in section 3.1 include

- As of 2025, the DA and IDA1 auctions account for about 97% of short-term trading volumes across DA, IDA1–3, and IC markets. References to market shares in section 5.5 relate to the study period (Nov 2018–Dec 2022), though percentages remain broadly consistent.
- To achieve a physical position in the Irish electricity market a participant needs to trade in one of the above markets, as a result the market structure therefore implies a non-trivial level of (DA) liquidity[77].
- Again, for clarity purposes we note that speculators that are active in the Irish electricity market do not consume or generate electricity.
- A structural market change occurred on 1st January 2021, prior to that date the Irish and British electricity markets were coupled in the DA and IDA1 auctions, post that date market coupling occurs at the IDA1 and IDA2 auctions.

Pricing Algorithms This section briefly introduces the DA pricing algorithm. Unlike a simple intersection of aggregated demand and supply curves, the actual market clearing process is considerably more complex and a full technical specification (including implementation details) is not publicly available. This combination of algorithmic complexity and limited transparency motivates our use of a deliberately straightforward observational approach when analysing speculator behaviour (section 5.4.2). We also note that our price inertia analysis (section 5.4.3) relies on a simplifying assumption that abstracts from the full market clearing process.

[77] We note that a subset of participants may also achieve a physical position by participating in the balancing market, but this does not impact the study.

EUPHEMIA, the *Pan-European Hybrid Electricity Market Integration Algorithm*, is used to calculate hourly DA prices in a number of European markets. According to the EUPHEMIA Public Description [5]:

- *"The algorithm can handle a large variety of order types at the same time"*, including *Aggregated Hourly Orders*, *Complex Orders* (e.g. with Minimum Income Condition or Load Gradient constraints), *Scalable Complex Orders*, and *Block Orders*.

- The algorithm solves a *Welfare Maximization Problem (Master Problem)* together with three interdependent sub-problems, one of which is the *Price Determination Sub-Problem*. In the Master Problem, *"EUPHEMIA searches among the set of solutions for a good selection of block and MIC orders that maximises social welfare. Once an integer solution has been found for this problem, EUPHEMIA moves on to determine the market clearing prices"*[78].

From this description it is clear that treating the DA auction simply as the crossing of bid and ask curves, while intuitive, omits the complexities of the actual clearing algorithm. EUPHEMIA is also subject to ongoing incremental changes [120, 121], making it a dynamic rather than static construct. In contrast, the algorithm governing the Irish Balancing Market is fully documented and publicly available [122][79].

Speculators In the following hypothetical example, the speculator buys in the DA and IDA1 markets and sells in the Balancing market. More generally, the opposite scenario also occurs, where a speculator sells in the ex-ante markets and buys in the Balancing market (speculation may also occur between the DA and intra-day markets). Consider a speculator active in a single delivery hour (10:00–11:00). Recall that the DA market settles in hourly blocks, while the IDA1 and Balancing markets settle in half-hourly blocks. Suppose the speculator forecasts prices of €X/MWh in the DA for 10:00-11:00 and in each IDA1 half-hour (i.e. 10:00-10:30 and 10:30-11:00), and expects the corresponding Balancing market price to be €$(X + 20)$/MWh in each half-hour. The speculator decides to buy 50MW in the DA, and an additional 25MW in each half-hourly block in the IDA1 market, giving a total ex-ante position of 75MW as shown in figure 29. Since the speculator neither generates nor consumes electricity, its position in the Balancing market is the opposite of these purchases. If its forecasts are correct, its income

[78]Here, social welfare denotes total economic surplus (consumer plus producer surplus). That is, the difference between willingness to pay and willingness to accept across all accepted orders, subject to market constraints.

[79]For the British electricity market, *Elexon* provides similar imbalance pricing guidance in [123].

from the Balancing market of €75$(X + 20)$ will be greater than its total costs of €75X from its DA and IDA1 positions, thereby earning a profit of €1500.

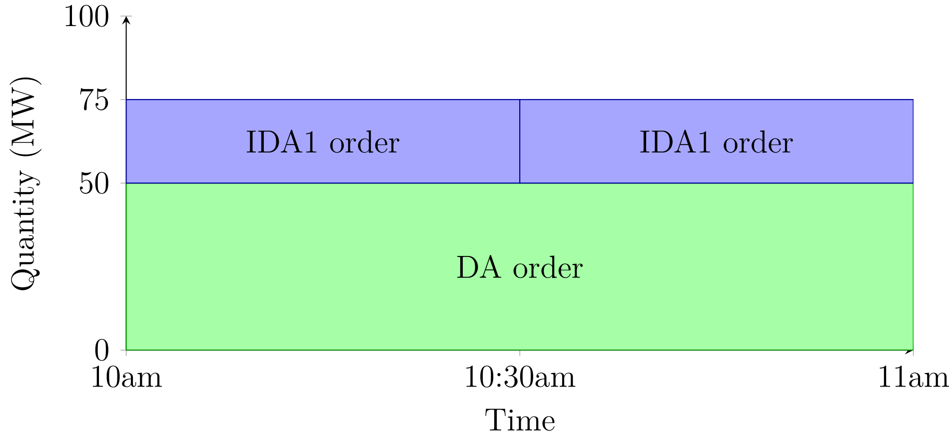

Figure 29: Speculator hypothetical example

5.4 Datasets and Methodology

In this section we describe the underlying datasets and the empirical analyses. For context, we note that the data used in the study coincides with a diverse and volatile set of energy price regimes including the Covid-19 pandemic[80] and the European energy crisis (described in section 1.2). Any reference to accompanying material in the appendices is for replicability and transparency purposes with the key details being captured in this section 5.4.

5.4.1 Datasets

SEMOpx and SEMO are two of the bodies involved in the operation of the Irish electricity market. Among the publicly available datasets they publish are:

- Granular Participant Data: As detailed in section 3.3 page 37, SEMOpx publishes a distinct ETS Bid File for each of the four ex-ante markets. These files provide participant and trading-period-level order and executed quantity data, where order quantities reflect submitted bids and offers (i.e. willingness to buy or sell), while executed quantities represent the portion of those orders that are matched in the market-clearing process. Positive values denote purchases, while negative values indicate sales (this convention applies to both order and executed quantities).

[80] Which resulted in a low price commodity environment at the outset of the pandemic [124, 125].

- Bid Ask Curve Data: As per section 3.3 page 37 for each ex-ante auction SEMOpx publish a *BidAskCurve* file containing a monotonic increasing (decreasing) and anonymised view of the sell (buy) orders for each trading period in the auction.

- Other Data:

  - *PUB_MnlyRegisteredCapacity* files which provide participant registration data such as registered plant capacity and *FuelType* (if applicable).
  - *PUB_30MinImbalCost* file containing the Balancing Market price, P_i^B, for trading period i. Similarly, *MarketResult* files containing DA, IDA1, IDA2 and IDA3 market prices (i.e. P_i^{DA}, P_i^{IDA1}, P_i^{IDA2}, and P_i^{IDA3}).

In the Irish electricity market, participants may register a variety of unit types, including *Supply*, *Generator*, *Assetless*, and *Trading*. Because these units can behave flexibly, we follow the steps outlined in appendix B to identify speculators. As explained there, the procedure was developed iteratively, including interrogation of participant samples to ensure robust cohort allocation. As a further check, we manually reviewed all participants classified as speculators to confirm the appropriateness of their assignment. The procedure does not capture every possible speculative action, and consistent with our conceptual definition in the introduction, does not include participants with obvious generation or consumption roles who may also act speculatively. Nevertheless, it provides a transparent and reproducible framework. Alternative classification approaches are possible, and these caveats are noted in both the Limitations section and appendix B.

5.4.2 Market and Speculator Analysis

Quantities Given the importance of the DA market in an Irish electricity market context (section 3.1), we examine how DA volumes have evolved for speculators and for the market as a whole, with the latter providing context for the former. Using information in the DA market ETS Bid Files we derive a per trading period time series that shows

- The quantities of energy that participants were willing to buy or sell, we refer to these as the *order quantities*.

- The quantities of energy bought or sold which we term the *matched quantities*.

Marginal Participants To understand which participants are marginal (i.e. participants whose submitted order prices are equal to the market-clearing price) and whether or not this pattern has changed over time for each trading period

1. From the MarketResult file (section 5.4.1), determine the DA market price for the trading period.

2. Using the ETS Bid File, select the rows where market participants have an active order in that trading period.

3. Iterating through each row in the previous step

   - if the market price equals any of the price points in the participant's order, we flag the participant
   - else do nothing

   If one or more participants have been flagged, then we have identified the marginal unit(s) and the process ends. If no units have been flagged, continue to the next step.

4. Utilising the BidAskCurve file, for that trading period, ascertain how the bid and ask curves intersect. If the curves intersect vertically, determine the two points at which the curves overlap. Denote the prices associated with the overlap as *lower_price* and *upper_price*.

5. Iterate through each of the rows selected in step 2. If either the lower_price or the upper_price identified in step 4 equals any of the price points in the participant's order, we flag the participant as being marginal.

A by-product of the marginal analysis is the identification of bid and ask curve intersections. As a result, in addition to identifying marginal participants we also examine how DA bid and ask curves intersect and how these intersection patterns have evolved over time. Given the stepwise nature of the bid and ask curves in the Irish electricity market there are three possible types of intersection: horizontal (the curves are horizontally parallel and intersect over ranges of quantities), vertical (the curves are vertically parallel and intersect over ranges of prices) and perpendicular (the curves intersect perpendicularly at a single price and quantity point). For a perpendicular intersection, if at the point of intersection the ask curve is horizontal, *ACH*, we denote it as a *perpendicular (ACH)* intersection. Similarly, if at the point of intersection the bid curve is horizontal, *BCH*, we denote it as a *perpendicular (BCH)* intersection. In appendix E.2 we provide examples of different bid ask curve intersection types.

Aggregate speculator behaviour and profitability To examine how aggregate speculator behaviour has evolved, we begin with their DA order (i.e. pre-clearing) quantities. Rather than comparing order quantity distributions across time (e.g. using kernel density estimates), we take a simpler approach that still captures the main trends. For trading period i, let the DA market price be P_i^{DA} and the standard deviation of DA prices for the corresponding hour over the previous 20 days be SD_i. We partition the price domain into three intervals:

- Interval 1: $[-500,\ P_i^{\mathrm{DA}} - \mathrm{SD}_i]$
- Interval 2: $[P_i^{\mathrm{DA}} - \mathrm{SD}_i,\ P_i^{\mathrm{DA}} + \mathrm{SD}_i]$
- Interval 3: $[P_i^{\mathrm{DA}} + \mathrm{SD}_i,\ 3000]$

Here -500 and 3000 are the minimum and maximum permissible DA prices[81]. For each interval we calculate the sum of speculator sell orders, and separately buy orders, falling within that range. These values are then aggregated by month, and we plot the monthly totals from November 2018 to December 2022.

Next, we consider aggregate matched (i.e. cleared) speculator quantities. Our main focus is DA, but motivated by the structural market change and the hypothetical example (both are presented in section 5.3), we also examine IDA1 and Balancing Market values. Let Speculator denote the set of speculator participants and $Q_{i,j}^m$ the matched quantity for trading period i by speculator j in market m. We define

$$\text{Speculator DA Matched Quantity}_i = \sum_{j \in \text{Speculator}} Q_{i,j}^{\mathrm{DA}} \tag{49}$$

$$\text{Speculator IDA1 Matched Quantity}_i = \sum_{j \in \text{Speculator}} Q_{i,j}^{\mathrm{IDA1}} \tag{50}$$

$$\text{Speculator Balancing Market Quantity}_i = \sum_{j \in \text{Speculator}} Q_{i,j}^{B} \tag{51}$$

We compute these for each trading period and use pair plots to compare behaviours before and after the structural change (IDA2/IDA3 are omitted given their negligible volumes). Note that the index i in equations (49–51) refers to different temporal granularities depending on the market (hourly for DA, half-hourly for IDA1 and BM), but since the values are expressed in power (MW), we are still comparing like with like in the pair plots.

Ignoring transaction costs which are reasonably small in an I-SEM context and therefore do not materially affect the first-order estimates considered here,

[81] The maximum DA price rose in stages from €3000/MWh to €5000/MWh; in our analysis all orders are capped at €3000/MWh.

speculator j's profit or loss is estimated by

$$-\frac{1}{2}\sum_{m\in\{\text{DA,IDA1,IDA2,IDA3,B}\}} P_i^m\, Q_{i,j}^m \tag{52}$$

where P_i^m is the price in market m for period i. Since we sum across markets with different temporal granularities, we take the lowest resolution, so i in equation (52) corresponds to a half-hourly settlement period. The negative sign reflects the buy/sell convention. Given that $Q_{i,j}^m$ is expressed in power (MW), when multiplying by prices (€/MWh) and accounting for the half-hour duration, we multiply by $\frac{1}{2}$ to convert from MW to MWh. We repeat this calculation for all speculators and periods to obtain aggregate profitability.

5.4.3 Price Inertia Analysis

We measure price inertia in the Irish DA market by applying a small horizontal shift to the bid or ask curve and calculating the resulting intersection price. The difference from the actual market price provides our inertia measure. In principle this would require a full rerun of the EUPHEMIA algorithm. However, given that our shifts are small (1 MW or 10 MW, minor relative to average forecast demand of 4300 MW per period over the study horizon) and the algorithm is not publicly available, we adopt the simplifying assumption that a rerun is unnecessary. This approach is consistent with existing literature where similar approximations are used in fundamental supply stack models [49, 50] and curve based simulations [51, 41]. For trading period i, the process of shifting the DA ask curve by X MW is:

1. Retrieve the bid and ask curves from the DA *BidAskCurve* file.
2. Add X MW to each $(\text{quantity}, \text{price})$ point on the ask curve.
3. Identify the intersection with the bid curve; call the resulting price the simulated DA price P_i^{SDA}.
4. With P_i^{DA} and SD_i as defined in section 5.4.2, compute:

$$\text{Price Difference}_i = P_i^{\text{SDA}} - P_i^{\text{DA}}, \tag{53}$$

$$\text{Custom Metric}_i = \frac{|\text{Price Difference}_i|}{SD_i}. \tag{54}$$

Interpreting the raw *Price Difference* series alone can be misleading given the very different price regimes over the study horizon (fig 8 section 3.3 presents a plot of the corresponding electricity market prices). To normalise, one option is to divide by P_i^{DA}, but this fails when prices equal €0/MWh. Instead, we divide by the standard deviation measure SD_i (section 5.4.2), though alternative normalisations are possible. While we report results for positive shifts to the ask curve, we also tested negative shifts and bid curve shifts (±1 MW, ±10 MW); results were qualitatively similar and are omitted for brevity.

Finally, we extend this baseline approach with two related methods: one that examines the distribution of price inertia and another that uses speculators' order data to estimate inertia (akin to applying shifts directly to the bid or ask curve). We note that in these extensions the small shift assumption may not always hold.

Price Inertia Distribution

- For each trading period adjust the ask curve in +1MW increments until the simulated market price, P_i^{SDA}, differs from the published market price. Let $\mathrm{Shift}_i^{+1\mathrm{MW}}$ denote the shift required to observe the price change where the +1MW superscript indicates that we have used +1MW increments. Use equation (53) to calculate the associated price difference which we denote as $\text{Price Difference}_i^{+1\mathrm{MW}}$.
- Repeat with −1MW decrements to derive the corresponding $\mathrm{Shift}_i^{-1\mathrm{MW}}$ and $\text{Price Difference}_i^{-1\mathrm{MW}}$ values.
- For each trading period, define and calculate the *Min Shift* and *Price Impact* via

$$\text{Min Shift}_i = \min\left(|\mathrm{Shift}_i^{+1\mathrm{MW}}|,\ |\mathrm{Shift}_i^{-1\mathrm{MW}}|\right)$$
$$\text{Price Impact}_i = \begin{cases} |\text{Price Difference}_i^{+1\mathrm{MW}}|, & \text{if Min Shift}_i = |\mathrm{Shift}_i^{+1\mathrm{MW}}| \\ |\text{Price Difference}_i^{-1\mathrm{MW}}|, & \text{if Min Shift}_i = |\mathrm{Shift}_i^{-1\mathrm{MW}}| \end{cases} \tag{55}$$

A graphical illustration of this procedure is provided in appendix E.7 fig 55. While alternative analytical derivations of Min Shift and Price Impact may be possible, we adopt this simple numerical approach for transparency and reproducibility. Figure 40 in section 5.5.2 presents presents density histograms of the resulting Min Shift and Price Impact values.

Order Example This approach involves removing the order data for a single (speculator) participant, reconstructing the bid and ask curves using the remaining participants' order data, and determining the new intersection price. Implementation details for this alternative approach are provided in appendix E.3 while in section 5.5.2 we present a trading period example where we omit the order data for a single speculator participant.

5.5 Results

5.5.1 Market and Speculator Analysis

Quantities Figure 30 shows stacked area charts of DA sell order and sell matched quantities (section 5.4.2) in TWh, grouped by FuelType[82] and month. The dashed vertical line marks the structural market change (section 5.3); sells are represented by negative quantities. On average, 52% of sell order quantities are matched. For conciseness we omit the buy side plots, but note that about 83% of buy orders are matched. The absence of Coal, Oil, and Distillate categories from matched quantities in certain periods reflects the DA merit order in the Irish market. The speculator sell order and matched quantities (red series) are seen to

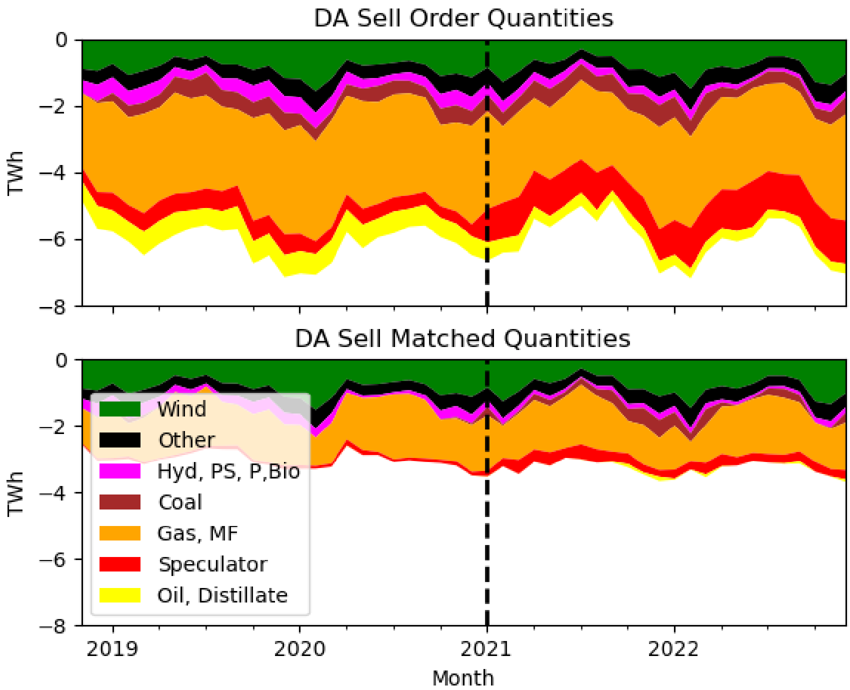

Figure 30: The top plot presents Irish electricity market DA Sell Order Quantities, the bottom plot presents DA Sell Matched Quantities. Values are in TWh and the data has been grouped by FuelType and month. The dashed vertical line indicates the point in time at which the structural market change related to Brexit occurred.

increase after the structural market change. Pre-change, speculators accounted for ∼10% of DA order quantities and ∼2% of matched quantities; post-change these shares rose to ∼20% and ∼6%. The number of active speculators also increased: before the change an average of 16 (14) submitted sell (buy) orders per day, rising to 30 (24) afterward, that is almost a doubling. For context, the DA

[82] See appendix E.8 for details on FuelType notation.

market hosted 472 distinct participants over the study period.

Table 11 summarises these statistics. We note in passing that speculators' relative share is larger in IDA1, averaging 54% (42%) of buy (sell) order quantities and 46% (33%) of buy (sell) matched quantities, though IDA1 is much smaller in size than the DA auction.

Table 11: Speculator participation in the Irish DA market. *Note:* over the study time period, DA accounts for $\geq$85% of ex-ante traded volumes, IDA1 up to ~11%, and the remainder is split across IDA2, IDA3, and the continuous market.

	Pre-change	Post-change
DA order quantities (share)	~10%	~20%
DA matched quantities (share)	~2%	~6%
Avg. active speculators/day (sell)	16	30
Avg. active speculators/day (buy)	14	24

Marginal Participants The top plot in figure 31 shows the percentage of marginal observations, grouped by FuelType and month. Gas and multi-fuel units exhibit a declining tendency to be marginal, falling from about 35% of trading periods before the structural change to around 20% after. In contrast, speculators appear marginal more often, rising from 49% to 66%. That speculators are marginal in a large proportion of trading periods contrasts with their overall minority share of market activity noted earlier. The bottom plot in figure 31 shows how DA bid and ask curve intersection types (vertical, horizontal, or perpendicular) evolved over time, with a marked change after the structural shift.

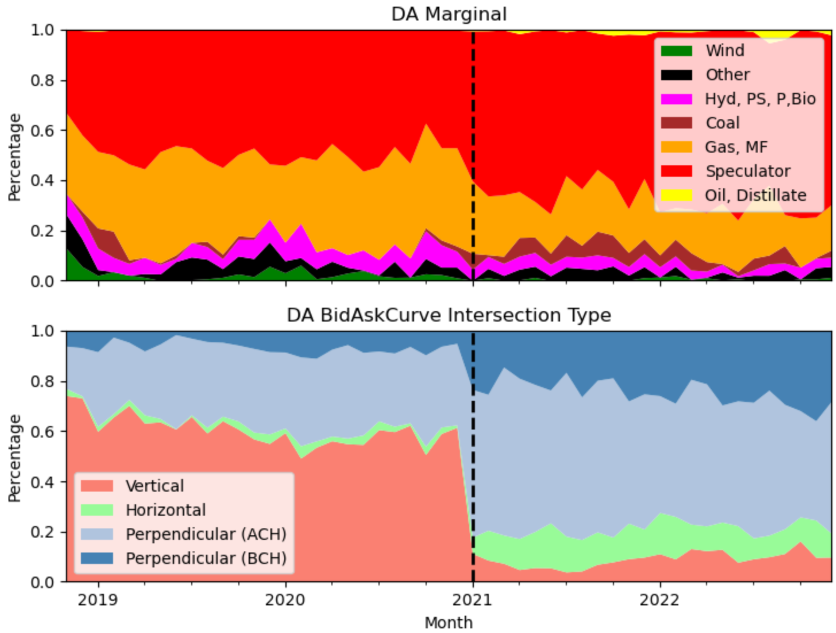

Figure 31: Top plot shows the the percentage of marginal observations in the DA market grouped by FuelType and month. Bottom plot shows the percentage of DA bid and ask curve intersection types by month.

Aggregate Speculator behaviour and profitability To examine aggregate pre-clearing behaviour, we sum DA speculator sell order quantities from the top plot in figure 30 (red series) and split them by price interval (section 5.4.2), shown in the top plot of figure 32. For comparison, the bottom plot shows the same breakdown for non-speculators. For speculators, much of the post-change increase in sell orders is concentrated in the $[P_i^{DA} - SD_i,\ P_i^{DA} + SD_i]$ interval (orange series), whereas the non-speculator data show no clear step change.

Plots of buy orders by interval in figure 33 show similar patterns, with only a slight, visually imperceptible uptick in non-speculator buy orders within the $[P_i^{DA} - SD_i,\ P_i^{DA} + SD_i]$ interval after the structural change. We note in passing that on average 97% of non-speculator buy volumes lie in $[3000,\ P_i^{DA} + SD_i]$, highlighting the inelasticity of demand in that a large proportion of demand is submitted at prices well above the prevailing market price and is therefore relatively insensitive to price variation.

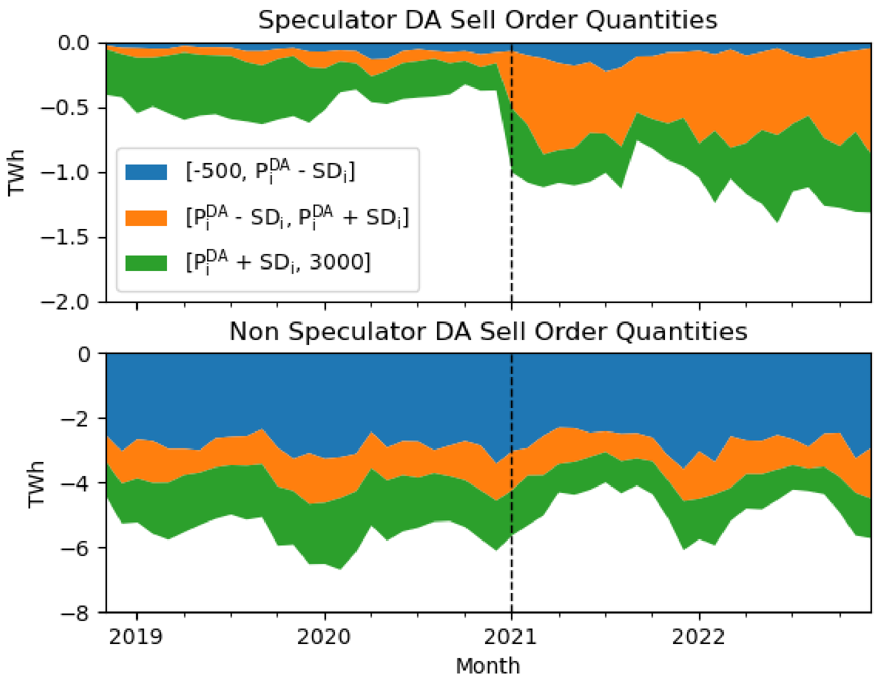

Figure 32: Top plot presents speculator sell order quantities grouped by price interval and month. Bottom plot presents corresponding data for non-speculators. Values are in TWh, note the different y-axis scales.

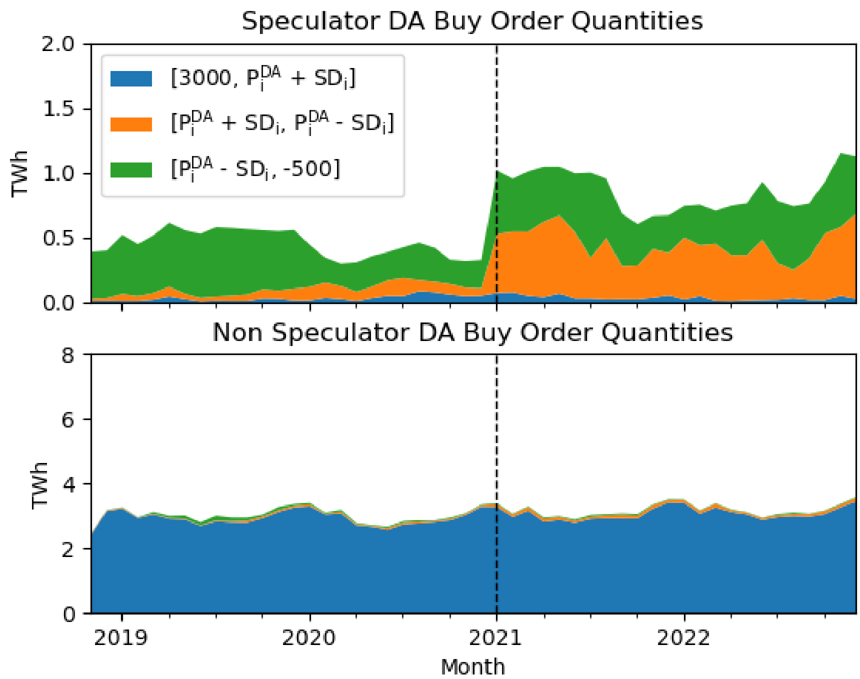

Figure 33: Top plot presents speculator buy order quantities grouped by price interval and month. Bottom plot presents corresponding data for non-speculators. Values are in TWh, note the different y-axis scales.

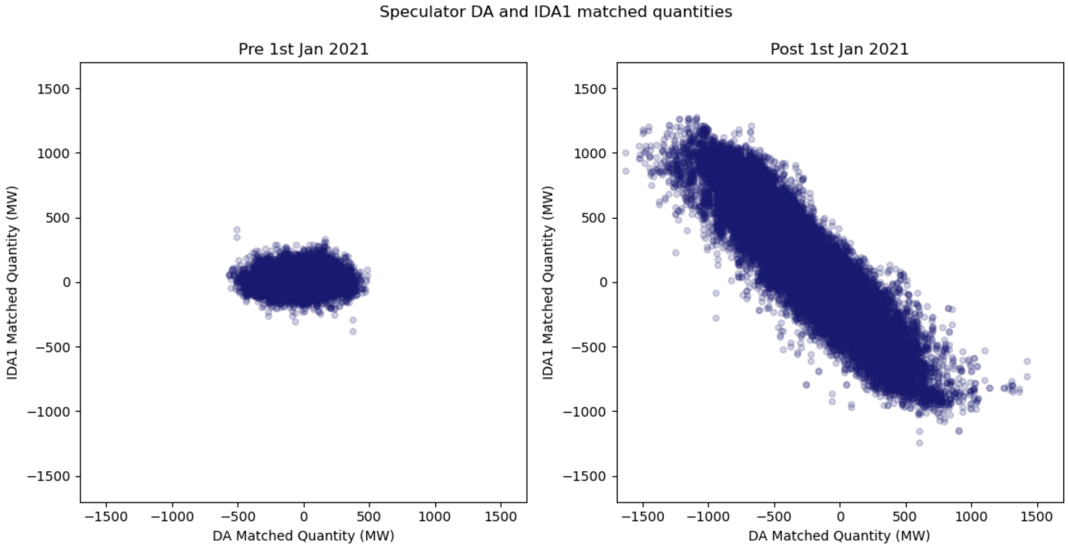

Figure 34: Scatter plot of speculator DA and IDA1 matched quantities which are defined by equations (49) and (50). Values are in MW.

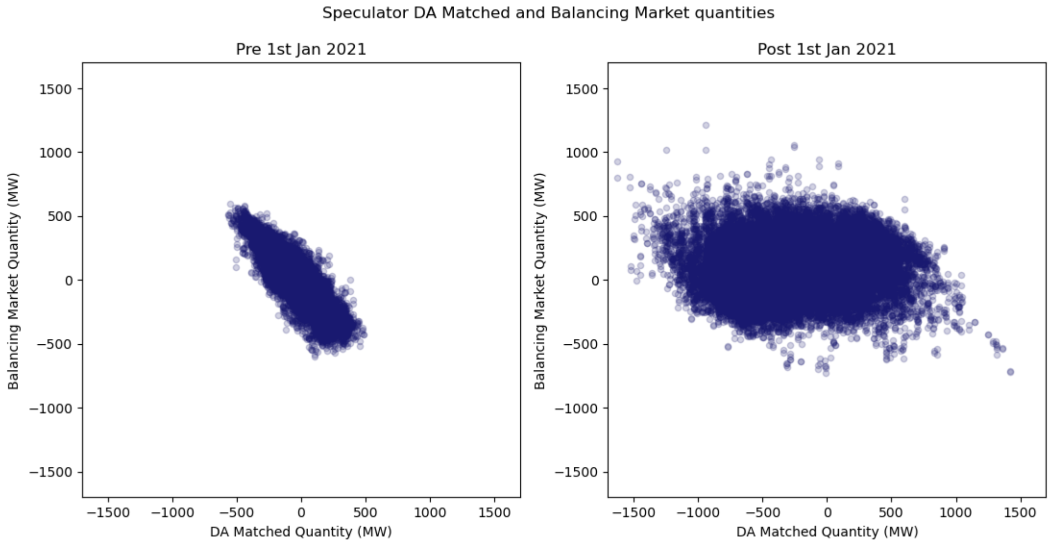

Figure 35: Scatter plot of speculator DA matched and Balancing Market quantities which are defined by equations (49) and (51). Values are in MW.

Figure 34 plots Speculator IDA1 Matched Quantities, equation (50), against DA Matched Quantities, equation (49), distinguishing pre- and post-change observations. Post-change, the magnitude of DA positions increased and a negative correlation emerges: DA matched quantities are often partially or fully unwound in IDA1. Similarly, figure 35 shows Balancing Market Quantities. equation (51), against DA Matched Quantities. Here, a negative correlation is evident before the change, whereas post-change the relationship is less clear.

To illustrate aggregate speculator behaviour after January 2021 and highlight their commercial dynamism, figure 36 presents a two-week granular time series. The top plot shows Speculator DA Matched Quantities (green) and IDA1 Matched Quantities multiplied by -1 (brown), denoted as *IDA1 Matched Quantity (inverted)*. The visible correlation indicates that DA positions are often, though not always, offset in IDA1. The middle plot shows Balancing Market Quantities, which are smaller in magnitude since DA positions are frequently unwound in IDA1. Finally, the bottom plot splits DA matched quantities into buys and sells, revealing heterogeneity across speculators. Taken together, these plots underscore the dynamic and non-stationary nature of speculator activity in the Irish electricity market.

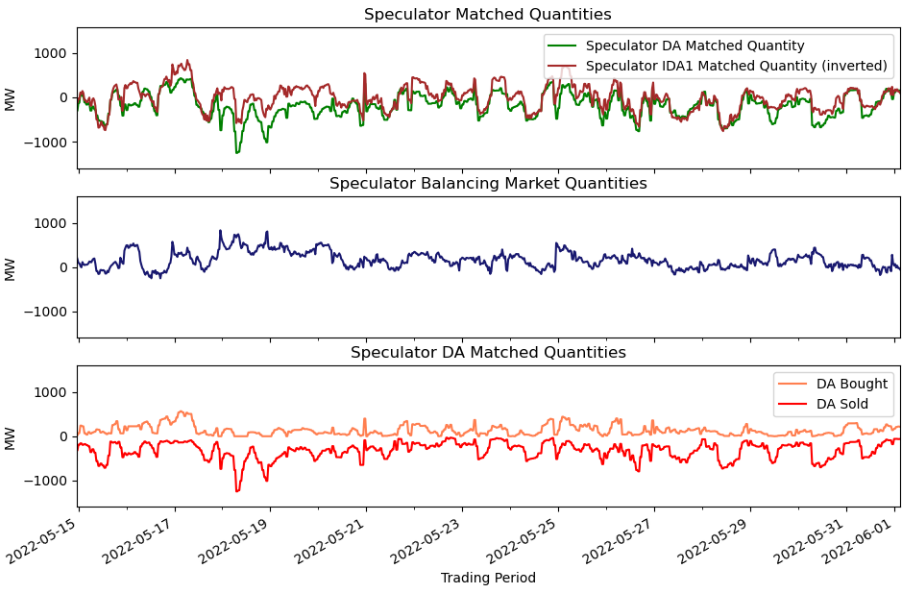

Figure 36: Speculator time series plots for a sample timeframe, units are in MW. The top plot relates to equations (49) and (50) with the latter being inverted (i.e. multiplied by -1). The middle plot is given by equation (51); in the bottom plot we distinguish between Speculator DA Buy Matched Quantities and Speculator DA Sell Matched Quantities.

Using equation (52), aggregate speculator profit and loss (*P&L*) from November 2018 to December 2022 is estimated at €76.8 million[83]. Of this total, €66.2 million relates to 2021–2022. Speculators taking opposing positions in ex-ante markets (see appendix E.4) contributed an estimated €29.9 million. At a granular level (by speculator and trading period), cash flows were positive in approximately 54% of cases and negative in 46%, indicating that at the population level speculators incur losses in a substantial fraction of trading periods and that participation in short-term electricity markets is therefore not risk-free.

[83]Close to €21.76 billion worth of energy has been traded in the ex-ante markets over the same timeframe.

5.5.2 Price Inertia Analysis

Using bid and ask curves from November 2018 to December 2022, we apply a +10 MW horizontal shift to the ask curve and calculate the percentage of trading periods per day with a Price Difference, equation (53), of €0/MWh. Figure 37 shows a clear step change in the share of zero price difference observations after the structural market change. Figure 38 presents the corresponding time series: the top plot shows Price Differences by trading period, the bottom the Custom Metric (equation (54)). The Custom Metric indicates a reduced impact of horizontal shifts from 2021 onward. Overall, a +10 MW shift altered the market price in about 72% of trading periods before 1 January 2021, compared with 36% after. Where prices did change, the average reduction in DA price was 3.6% pre-change and 0.8% post-change. A two-proportion z-test rejects equality of the pre- and post-change proportions[84].

Figure 39 plots the Price Difference and Custom Metric values against DA prices, highlighting shifts in the distributions before and after the structural change. Supporting material includes histograms of these time series (appendix E.5) and the corresponding +1 MW analysis (appendix E.6).

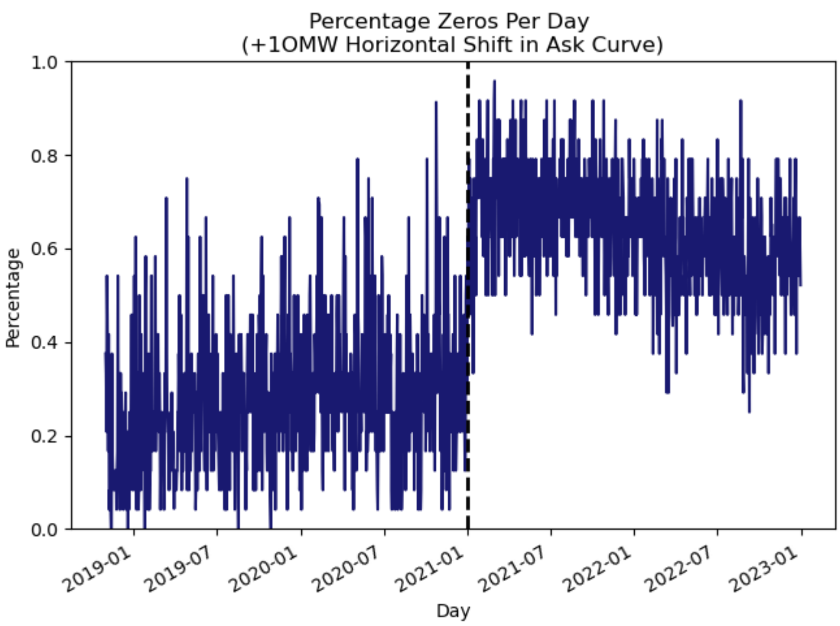

Figure 37: Percentage of trading periods per day where the Price Difference values, defined by equation (53), are €0/MWh.

[84]The associated p-value is $\ll 0.01$.

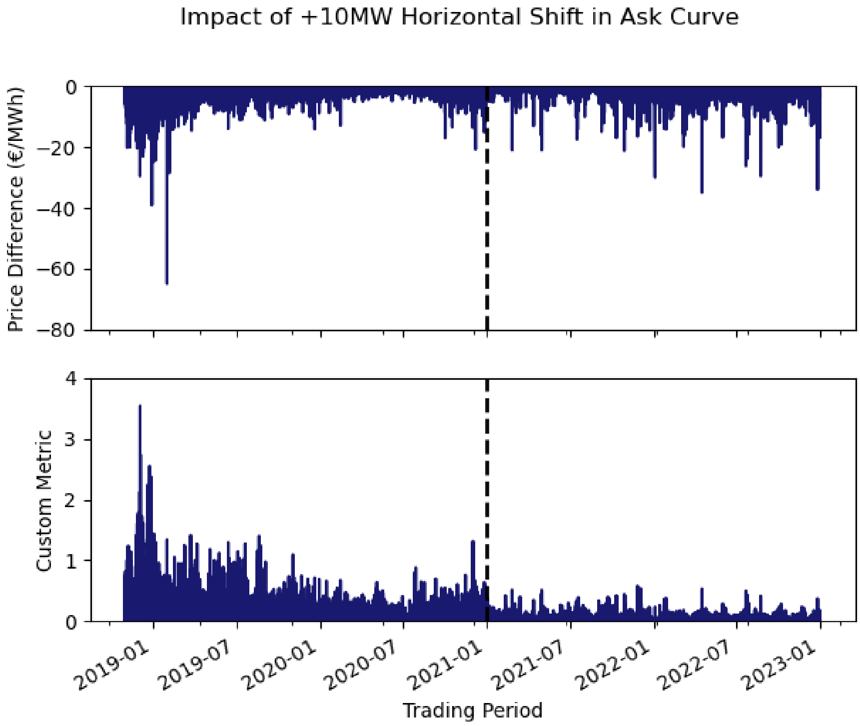

Figure 38: Impact, per trading period, of a +10MW horizontal shift in the ask curve. Values in top plot are given by equation (53), values in bottom plot are given by equation (54).

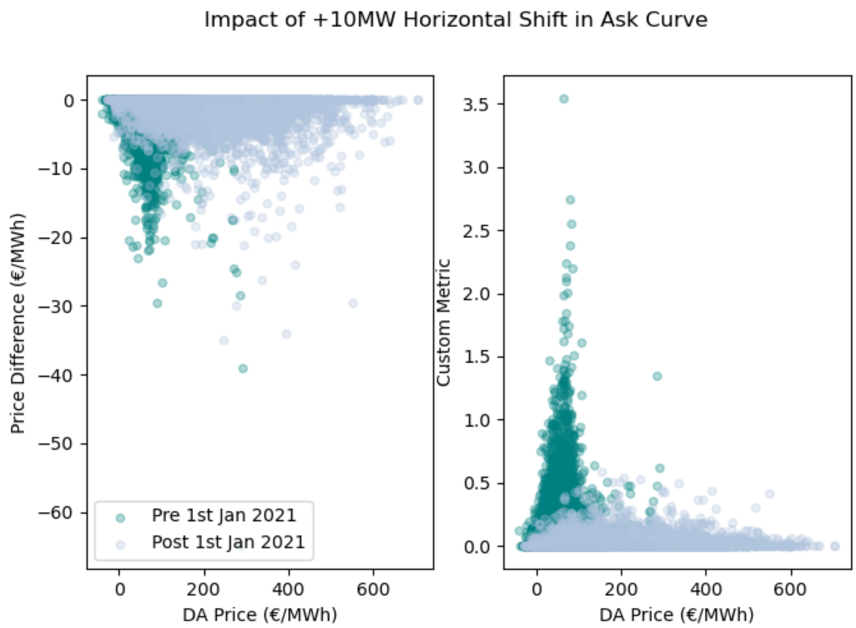

Figure 39: Scatter plot showing the impact, per trading period, of a +10MW horizontal shift in the ask curve distinguishing between observations pre and post 1st January 2021.

Next, Figure 40 presents a density histogram of the Min Shift[85], Price Impact, and Price Impact normalised by the standard deviation, as discussed in Section 5.4.3. Notably, in the left-hand plot, a higher proportion of small shifts resulted in price changes prior to 1 January 2021 compared with the period thereafter.

In figure 41 we analyse the impact of excluding a speculator's orders on DA bid and ask curves for two trading periods on 10th January 2019. Blue lines exclude the speculator, while red lines include them. Since the speculator only submitted sell orders, bid curves remain unchanged. For 4–5 AM, the speculator placed 4MW sell orders at €40/MWh and €50/MWh; for 5–6 AM, similar orders were at €44/MWh and €54/MWh. DA clearing prices were €59.16/MWh and €64.09/MWh, respectively. Excluding the speculator shifted the ask curve left, raising the 4–5 AM price to €61.77/MWh, while the 5–6 AM price remained unchanged. Figure 51 in appendix E.3 extends this analysis across all relevant periods, though, as noted in section 5.4.3, small bid/ask shifts do not always hold.

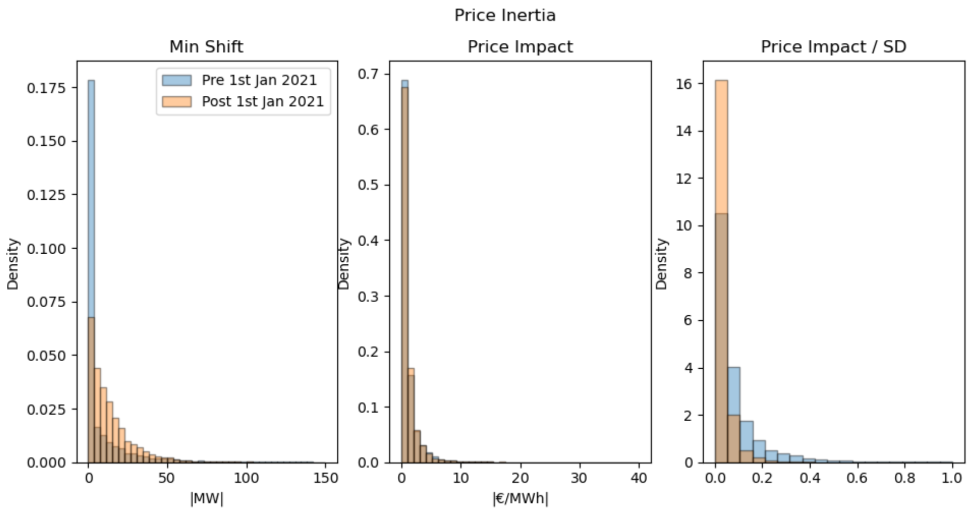

Figure 40: Density histograms of (a) Min Shift (b) Price Impact and (c) Price Impact divided by SD (see section 5.4.3 page 113). (a) and (b) represent the minimum absolute shift needed to trigger a market price change and the corresponding absolute price difference respectively.

[85] The x-axis range in the corresponding density histogram is restricted for display purposes.

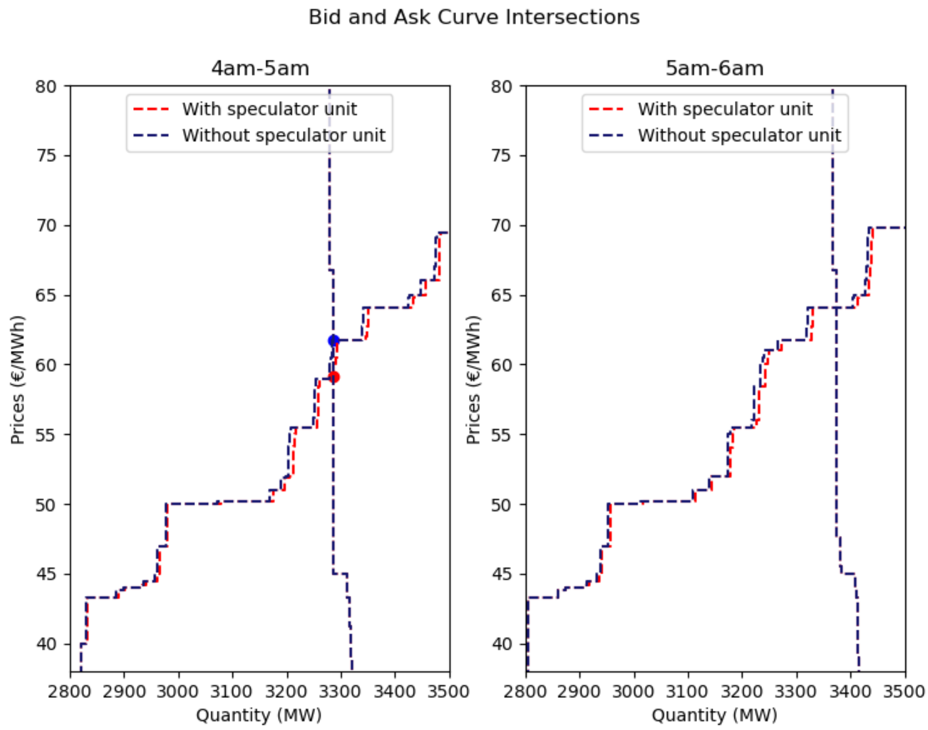

Figure 41: DA bid and ask curve intersections, with and without, a single speculator, 10th January 2019.

5.6 Discussion

The analysis shows that speculators in the Irish DA market exhibit both dynamism and heterogeneity, with participation undergoing a clear step change after the structural market shift. In particular, buy and sell order volumes within the $[P_i^{\text{DA}} - \text{SD}_i, P_i^{\text{DA}} + \text{SD}_i]$ interval increased. At first sight, the frequency with which speculators were marginal, exceeding their overall market share both before and after the change, may seem counterintuitive, but this is consistent with strategies that place orders close to the market clearing price.

The price inertia analysis provided a simple and transparent metric. A relatively high share of trading periods showed that even small shifts in the ask curve could change the market price, and in some cases the shifts produced substantial movements. This underscores the susceptibility of DA prices to local shocks, a feature familiar to market participants. The shift in inertia levels around the structural change was also notable. Conceptually, the approach is similar to liquidity and depth measures in other markets.

A plausible link between these findings is that, after the structural change, speculators began supplying DA liquidity previously provided through EUPHEMIA's interconnector scheduling with Great Britain. This resembles one of the motivations for Virtual Bidding noted by [63], namely liquidity provision and price transparency. Much of this additional liquidity occurred near the clearing price, potentially contributing to higher observed inertia. While suggestive, this remains correlation rather than causation, with two main challenges to formal testing:

- Pricing algorithm access: the market clearing process is complex and not fully disclosed, limiting counterfactual analysis [see also 62].
- Market dynamics: outcomes also reflect evolving behaviours of non-speculators [64]; even modest shifts in their bidding or trading strategies, individually or in aggregate, could contribute to the observed changes.

Further research could test this hypothesis, for example by first classifying trading periods as exhibiting low, medium, or high price inertia, and subsequently estimating a regression or predictive model for these classifications. Explanatory variables could capture market conditions, including the presence and activity of speculators. Statistical significance and explanatory power in such a framework would help assess whether speculator activity contributes materially to observed inertia, conditional on broader market dynamics. More broadly, the results highlight potential limitations in applying fundamental forecasting to short-term markets: speculators' adaptive behaviour complicates model input assumptions, and price inertia implies that small demand or supply changes can trigger non-trivial price movements, nonlinearities that may be beyond the reach of fundamental or other price forecasting approaches.

5.6.1 Limitations

This study has several limitations that should be acknowledged. First, the analysis focuses on the Irish DA market which, while representative of European auction based designs, reflects a relatively small system. Second, while our cohort classification provides a transparent framework for identifying speculators, participants in other cohorts may also engage in speculative activity. Third, the market environment is inherently dynamic, making it difficult to isolate speculator effects from those of other participants; our results should therefore be interpreted as correlational rather than causal. In addition, the DA pricing algorithm is complex and only partly disclosed (section 5.3), preventing full replication of the market clearing process and requiring the simplifying assumptions outlined in section 5.4.3. Finally, alternative participant classification approaches are possible; however, given the robustness checks conducted, we consider our process reasonable and sufficient and do not expect materially different results. These limitations constrain the scope of our conclusions but also indicate promising directions for future research.

5.7 Conclusion

This chapter presented two empirical analyses of the Irish DA electricity market, a setting representative of short-term European markets, characterised by non-trivial liquidity, high renewable penetration, and diverse commodity price regimes over the study horizon. The first analysis examined the presence and behaviour of speculators; the second addressed market price inertia.

For speculators, we found that their share of marginal periods was well above their overall market share, that activity increased following a structural market change, and that their behaviour was both dynamic and heterogeneous. The price inertia analysis introduced an intuitive method that highlighted the sensitivity of short-term prices to small shifts in demand or supply. Inertia levels also increased after the structural change, though the underlying causes remain conjectural given the complexity of the auction mechanism and lack of access to the pricing algorithm.

The research has two key implications. First, both sets of findings point to challenges for short-term fundamental modelling approaches in forecasting electricity prices. Second, the price inertia metric provides a useful additional lens for evaluating structural market change proposals in auction based electricity markets. Overall, our analysis reinforces the view that short-term electricity markets are complex, non-stationary environments shaped by numerous interacting factors.

6 Conclusion

In this final chapter, we briefly revisit the key insights and contributions developed throughout the thesis. Our reflections focus on two distinct but interrelated themes within short-term electricity market settings. First, we consider the role and effectiveness of fundamental modelling approaches. Second, we reflect on the appeal and intuitiveness of the price inertia methodology. Together, these aspects highlight some of the complexities inherent in short-term electricity markets, an issue we touch upon in closing.

The research in chapter 4 focused on addressing key limitations in the literature regarding the use of fundamental modelling approaches in DA electricity market settings. Specifically, the heavy reliance on expert assumptions and the tendency of fundamental models to exclude strategic bidding and trading behaviours. Previous literature indicated that this latter omission contributed to flatter and less volatile price forecasts. This chapter also introduced an explainability framework for forecast performance, addressing another gap in several prior studies. The findings showed that although the granular, bottom-up modelling approach effectively captured market dynamics, it lagged behind top-down methods in price forecast accuracy, largely due to the influence of the speculator cohort. The challenge with the speculator cohort lay in the dynamic and versatile nature of participant behaviour, which proved difficult to capture accurately in the individual-level forecasts; it was also the case that small aggregate deviations between forecasted and actual behaviour could lead to non-trivial impacts on the resulting price forecast. However, when this uncertainty was removed (for example by treating speculator orders as known) the bottom-up approach saw a significant improvement in overall price forecast accuracy. A comparable level of improvement was not observed when similar treatment was applied to other participant cohorts. More broadly, based on industry experience and conversations with market participants, we are aware of the widespread use of fundamental modelling approaches to forecast market prices, analyse asset dispatch, and conduct scenario analysis[86]. Among these capabilities, the ability to conduct scenario-based simulations, such as assessing shifts in the merit order due to changes in fuel input costs or the introduction of new technologies, is particularly valued. We also note the application of fundamental models in short-term electricity market settings; for example [126] provide software solutions that support forecasting across short, medium and long-term horizons. Hence, the research in chapter 4 shows some of the limitations, and the scenarios in which they arise, of

[86] We view fundamental models as part of a broader suite of modelling approaches employed by market participants.

fundamental modelling approaches in short-term electricity market settings.

In chapter 5, the simplicity and intuitiveness of the proposed price inertia methodology offered valuable insight into price sensitivity, that is the market price's responsiveness to small changes in demand or supply. More generally, the concept of price inertia is well understood by market participants, who often monitor or measure it in continuous trading environments. The methodology introduced here shares conceptual similarities with such practices but applies them in the context of electricity auction markets, a setting in which price inertia had not previously been explored in the literature. In the case of the Irish electricity market, applying a small shift to the supply curve resulted in a market price change in approximately 72% of trading periods prior to 1 January 2021 and in around 36% of periods thereafter. This finding may initially seem surprising. Intuitively, one might expect the market to exhibit a degree of robustness, that is sufficient volume would be available at the marginal price to absorb small fluctuations in demand or supply without affecting the market-clearing price. However, the analysis suggests that this is not necessarily the case. Upon further reflection, the result is perhaps less unexpected. When one considers the variation in demand over the course of the trading day, differing wind forecasts, and other shifting factors, it becomes evident that supply and demand curves intersect at significantly different price and quantity points throughout the day. Accordingly, it may follow that the slopes of these curves, and thus the degree of price inertia, vary depending on local market conditions around each intersection point.

A subtle theme linking the research in chapters 4 and 5, and more generally characterising the domain, is the inherent complexity of short-term electricity markets. This complexity manifests in several ways. For example, the price inertia analysis revealed frequent instances of low inertia, implying that small shocks to demand or supply can result in changes to the market-clearing price. The growing presence of speculators, and the clear dynamism in many of their commercial behaviours, adds another layer of unpredictability. A further complexity, which fundamental and demand and supply curve-based modelling approaches tend to simplify or approximate, is the market-clearing algorithm itself. As noted in [62], the pricing mechanisms underpinning DA markets are non-trivial. In contrast, top-down models do not rely on the market-clearing mechanism as they typically forecast prices directly rather than deriving them from underlying bid and offer curves. Additional sources of complexity, and factors contributing to increased dynamism in short-term electricity markets, include the rising share of variable renewable generation, a growing number of market participants and the frequent and rapid shifting in the demand and supply curve stacks. In short, it is this dynamism and evolving market environment that continue to make short-term electricity markets such an active and varied area of research.

We conclude with an observation that links the above findings to ongoing and future research directions. Insights from both the price inertia analysis and the granular fundamental modelling approach resonate with the growing interest in probabilistic electricity price forecasting. In particular, the frequent occurrence of low price inertia periods, combined with increasingly dynamic commercial behaviours, suggests that market participants may place value not only on point forecasts, but also on probabilistic forecasts that capture the likely direction of prices as well as the probability of extreme events. In terms of ongoing work, current efforts are focused on improving the accuracy of forecasts for speculator order behaviour. This includes investigating the use of customised cost functions during model training to better capture the nuances of these participants and enhance overall forecast performance. Looking ahead, a promising direction for further study involves the application of unsupervised learning techniques (e.g.self-organising maps) to identify the commercial strategies of market participants. Preliminary analyses earlier in the research demonstrated that this approach was effective in capturing behavioural patterns in low-dimensional settings (e.g. Approach 1 transformations with $n = 1$). Time permitting, it would be of interest to extend this analysis to higher-dimensional transformations.

A Additional Material: Chapter 3

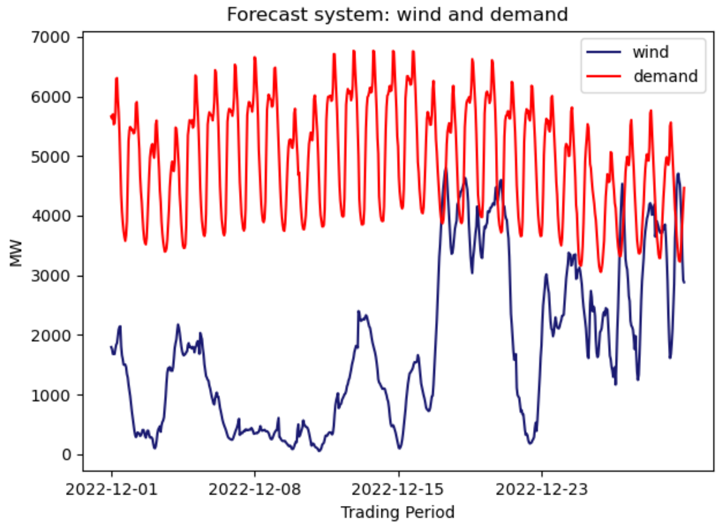

Figure 42: Forecast system wind and demand in December 2022.

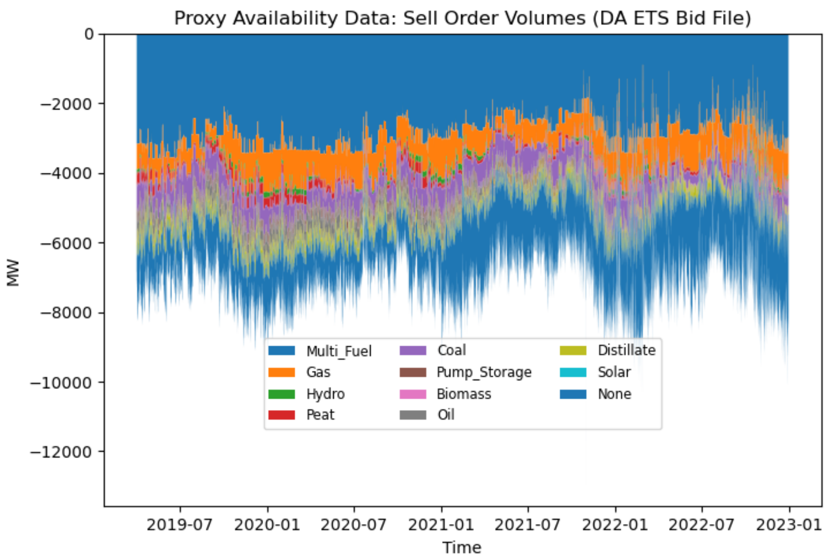

Figure 43: Time series plot of proxy aggregate availability data.

B Participants and Cohorts

The process for splitting participants into cohorts, in addition to identifying speculators, is as follows

- The Registered Capacity Report [78] which contains a list of registered participants is sourced from sem-o. Using the report, if a participant has a

FuelType specified, we assign it to the relevant cohort. To limit the number of cohorts, where appropriate we merge some of these cohorts, see appendix E.8.

- For participants that do not have a specified FuelType, we examine their order data.
  - *buy_sell_units*: if a participant only submits buysell orders, then it is assigned to this cohort.
  - *sell_units*: if a participant only submits sell orders, then it is assigned to this cohort. We also assign participants to this cohort if $\geq 75\%$ of its orders are sells and $\leq 7.5\%$ of it orders are buys.
  - *buy_units*: participants are assigned to this cohort if either (a) participant only submits buy orders or (b) if $\geq 75\%$ of the participant's orders are buys and $\leq 7.5\%$ of it orders are sells.
  - Note: in relation to buy_units and sell_units, the former are likely to correspond to supplier units while the latter are likely to correspond to generator units. We note, however, that speculators could also adopt such strategies (see discussion further below).
  - *wind_demand_units*: This cohort is a comprised of participants that are likely (a) wind participants or (b) wind and demand participants[87]. A participant is deemed to match categorisation (a) if $\geq 75\%$ of its orders are sell and $\leq 30\%$ of its orders have a similar sell quantity in the same trading day[88]. A participant is deemed to match categorisation (b) if it has ≥ 150 buy rows, ≥ 150 sell rows, and if the proportion of similar buy and similar sell quantity observations in the same trading day is $< 7.5\%$.
  - *speculators*: participants not already assigned to another cohort (or if the participant has ≤ 300 orders) are assigned to this cohort.

To strengthen the robustness of the approach, once the population of speculators had been identified, we performed a cursory manual inspection of each unit (i.e., their positions in the ex-ante markets), removing from consideration any units that did not exhibit speculator-type behaviour. The heuristics described above were developed through trial and error and are considered sufficiently robust for our purposes, particularly given that a participant can only belong to one cohort.

[87] Some market participants register a demand portfolio and a wind asset/portfolio under the same participant identifier.

[88] The latter requirement is designed to capture the variable nature of wind

We note, however, two shortcomings of this approach. First, it does not capture speculative actions undertaken by participants with clear generation or consumption roles (e.g., thermal generators, demand). Second, a more comprehensive algorithm could have incorporated additional ex-ante participant information directly into the classification. Nevertheless, in a subsequent review of all participants without an associated FuelType, the procedure performed well with very few exceptions. One such exception was a participant active only from August 2022 (towards the end of the study period), which was incorrectly assigned to the buy cohort but should have been classified as a speculator. Given the small scale of this participant's activity, the misclassification has no material impact on the analysis.

The cohorts and number of participants per cohort are as follows: battery: 8, buy_sell_units: 5, buy_units: 68, coal: 5, hydro_peat_biomass: 16, multi_fuel_gas: 19, oil_distillate: 19, pump_storage: 4, speculator: 50, sell_units: 47, solar: 4, wind: 161,
wind_demand_units: 6.

C Order Data Reconciliation

This appendix provides a transparency and validation exercise illustrating how published BidAskCurve data can be reconciled with participant level order data; it is included for completeness and does not affect the analysis or results presented in the main body of the thesis.

Following back-solving and experimentation effort, it has been possible to use participant stepwise order data to replicate the published BidAskCurve data once the following trading period level adjustments are made

1. Complex Orders are not part of the ask curve, unless the Complex Order is matched. If a Complex Order is matched then the matched quantity is included in the ask curve at the minimum price point.

2. Filter on participant orders which have settlement currency of €. Calculate the difference between the matched buy quantities and matched sell quantities; depending on the sign, the difference needs to be added to either the bid or ask curve at the maximum or minimum price point. Repeat, but for orders which have settlement currency of £.

As of July 2021 onwards, when SEMOpx commenced publishing a single BidAskCurve file, with the bid and ask curve information per trading presented in €/MWh, the adjustments described in step 2 are no longer required.

D Additional Material: Chapter 4

D.1 Modelling Approach details

For replicability purposes we note the following additional details relating to table 2 in section 4.4.1: the RF models have hyperparameter settings {n_estimators, max_depth, min_samples_leaf} set to $\{100, 8, 8\}$. The LASSO models use an *alpha* value of 0.1 with *fit_intercept* set to False. Both SVM models use the RBF kernel, with SVM No. 5 having parameter values $\{C, \text{epsilon}\} = \{1.0, 0.1\}$ and SVM No.6 having $\{10.0, 0.1\}$. The FFNN has 64 neurons in both the first and second layers, a learning rate of 0.01, L2 regularization with a penalty coefficient of 0.01 is applied to both layers. In relation to data scaling, the LASSO and FFNN models scale both the target and explanatory variables whereas in the SVM models only the explanatory data is scaled. For model training, the FFNN models use MAE loss while the RF models use MSE loss, the latter was chosen due to significantly faster compute time compared to RF models trained with MAE loss (as a sense check RF No. 2 was rerun using MAE loss in model training and the test set forecast error metrics showed only minor differences).

D.2 Additional Info

Table 12: Transformation error metrics for different transformation approaches (the corresponding normalised transformation error metrics are shown in section 4.5.1 table 4). The metrics are described in section 4.3.4.

Metric	**Approach 1**						**Approach 2**	
	$n=7$ $(dim=28)$	$n=5$ $(dim=20)$	$n=4$ $(dim=16)$	$n=3$ $(dim=12)$	$n=2$ $(dim=8)$	$n=1$ $(dim=4)$	χ_P $(dim=25)$	χ_P $(dim=15)$
Price MAE	0.4	1.1	2.3	3.9	10.4	32.4	2.4	5.2
Quantity MAE	2.3	6.7	14.5	35.4	93.9	208.4	15.8	31.5
Bid Curve MAE	330	858	1,607	5,053	13,451	29,810	6,372	10,387
Ask Curve MAE	824	2,558	5,430	11,241	29,839	71,725	15,313	24,750

Table 13: The columns indicate the normalised price and quantity MAE values over a January 2021 to December 2022 timeframe that result when we compare price and quantity intersections from published bid and ask curves to those that result from a curve constructed from participant stepwise order data plus transformed order data for the cohort of interest using transformation Approach 1 with $n = 4$. The impact of the remaining cohorts is negligible and hence are not shown. This approach has parallels with that described in the second bullet point in section 4.4.4. Note: multi-fuel gas and coal participants will frequently utilise Complex Order types and hence the caveats referenced in the discussion section 4.6 apply, that is, the table only considers the impact of transforming Simple Order types.

Metric	**Buy**	**Speculator**	**Multi-Fuel Gas**	**Coal**
Price MAE	0.1%	0.9%	0.6%	0.3%
Quantity MAE	0%	0.2%	0.1%	0.1%

Table 14: Forecast error metrics for the modelling approaches described in section 4.4.1, table 2 (the corresponding normalised forecast error metrics are shown in section 4.5.2, table 5). The top-down modelling approach is described in section 4.5.2.

Metric	**Naive** (No. 0)	**RF** (No. 1)	**RF** (No. 2)	**LASSO** (No. 3)	**LASSO** (No. 4)	**SVM** (No. 5)	**SVM** (No. 6)	**FFNN** (No. 7)	**Top-Down** (DNN)
Price MAE	101.8	51.5	54.2	34.3	33.0	50.0	35.8	32.4	24.0
Quantity MAE	380	186	192	163	157	268	171	153	
Bid Curve MAE	95,262	50,107	50,204	52,378	47,054	88,394	53,347	46,662	
Ask Curve MAE	261,052	151,451	157,381	83,416	82,931	136,403	94,531	79,765	

Table 15: Over the Jan 2021 to Dec 2022 horizon, taking the published bid-ask curve and replacing the participant stepwise order data with the corresponding LASSO (No. 4) forecast order data described in table 2, we observe the following (normalised) price and quantity cohort impacts. See section 4.5.3 and table 9 for additional context.

Metric	**Sell**	**Hydro Peat Biomass**	**Battery**	**Buy-Sell**	**Oil-Distillate**	**Solar**	**Wind-Demand**
Price MAE	0.7%	0.5%	0%	0.4%	0.4%	0.1%	0.6%
Quantity MAE	0.2%	0.1%	0%	0.1%	0.1%	0%	0.2%

Table 16: This table is the counterpart to table 10 section 4.5.4, that table reports the rMAE for four distinct participants over the two year forecast horizon. Here we report the MAE for each element of the target vectors.

Participant	MAE τ_Q^b	MAE τ_P^b	MAE τ_Q^s	MAE τ_P^s
Demand	$[21.8, 21.5, 21.5, 21.9]$	$[16.7, 37.5, 29.4, 22.8]$		
Thermal			$[2.1, 2.6, 3.1, 3.6]$	$[267, 259, 259, 261]$
Speculator	$[8.03, 12.9, 21.5, 34.4]$	$[51.1, 42.8, 45.9, 49.9]$	$[8.35, 12.5, 20.8, 30.9]$	$[57.3, 67.8, 76.6, 105.1]$
Pump Storage	$[6.8, , ,]$	$[19.8, , ,]$	$[0.6, 4.9, 9.0,]$	$[65.0, 62.0, 62.0,]$

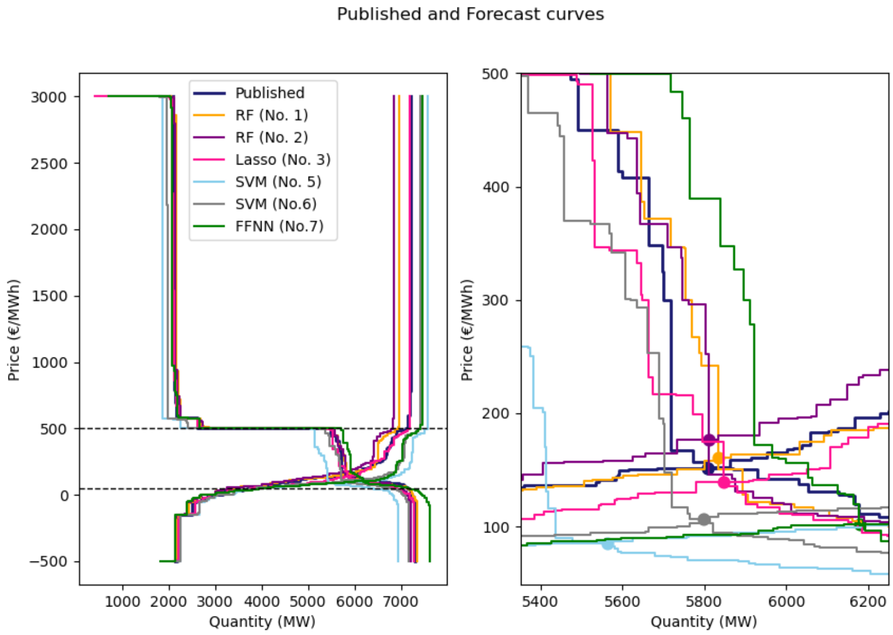

Figure 44: Published and forecast bid and ask curves for modelling approaches in table 2 for trading period 5pm 8th March 2021. Plot on lhs shows the full scale with horizontal lines at the 0th and 100th price quantile; plot on rhs zooms in on the curve intersection regions with the points indicating the intersections. See also section 4.5.2 figure 20.

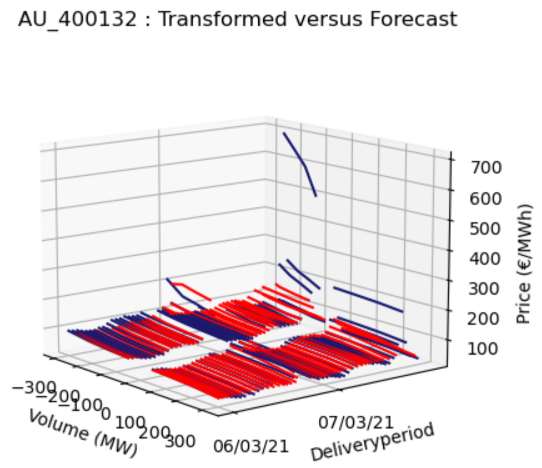

Figure 45: Plot of transformed (blue) and forecast (red) order data for one of the speculator participants (AU_400132) alluded to in section 4.5.3 figure 26. The horizon is from 11pm 6th March to 11pm 8th March 2021. Note the high prices associated with the transformed order data around the 5/6/7pm trading periods on the 8th March.

Table 17: Buy and Sell columns show each cohort's contribution to matched DA quantities (rounded to the nearest percent). For multi-fuel gas, coal, and oil distillate cohorts, approximately 64%, 23%, and 92% of matched sell quantities are Complex Orders.

Cohort	**Buy**	**Sell**
buy	94%	0%
speculator	5%	12%
multi-fuel gas	0%	42%
wind	0%	25%
coal	0%	7%
sell	0%	7%
hydro peat biomass	0%	4%
battery	0%	0%
buy-sell	0%	0%
oil distillate	0%	1%
pump storage	0%	0%
solar	0%	0%
wind demand	0%	2%

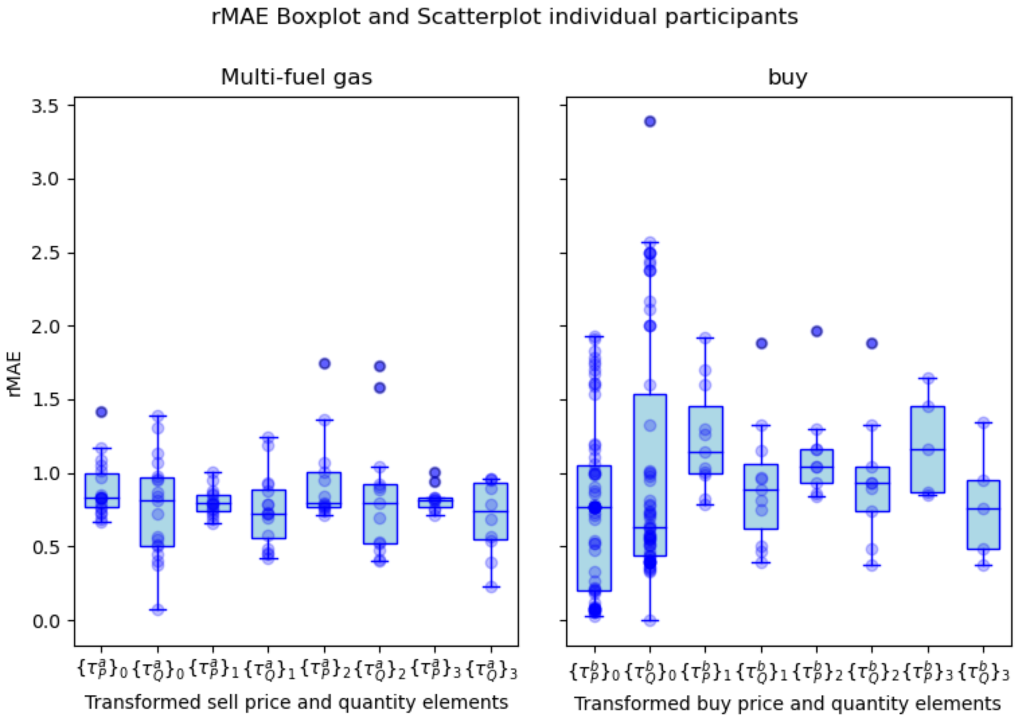

Figure 46: Boxplot and scatterplot of the rMAE's for the individual participants from the (a) multi-fuel gas and (b) buy cohorts for the RF (No.1) modelling approach defined in table 2. For the multi-fuel cohort we show the rMAEs for each element of $\tau_P^a, \tau_Q^a \in \mathbb{R}^4$, for the buy cohort we show the rMAEs for each element of $\tau_P^b, \tau_Q^b \in \mathbb{R}^4$. For the buy cohort there are significantly more observations at index 0 than at subsequent indices.

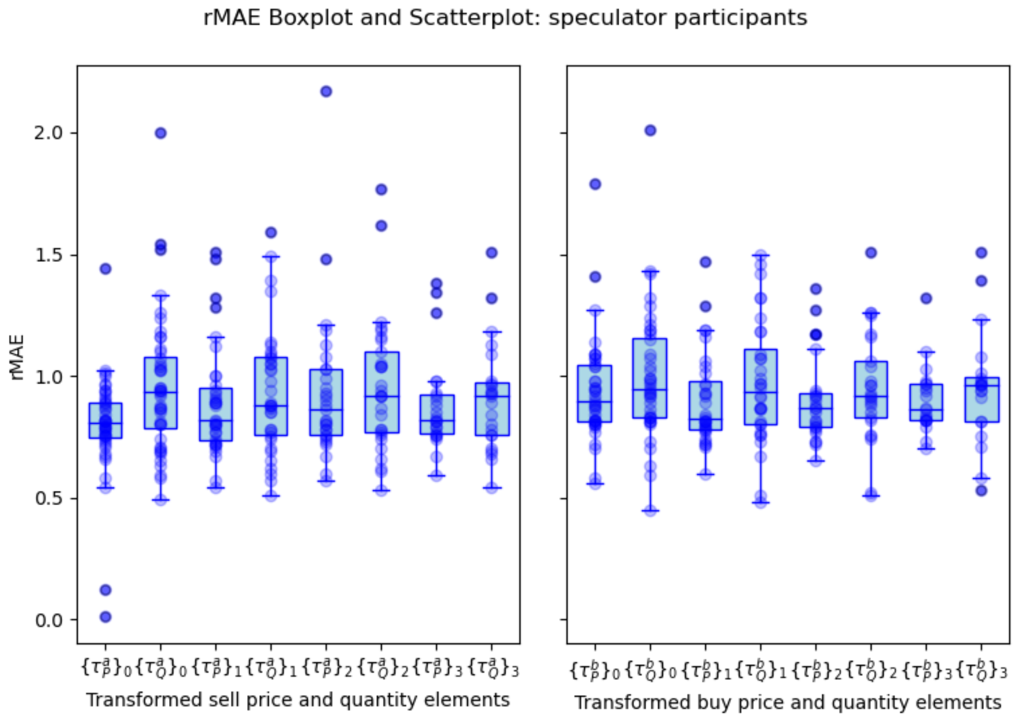

Figure 47: Similar to figure 46 except we present bid and ask rMAE's for individual participants from the speculator cohort for RF (No. 1) modelling approach from table 2.

E Additional Material: Chapter 5

E.1 Price Inertia and Price Volatility

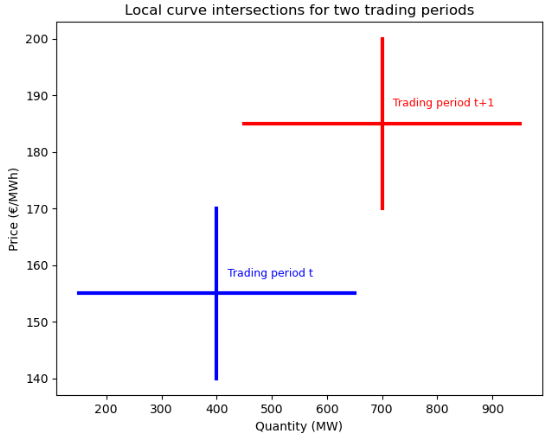

Figure 48: Stylised illustration of the distinction between price inertia and price volatility as used in this thesis. The figure shows hypothetical local supply and demand curve intersections (i.e. portions of the curves in the region of the intersection point) for two consecutive trading periods. The local geometry of the intersections is effectively identical across periods, implying similarly high price inertia: small changes in supply or demand quantities would not alter the market-clearing price. Across periods, however, the intersection occurs at a substantially different price level (approximately €155/MWh in period t and €185/MWh in period $t + 1$), illustrating high price volatility.

E.2 Bid Ask Curve Intersection Examples

Figures 49 and 50 provide examples of Bid Ask Curve intersections. In each plot we zoom in on the region where the bid (blue) and ask (red) curves intersect. Figure 49 presents perpendicular (i.e. point) intersections where it can be seen that either the ask or bid curve is horizontal at the point of intersection. The plot on the left in figure 50 shows a vertical intersection and the DA market price is represented by a black dot.

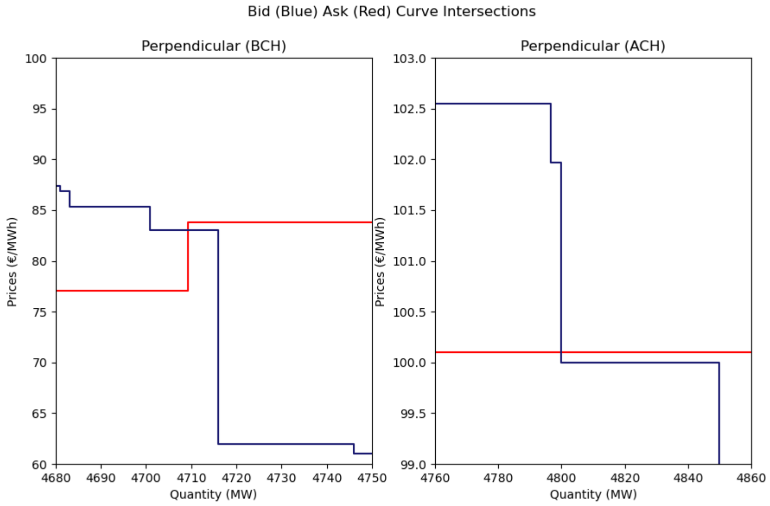

Figure 49: Plot on the left (right) shows the Bid Ask Curve intersection for the I-SEM DA trading period 3pm-4pm (4pm-5pm) 31st October 2018.

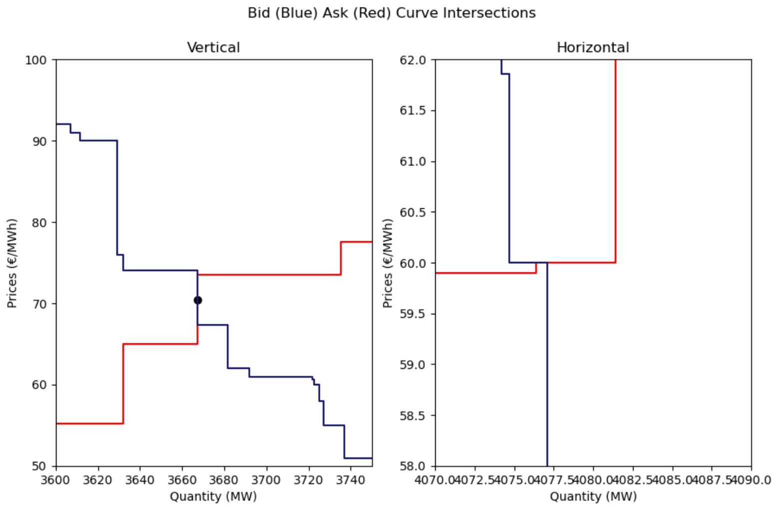

Figure 50: Plot on the left (right) shows the Bid Ask Curve intersection for the I-SEM DA trading period 11pm-12am 30th Oct 2018 (12am-1am 31th Oct 2018).

E.3 Simulating Speculator DA Price Impact

To estimate the impact of a single speculator, j, on the DA price for trading period i, the process is as follows

1. Retrieve all participant orders for trading period i from the relevant DA ETS Bid File. Convert each of the orders into price and quantity pairs.

2. Take the buy price and quantity pairs from step 1 and combine them to produce an aggregated stepwise bid curve. Similarly, take the sell price quantity pairs and combine them to produce an aggregated stepwise ask curve.

3. Adjust the stepwise bid and ask curves from step 2 as described in appendix C.

4. Use the bid and ask curves from step 3 to determine the intersection point/price.

5. Repeat steps 1 to 4 but this time exclude the order data for speculator j.

We note that steps 1 to 4 can essentially be viewed as a validation step.

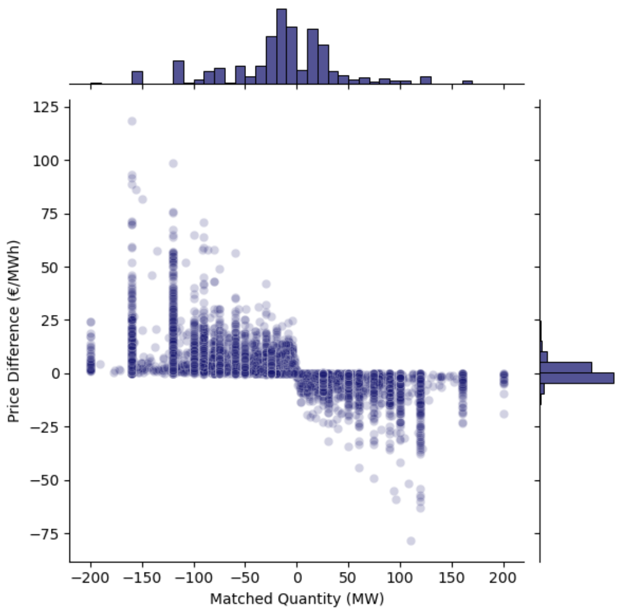

Figure 51: Impact of a single speculator on DA bid and ask curve intersection; negative (positive) x-axis values indicate sell (buy) speculator matched quantities; y-axis equals the intersection price without the speculator minus the intersection price with the speculator.

E.4 Speculator Ex-Ante Trading

Consider the following hypothetical scenarios for speculator j in trading period i:

- Scenario 1: the speculator buys 100MW in DA and sells 100MW in IDA1; in this situation the profit and loss equals $-\frac{1}{2}\left(100P_i^{\mathrm{DA}} - 100P_i^{\mathrm{IDA1}}\right)$.

- Scenario 2: the speculator buys 50MW in DA and sell 150MW in IDA1; using equation (52) from section 5.4.2 the profit and loss is given by $-\frac{1}{2}\left(50P_i^{\mathrm{DA}} - 150P_i^{\mathrm{IDA1}} +\right.$ Rearranging, this is equivalent to $\frac{1}{2}\left(50P_i^{\mathrm{IDA1}} - 50P_i^{\mathrm{DA}}\right) - \frac{1}{2}\left(100P_i^{\mathrm{IDA}} - 100P_i^{B}\right)$.

That is, in both scenarios a proportion of the profit and loss is attributable to the loss or gain associated with taking opposing positions in ex-ante markets.

E.5 +10MW Parallel Shift in Ask Curve

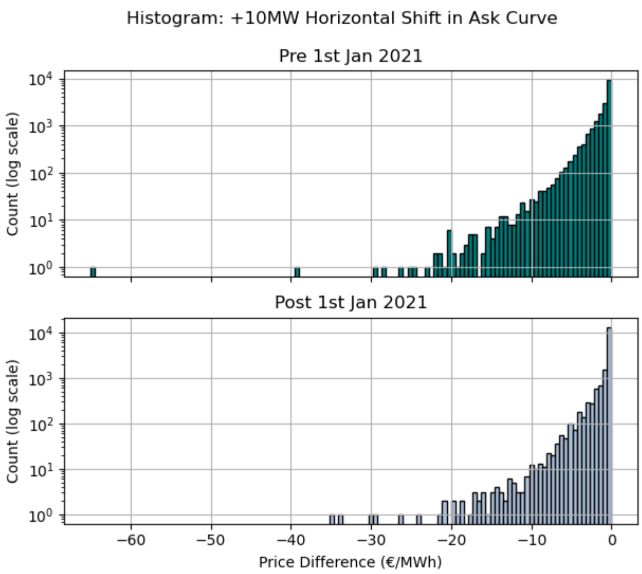

Figure 52: Histogram of Price Difference (€/MWh), equation (53), for a +10MW parallel shift in the ask curve distinguishing between observations pre and post the structural market change. Note that Count is on a log scale.

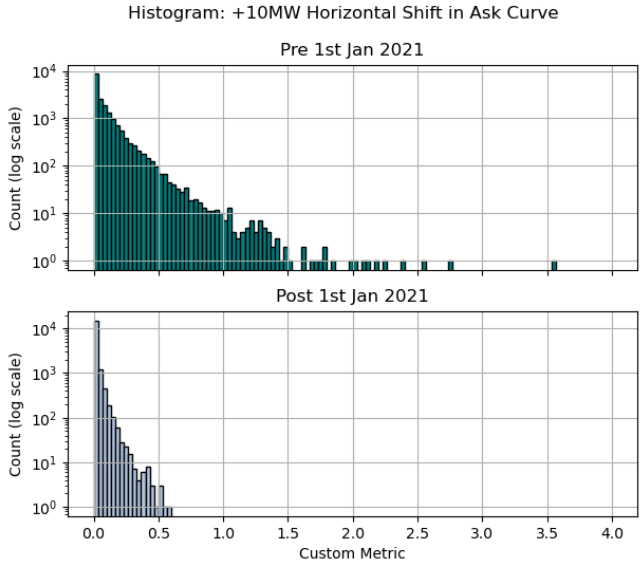

Figure 53: Histogram of Custom Metric, equation (54), for a +10MW parallel shift in the ask curve distinguishing between observations pre and post the structural market change. Note that the Count is on a log scale.

E.6 +1MW Parallel Shift in Ask Curve

The top (bottom) plot in figure 54 shows the Price Difference (Custom Metric) values per trading period associated with a +1MW parallel shift in the ask curve. From the underlying data it is observed that in approximately 62% (12%) of all trading periods pre (post) 1st January 2021 application of a +1MW horizontal shift in the ask curve results in a changed market price. For those trading periods where the market price changes, there is an average reduction of 2.3% (0.7%) in the DA price pre (post) 1st January 2021.

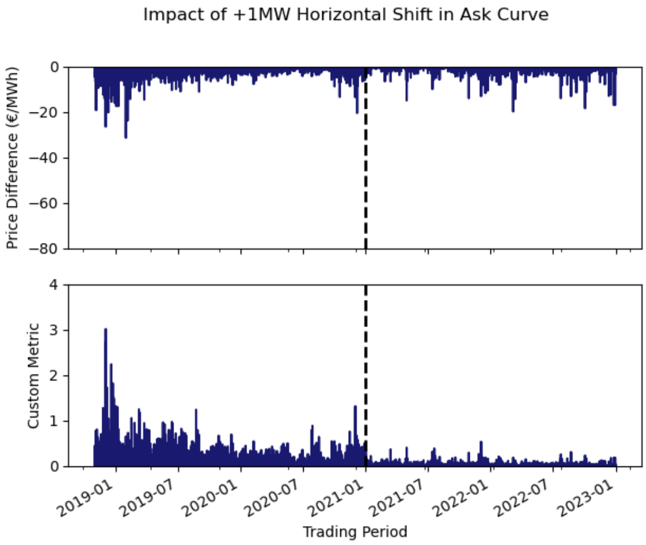

Figure 54: Impact, per trading period, of a +1MW horizontal shift in the ask curve. Top plot shows the Price Difference values defined by equation (53), bottom plot shows the Custom Metric values defined by equation (54).

E.7 Price Inertia Distribution

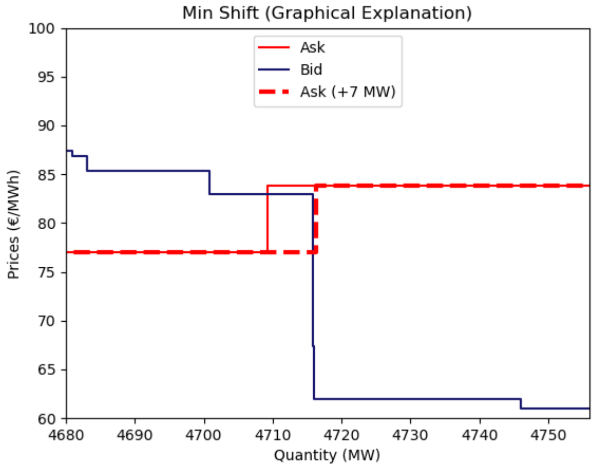

Figure 55: Graphical illustration of the minimum shift approach described in section 5.4.3. Solid lines show the original bid and ask curves; the dashed curve shows the ask curve after the minimum horizontal shift needed to change the market clearing price.

E.8 FuelType Notation

Categories for the FuelType field in the PUB_MnlyRegisteredCapacity file (section 5.4.1) include wind, multi_fuel, gas, hdyro, peat, coal, pump_storage, biomass, oil, distillate and solar (this field can also be unpopulated). When presenting stacked area chart plots in section 5.5, we combine some of the FuelType categories and utilise the following notation:

- Wind category represents those ResourceNames (market participants) where the FuelType is wind.
- Other category represents the units which don't have a FuelType (and from their commercial behaviours appear in the main to be either demand or wind participants).
- Gas, MF category corresponds to gas and multi-fuel thermal generation units.
- Hyd, PS, P, Bio denote hydro, pumped storage, peat and biomass units.
- Speculator represents those units specified in appendix B.

Bibliography

[1] Davis W. Edwards. *Energy Trading and Investing: Trading, Risk Management and Structuring Deals in the Energy Markets*. Accessed: 2025-12-29. McGraw-Hill Education, 2009. ISBN: 9780071629065. URL: https:

//www.mheducation.com/highered/mhp/product/energy-trading-investing.html.

[2] Trafigura. *Commodities Demystified: A Guide to Trading and the Global Supply Chain*. Accessed: 2025-12-29. 2021. URL: https://www.trafigura.com/brochure/commodities-demystified-a-guide-to-trading-and-the-global-supply-chain-second-edition.

[3] Paul L. Joskow. "Lessons Learned from Electricity Market Liberalization". In: *The Energy Journal* 29.2_suppl (2008), pp. 9–42. DOI: 10.5547/ISSN0195-6574-EJ-Vol29-NoSI2-3.

[4] International Energy Agency. *World Energy Outlook*. Accessed: 2025-12-30. 2023. URL: https://www.iea.org/reports/world-energy-outlook-2023.

[5] EUPHEMIA. *EUPHEMIA Public Description: Single Price Coupling Algorithm*. Accessed: 2025-12-29. 2020. URL: http://www.nemo-committee.eu/assets/files/euphemia-public-description.pdf.

[6] Natalia Fabra. "Reforming European electricity markets: Lessons from the energy crisis". In: *Energy Economics* 126 (2023), p. 106963. DOI: 10.1016/j.eneco.2023.106963.

[7] Frank A. Wolak. "An Empirical Analysis of the Impact of Hedge Contracts on Bidding Behavior in a Competitive Electricity Market". In: *International Economic Journal* 14.2 (2000), pp. 1–39.

[8] Koichiro Ito and Mar Reguant. "Sequential Markets, Market Power, and Arbitrage". In: *American Economic Review* 106.7 (2016), pp. 1921–1957. DOI: 10.1257/aer.20141529.

[9] Klaus Mayer and Stefan Trück. "Electricity Markets around the World". In: *Journal of Commodity Markets* 9 (2018), pp. 77–100. DOI: 10.1016/j.jcomm.2018.02.001.

[10] Sanjeev Kumar Aggarwal, Lalit Mohan Saini, and Ashwani Kumar. "Electricity price forecasting in deregulated markets: A review and evaluation". In: *International Journal of Electrical Power & Energy Systems* 31.1 (2009), pp. 13–22. DOI: 10.1016/j.ijepes.2008.09.003.

[11] Rafał Weron. "Electricity price forecasting: A review of the state-of-the-art with a look into the future". In: *International Journal of Forecasting* 30.4 (2014), pp. 1030–1081. DOI: 10.1016/j.ijforecast.2014.08.008.

[12] International Energy Agency. *Gas Market Report: Lessons from the 2022–2023 Energy Crisis*. Accessed: 2025-12-30. 2023. URL: https://www.iea.org/reports/gas-market-lessons-from-the-2022-2023-energy-crisis.

[13] Eurelectric. *How Are EU Electricity Prices Formed and Why Have They Soared?* Accessed: 2025-03-17. 2022. URL: https://www.eurelectric.org/in-detail/electricity_prices_explained/.

[14] Tooraj Jamasb and Michael Pollitt. "Electricity Market Reform in the European Union: Review of Progress toward Liberalization & Integration". In: *The Energy Journal* 26.1_suppl (2005), pp. 11–41. DOI: 10.5547/ISSN0195-6574-EJ-Vol26-NoSI-2.

[15] Michael G. Pollitt. "The role of policy in energy transitions: Lessons from the energy liberalisation era". In: *Energy Policy* 50 (2012), pp. 128–137. DOI: 10.1016/j.enpol.2012.03.004.

[16] European Parliament. *Internal Energy Market*. Accessed: 2025-02-07. 2024. URL: https://www.europarl.europa.eu/factsheets/en/sheet/45/internal-energy-market.

[17] EirGrid Group. *Quick Guide to the Integrated Single Electricity Market*. Accessed: 2025-12-19. URL: https://www.eirgridgroup.com/site-files/library/EirGrid/EirGrid-Group-I-SEM-QuickGuide.pdf.

[18] Giulia Ragosa. "The political economy of electricity market coupling: Comparing experiences from Europe and the United States". In: *The Electricity Journal* 37.7–10 (2024), p. 107430. DOI: 10.1016/j.tej.2024.107430.

[19] ENTSO-E. *Single Day-Ahead Coupling (SDAC)*. Accessed: 2025-02-10. URL: https://www.entsoe.eu/network_codes/cacm/implementation/sdac/.

[20] European Commission. *Energy Efficiency Targets, Directive and Rules*. Accessed: 2025-02-10. 2025. URL: https://energy.ec.europa.eu/topics/energy-efficiency/energy-efficiency-targets-directive-and-rules/energy-efficiency-targets_en.

[21] Sustainable Energy Authority of Ireland. *Ireland's Energy Targets*. Accessed: 2025-02-10. 2025. URL: https://www.seai.ie/about/irelands-energy-targets.

[22] Eurostat. *Renewable Energy Statistics*. Accessed: 2024-02-10. 2024. URL: https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Renewable_energy_statistics.

[23] Beyond Fossil Fuels. *Coal Phase-Out in Europe*. Accessed: 2025-02-10. URL: https://beyondfossilfuels.org/coal/.

[24] Timera Energy. *Ratcliffe closure completes phase out of coal in the UK*. Accessed: 2025-02-10. URL: https://timera-energy.com/blog/ratcliffe-closure-completes-phase-out-of-coal-in-the-uk.

[25] Eurelectric. *Understanding Ultra-Low and Negative Power Prices: Causes, Impacts and Improvements*. Accessed: 2025-02-10. 2024. URL: https://www.eurelectric.org/wp-content/uploads/2024/11/Eurelectric-explainer-on-negative-prices.pdf.

[26] Netherlands Authority for Consumers and Markets (ACM). *Algorithmic Trading in Wholesale Energy Markets*. Accessed: 2025-02-10. 2024. URL: https://www.acm.nl/system/files/documents/rapport-marktstudie-algoritmische-handel-energiemarkt-public_0.pdf.

[27] Gavin Finch, Jason Grotto, Todd Gillespie, and Sam Dodge. *Consumers Foot the Bill for Traders 'Manipulating' the UK Power Market*. Accessed: 2025-12-19. 2023. URL: https://www.bloomberg.com/graphics/2023-uk-power-electricity-market-manipulating/.

[28] ElectroRoute. *The First 150 Days of I-SEM: Part 1 – The Price Spike of January 24th*. Accessed: 2025-02-18. 2019. URL: https://electroroute.com/the-first-150-days-of-i-sem-part-1-the-price-spike-of-january-24th/.

[29] Joseph Collins, Andreas Amann, and Kieran Mulchrone. "A new approach to fundamental electricity market modelling: Exploring market dynamics and speculator influence". In: *Energy and AI* 22 (2025), p. 100639. DOI: 10.1016/j.egyai.2025.100639.

[30] Joseph Collins, Andreas Amann, and Kieran Mulchrone. "Speculators and Price Inertia in a Day-Ahead Electricity Market: An Irish Case Study". In: *Energies* 18.17 (2025), p. 4764. DOI: 10.3390/en18174764.

[31] Joseph Collins, Andreas Amann, and Kieran Mulchrone. *Why Gas and Electricity Prices in Ireland Are So Closely Linked*. Accessed: 2025-12-19. 2022. URL: https://www.rte.ie/brainstorm/2022/0727/1312565-ireland-gas-electricity-prices-correlation/.

[32] Ciaran O'Connor, Joseph Collins, Steven Prestwich, and Andrea Visentin. "Electricity Price Forecasting in the Irish Balancing Market". In: *Energy Strategy Reviews* 54 (2024), p. 101436. DOI: 10.1016/j.esr.2024.101436.

[33] Ciaran O'Connor, Joseph Collins, Steven Prestwich, and Andrea Visentin. "Optimising Quantile-Based Trading Strategies in Electricity Arbitrage". In: *Energy and AI* (2025), p. 100476. DOI: 10.1016/j.egyai.2025.100476.

[34] Philip Beran, Arne Vogler, and Christoph Weber. *Multi-Day-Ahead Electricity Price Forecasting: A Comparison of Fundamental, Econometric and Hybrid Models*. HEMF Working Paper No. 02/2021. Accessed: 2025-12-29. House of Energy Markets and Finance (HEMF), University of Duisburg-Essen, 2021. URL: https://ideas.repec.org/p/dui/wpaper/2102.html.

[35] Jakub Nowotarski and Rafał Weron. "Recent advances in electricity price forecasting: A review of probabilistic forecasting". In: *Renewable and Sustainable Energy Reviews* 81 (2018), pp. 1548–1568. DOI: 10.1016/j.rser.2017.05.234.

[36] Tao Hong, Pierre Pinson, Shu Fan, Hamidreza Zareipour, Alberto Troccoli, and Rob J. Hyndman. "Probabilistic energy forecasting: Global Energy Forecasting Competition 2014 and beyond". In: *International Journal of Forecasting* 32.3 (2016), pp. 896–913. DOI: 10.1016/j.ijforecast.2016.02.001.

[37] Mariano Ventosa, Álvaro Baíllo, Andrés Ramos, and Michel Rivier. "Electricity market modeling trends". In: *Energy Policy* 33.7 (2005), pp. 897–913. DOI: 10.1016/j.enpol.2003.10.013.

[38] Jesus Lago, Grzegorz Marcjasz, Bart De Schutter, and Rafał Weron. "Forecasting Day-Ahead Electricity Prices: A Review of State-of-the-Art Algorithms, Best Practices and an Open-Access Benchmark". In: *Applied Energy* 293 (2021), p. 116983. DOI: 10.1016/j.apenergy.2021.116983.

[39] Jesus Lago, Fjo De Ridder, and Bart De Schutter. "Forecasting spot electricity prices: Deep learning approaches and empirical comparison of traditional algorithms". In: *Applied Energy* 221 (2018), pp. 386–405. DOI: 10.1016/j.apenergy.2018.02.069.

[40] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. 2nd ed. MIT Press, 2018.

[41] Dominik Liebl. "Modeling and Forecasting Electricity Spot Prices: A Functional Data Perspective". In: *The Annals of Applied Statistics* 7.3 (2013), pp. 1562–1592. DOI: 10.1214/13-AOAS652.

[42] John Bower and Derek W. Bunn. “Model-Based Comparisons of Pool and Bilateral Markets for Electricity”. In: *The Energy Journal* 21.3 (2000), pp. 1–29. URL: https://www.jstor.org/stable/41322889.

[43] Christopher Day and Derek Bunn. “Divestiture of Generation Assets in the Electricity Pool of England and Wales: A Computational Approach to Analyzing Market Power”. In: *Journal of Regulatory Economics* 19 (2001), pp. 123–141. DOI: 10.1023/A:1011141105371.

[44] Nick Harder, Ramiz Qussous, and Anke Weidlich. “Fit for Purpose: Modeling Wholesale Electricity Markets Realistically with Multi-Agent Deep Reinforcement Learning”. In: *Energy and AI* 14 (2023), p. 100295. DOI: 10.1016/j.egyai.2023.100295.

[45] Maarten Brinkerink, Paul Deane, Seán Collins, and Brian Ó Gallachóir. “Developing a Global Interconnected Power System Model”. In: *Global Energy Interconnection* 1.3 (2018), pp. 330–343. DOI: 10.14171/j.2096-5117.gei.2018.03.006.

[46] Barry Johnston, Dlzar Al Kez, and Aoife Foley. “Assessing the Effects of Increasing Offshore Wind Generation on Marginal Cost in the Irish Electricity Market”. In: *Applied Energy* 374 (2024), p. 123892. DOI: 10.1016/j.apenergy.2024.123892.

[47] Tom Brown, Jonas Hörsch, and David Schlachtberger. *PyPSA: Python for Power System Analysis*. arXiv preprint arXiv:1707.09913. 2017. URL: https://arxiv.org/abs/1707.09913.

[48] Paul Ghelasi and Florian Ziel. *A Data-Driven Merit Order: Learning a Fundamental Electricity Price Model*. arXiv preprint arXiv:2501.02963. 2025. URL: https://arxiv.org/abs/2501.02963.

[49] Christian Pape, Simon Hagemann, and Christoph Weber. “Are Fundamentals Enough? Explaining Price Variations in the German Day-Ahead and Intraday Power Market”. In: *Energy Economics* 54 (2016), pp. 376–387. DOI: 10.1016/j.eneco.2015.12.013.

[50] Philip Beran, Christian Pape, and Christoph Weber. “Modelling German Electricity Wholesale Spot Prices with a Parsimonious Fundamental Model – Validation & Application”. In: *Utilities Policy* 58 (2019), pp. 27–39. DOI: 10.1016/j.jup.2019.01.008.

[51] Florian Ziel and Rick Steinert. “Electricity price forecasting using sale and purchase curves: The X-Model”. In: *Energy Economics* 59 (2016), pp. 435–454. DOI: 10.1016/j.eneco.2016.08.008.

[52] Simon Schnürch and Andreas Wagner. *Machine Learning on EPEX Order Books: Insights and Forecasts*. arXiv preprint arXiv:1906.06248. 2019. URL: https://arxiv.org/abs/1906.06248.

[53] Hamid Zareipour. "Short-Term Electricity Market Prices: A Review of Characteristics and Forecasting Methods". In: *Handbook of Networks in Power Systems I*. Ed. by Alexey Sorokin, Steffen Rebennack, Panos M. Pardalos, Niko A. Iliadis, and Mario V. F. Pereira. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 89–121. ISBN: 978-3-642-23193-3. DOI: 10.1007/978-3-642-23193-3_4.

[54] Shutang Fan, Cheng Mao, and Lijun Chen. "Next-day electricity-price forecasting using a hybrid network". In: *IET Generation, Transmission & Distribution* 1 (2007), pp. 176–182. DOI: 10.1049/iet-gtd:20060006.

[55] Teuvo Kohonen. *Self-Organizing Maps*. Springer-Verlag Berlin Heidelberg, 1997. ISBN: 978-3-540-67921-9. DOI: 10.1007/978-3-642-56927-2.

[56] Dogan Keles, Jonathan Scelle, Florentina Paraschiv, and Wolf Fichtner. "Extended forecast methods for day-ahead electricity spot prices applying artificial neural networks". In: *Applied Energy* 162 (2016), pp. 218–230. DOI: 10.1016/j.apenergy.2015.09.087.

[57] Umut Ugurlu, Ilkay Oksuz, and Oktay Tas. "Electricity Price Forecasting Using Recurrent Neural Networks". In: *Energies* 11.5 (2018), p. 1255. DOI: 10.3390/en11051255.

[58] Bartosz Uniejewski, Jakub Nowotarski, and Rafał Weron. "Automated Variable Selection and Shrinkage for Day-Ahead Electricity Price Forecasting". In: *Energies* 9.8 (2016), p. 621. DOI: 10.3390/en9080621.

[59] Florian Ziel and Rafał Weron. "Day-ahead electricity price forecasting with high-dimensional structures: Univariate vs. multivariate modeling frameworks". In: *Energy Economics* 70 (2018), pp. 396–420. DOI: 10.1016/j.eneco.2017.12.016.

[60] Marco Haase, Yvonne Seiler Zimmermann, and Heinz Zimmermann. "The Impact of Speculation on Commodity Futures Markets – A Review of the Findings of 100 Empirical Studies". In: *Journal of Commodity Markets* 3.1 (2016), pp. 1–15. DOI: 10.1016/j.jcomm.2016.07.006.

[61] Julie Creswell and Robert Gebeloff. *Traders Profit as Power Grid Is Overworked*. Accessed: 2025-02-22. 2014. URL: https://www.nytimes.com/2014/08/15/business/energy-environment/traders-profit-as-power-grid-is-overworked.html.

[62] John Parsons, Cathleen Colbert, Jeremy Larrieu, Taylor Martin, and Erin Mastrangelo. *Financial Arbitrage and Efficient Dispatch in Wholesale Electricity Markets*. SSRN working paper. 2015. DOI: 10.2139/ssrn.2574397.

[63] Akshaya Jha and Frank A. Wolak. *Can Financial Participants Improve Price Discovery and Efficiency in Multi-Settlement Markets with Trading Costs?* Tech. rep. 24713. National Bureau of Economic Research, 2019. DOI: 10.3386/w25851.

[64] Ignacia Mercadal. "Dynamic Competition and Arbitrage in Electricity Markets: The Role of Financial Players". In: *American Economic Journal: Microeconomics* 14.3 (2022), pp. 665–699. DOI: 10.1257/mic.20190276.

[65] John R. Birge, Ali Hortaçsu, Ignacia Mercadal, and J. Michael Pavlin. "Limits to Arbitrage in Electricity Markets: A Case Study of MISO". In: *Energy Economics* 75 (2018), pp. 518–533. DOI: 10.1016/j.eneco.2018.08.024.

[66] Jörgen Hellström, Jens Lundgren, and Haishan Yu. "Why Do Electricity Prices Jump? Empirical Evidence from the Nordic Electricity Market". In: *Energy Economics* 34.6 (2012), pp. 1774–1781. DOI: 10.1016/j.eneco.2012.07.006.

[67] Katarzyna Maciejowska. "Fundamental and Speculative Shocks, What Drives Electricity Prices?" In: *Proceedings of the 11th International Conference on the European Energy Market (EEM 2014)*. 2014, pp. 1–5.

[68] Mikael Bask and Anna Widerberg. "Market structure and the stability and volatility of electricity prices". In: *Energy Economics* 31.2 (2009), pp. 278–288. DOI: 10.1016/j.eneco.2008.11.006.

[69] Carl J. Ullrich. "Realized Volatility and Price Spikes in Electricity Markets: The Importance of Observation Frequency". In: *Energy Economics* 34.6 (2012), pp. 1809–1818. DOI: 10.1016/j.eneco.2012.07.003.

[70] Erkan Erdogdu. "Asymmetric Volatility in European Day-Ahead Power Markets: A Comparative Microeconomic Analysis". In: *Energy Economics* 56 (2016), pp. 398–409. DOI: 10.1016/j.eneco.2016.04.002.

[71] EirGrid. *I-SEM Industry Guide*. Accessed: 2025-03-03. 2017. URL: https://www.sem-o.com/documents/general-publications/I-SEM-Industry-Guide.pdf.

[72] SEMOpx. *Monthly Market Report – January 2025*. Accessed: 2025-12-29. 2025. URL: https://www.semopx.com/sites/semopx/files/2025-03/MONTHLY%20mkt%20report%20January%202025.pdf.

[73] Our World in Data. *Energy Data: Electricity Only (2023)*. Accessed: 2025-02-25. 2025. URL: https://ourworldindata.org/energy.

[74] NEMO Committee. *Euphemia Public Description*. Accessed: 2025-02-25. 2025. URL: https://www.nemo-committee.eu/assets/files/euphemia-public-description.pdf.

[75] ElectroRoute. *Interconnecting the Dots on Future Brexit Arrangements*. Accessed: 2025-02-25. 2025. URL: https://electroroute.com/interconnecting-the-dots-on-future-brexit-arrangements/.

[76] SEMOpx. *SEMOpx Bidding*. Accessed: 2025-02-25. 2025. URL: https://www.semopx.com/documents/training/SEMOpx-Bidding.pdf.

[77] SEMOpx. *SEMOpx Data Publication Guide Issue 7.7*. Accessed: 2025-02-27. 2025. URL: https://www.semopx.com/market-data/document-library/index.xml?keyword=SEMOpx+Data+Publication+Guide.

[78] SEMO. *SEMO Data Publication Guide*. Accessed: 2025-02-25. 2025. URL: https://www.sem-o.com/market-messages/semo-data-publication-guide.

[79] Smart Grid Dashboard. *Interconnection Data*. Accessed: 2025-02-25. 2025. URL: https://www.smartgriddashboard.com/#all/interconnection.

[80] Gas Networks Ireland. *Imbalance Prices - Balancing Actions and Prices Dashboard*. Accessed: 2025-02-25. 2025. URL: https://www.gasnetworks.ie/corporate/gas-regulation/transparency-and-publicat/dashboard-reporting/balancing-actions-and-prices/imbalance-prices/.

[81] Investing.com. *Carbon Emissions Historical Data*. [Accessed: 2025-02-27]. 2025. URL: https://www.investing.com/commodities/carbon-emissions-historical-data.

[82] SEMOpx. *SEMOpx Operating Procedures*. Accessed: 2025-02-26. 2025. URL: https://www.semopx.com/rules-and-monitoring/market-rules/.

[83] SEM Committee. *Participant Technical Data 2019-2025 - Public*. Accessed: 2025-02-27. 2020. URL: https://www.semcommittee.com/publications/sem-20-003-sem-plexos-forecast-model-2020-2025.

[84] EirGrid. *EirGrid Smart Grid Dashboard: Interconnection*. Accessed: 2025-12-19. URL: https://www.smartgriddashboard.com/all/interconnection/.

[85] Nord Pool Group. *About the N2EX Day-Ahead Auction*. Accessed: 2025-02-27. 2025. URL: https://support.nordpoolgroup.com/support/solutions/articles/8000088463-about-the-n2ex-day-ahead-auction.

[86] Aurélien Géron. *Hands-On Machine Learning with Scikit-Learn and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems.* 1st ed. O'Reilly Media, Inc., 2017. ISBN: 978-1491962299.

[87] Trevor Hastie, Robert Tibshirani, and Jerome H. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction.* 2nd ed. Springer, 2009. URL: https://web.stanford.edu/~hastie/ElemStatLearn/.

[88] Adele Cutler, David Cutler, and John Stevens. "Random Forests". In: *Ensemble Machine Learning.* Springer, 2011, pp. 157–176. ISBN: 978-1-4419-9325-0. DOI: 10.1007/978-1-4419-9326-7_5.

[89] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning.* MIT Press, 2016. URL: https://www.deeplearningbook.org.

[90] Michael A. Nielsen. *Neural Networks and Deep Learning.* Determination Press, 2018. URL: https://neuralnetworksanddeeplearning.com/.

[91] François Chollet. *Deep Learning with Python.* Shelter Island, NY: Manning Publications, 2017. ISBN: 978-1-61729-443-3. URL: https://www.manning.com/books/deep-learning-with-python.

[92] Catherine F. Higham and Desmond J. Higham. "Deep Learning: An Introduction for Applied Mathematicians". In: *SIAM Review* 61.3 (2019), pp. 860–891. DOI: 10.1137/18M1165748.

[93] Xavier Glorot and Yoshua Bengio. "Understanding the Difficulty of Training Deep Feedforward Neural Networks". In: *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics (AISTATS).* 2010.

[94] Robert Tibshirani, Trevor Hastie, and Jerome Friedman. "Regularized Paths for Generalized Linear Models via Coordinate Descent". In: *Journal of Statistical Software* 33.1 (2010), p. 1. DOI: 10.18637/jss.v033.i01.

[95] Alex J. Smola and Bernhard Schölkopf. "A Tutorial on Support Vector Regression". In: *Statistics and Computing* 14 (2004), pp. 199–222. DOI: 10.1023/B:STCO.0000035301.49549.88.

[96] Steve R. Gunn. *Support Vector Machines for Classification and Regression.* Tech. rep. Accessed: 2025-12-29. University of Southampton, Department of Electronics and Computer Science, 1997. URL: https://eprints.soton.ac.uk/256459/.

[97] Brian Publicover. *Europe Posts Record Negative Power Prices for 2024 as Renewables Rise*. PV Magazine; accessed: 2025-03-10. 2025. URL: https://www.pv-magazine.com/2025/01/21/europe-posts-record-negative-power-prices-for-2024-as-renewables-rise/.

[98] Matthew Dalton. *Electricity That Costs Nothing—or Even Less? It's Happening More and More*. The Wall Street Journal; accessed: 2025-03-10. 2024. URL: https://www.wsj.com/business/energy-oil/electricity-that-costs-nothingor-even-less-its-happening-more-and-more-53f16e49.

[99] Chris Johnathon, Ashish Prakash Agalgaonkar, Joel Kennedy, and Chayne Planiden. "Analyzing Electricity Markets with Increasing Penetration of Large-Scale Renewable Power Generation". In: *Energies* 14.22 (2021), p. 7618. DOI: 10.3390/en14227618.

[100] Johanna Cludius, Hauke Hermann, Felix Chr. Matthes, and Verena Graichen. "The Merit Order Effect of Wind and Photovoltaic Electricity Generation in Germany 2008–2016: Estimation and Distributional Implications". In: *Energy Economics* 44 (2014), pp. 302–313. DOI: 10.1016/j.eneco.2014.04.020.

[101] Carlo Brancucci Martinez-Anido, Greg Brinkman, and Bri-Mathias Hodge. "The Impact of Wind Power on Electricity Prices". In: *Renewable Energy* 94 (2016), pp. 474–487. DOI: 10.1016/j.renene.2016.03.053.

[102] Katarzyna Maciejowska. "Assessing the Impact of Renewable Energy Sources on the Electricity Price Level and Variability – A Quantile Regression Approach". In: *Energy Economics* 85 (2020), p. 104532. DOI: 10.1016/j.eneco.2019.104532.

[103] NERA Economic Consulting. *SEM-21-086 SEM PLEXOS Model Validation (2021–2029) and Backcast Report*. Accessed: 2025-03-10. 2021. URL: https://www.semcommittee.com/publications/sem-21-086-sem-plexos-model-validation-2021-2029-and-backcast-report.

[104] Derek W. Bunn and Francisco S. Oliveira. "Agent-Based Simulation—An Application to the New Electricity Trading Arrangements of England and Wales". In: *IEEE Transactions on Evolutionary Computation* 5.5 (2001), pp. 493–503. DOI: 10.1109/4235.956713.

[105] SEMOpx. *SEMOpx DAM Information Note: Maximum Price Increase (12 April 2022)*. Accessed: 2025-03-10. 2022. URL: https://www.semopx.com/documents/general-publications/SEMOpx-Info-Maximum-Price-increase.pdf.

[106] SEMOpx. *SEMOpx DAM Information Note: Maximum Price Increase (30 August 2022)*. Accessed: 2025-03-19. 2022. URL: https://www.semopx.com/sites/semo/files/documents/general-publications/SEMOpx-Info-Maximum-Price-increase-310822.pdf.

[107] Zsuzsanna Csereklyei. "Price and income elasticities of residential and industrial electricity demand in the European Union". In: *Energy Policy* 137 (2020), p. 111079. DOI: 10.1016/j.enpol.2019.111079.

[108] Lion Hirth. "The Market Value of Variable Renewables: The Effect of Solar and Wind Power Variability on Their Relative Price". In: *Energy Economics* 38 (2013), pp. 218–236. DOI: 10.1016/j.eneco.2013.02.004.

[109] Michał Narajewski and Florian Ziel. "Econometric Modelling and Forecasting of Intraday Electricity Prices". In: *Journal of Commodity Markets* 19 (2020), p. 100107. DOI: 10.1016/j.jcomm.2019.100107.

[110] Mariusz Kaleta. "Robust Co-Optimization of Medium- and Short-Term Electrical Energy and Flexibility in Electricity Clusters". In: *Energies* 18.3 (2025), p. 479. DOI: 10.3390/en18030479.

[111] Emanuel Canelas, Tânia Pinto-Varela, and Bartosz Sawik. "Electricity portfolio optimization for large consumers: Iberian electricity market case study". In: *Energies* 13.9 (2020), p. 2249. DOI: 10.3390/en13092249.

[112] Elexon. *BSC Insights: What Is Driving Increases in Electricity Imbalance Volumes*. Elexon Insights blog, July 2019, Accessed: 2025-03-17. 2019. URL: https://www.elexon.co.uk/operations-settlement/balancing-and-settlement/elexon-insights-what-is-driving-increases-in-electricity-imbalance-volumes-july-2019/.

[113] Elexon. *Investigations into Net Imbalance Volume Chasing (NIV Chasing) in the GB Electricity Market*. BSC Panel Paper 294/09, Accessed: 2025-03-17. 2019. URL: https://www.elexon.co.uk/documents/groups/panel/2019-meetings/294-september/294-09-investigations-into-niv-chasing-in-gb-electricity-market-shadrack-nyarko/.

[114] Kun Li, Joseph D. Cursio, Mengfei Jiang, and Xi Liang. "The Significance of Calendar Effects in the Electricity Market". In: *Applied Energy* 235 (2019), pp. 487–494. DOI: 10.1016/j.apenergy.2018.10.124.

[115] Frank Sensfuß, Mario Ragwitz, and Massimo Genoese. "The Merit-Order Effect: A Detailed Analysis of the Price Effect of Renewable Electricity Generation on Spot Market Prices in Germany". In: *Energy Policy* 36.8 (2008), pp. 3086–3094. DOI: 10.1016/j.enpol.2008.03.035.

[116] Aitor Ciarreta and María Paz Espinosa. "Market power in the Spanish electricity auction". In: *Journal of Regulatory Economics* 37.1 (2010), pp. 42–69. DOI: 10.1007/s11149-009-9102-7.

[117] Energy Traders Europe (EFET). *Insight into Marginal Pricing in Wholesale Electricity Markets*. Accessed: 2025-03-17. 2022. URL: https://cms.energytradersdeutschland.org/storage/uploads/media/220601-energy-traders-europe-insight-14-spot-trading-electricity.pdf.

[118] UK Department for Energy Security and Net Zero. *Review of Electricity Market Arrangements (REMA)*. Government consultation. 2022. URL: https://www.gov.uk/government/consultations/review-of-electricity-market-arrangements.

[119] The Guardian. *'I've fallen out with people': the bruising debate over UK zonal energy pricing*. Accessed: 2025-05-23. Oct. 2024. URL: https://www.theguardian.com/money/2024/oct/07/ive-fallen-out-with-people-the-bruising-debate-over-uk-zonal-energy-pricing.

[120] Mehdi Madani. *Single Day-Ahead Coupling: Important Successes for the Euphemia Lab*. Accessed: 2025-03-17. 2020. URL: https://energy.n-side.com/blog/single-day-ahead-coupling-important-successes-for-the-euphemia-lab.

[121] NEMO Committee. *CACM Annual Report 2021*. Accessed: 2025-03-17. 2021. URL: https://www.nemo-committee.eu/assets/files/nemo_CACM_Annual_Report_2021_220630-4e7321983974b812f28730a301c9f7d9.pdf.

[122] V. Naren Bharatwaj and Aodhagan Downey. "Real-Time Imbalance Pricing in I-SEM: Ireland's Balancing Market". In: *2018 20th National Power Systems Conference (NPSC)*. 2018, pp. 1–6. DOI: 10.1109/NPSC.2018.8771848.

[123] Elexon. *A Guide to Electricity Imbalance Pricing in Great Britain*. Accessed: 2025-03-17. 2020. URL: https://bscdocs.elexon.co.uk/guidance-notes/imbalance-pricing-guidance.

[124] IEA. *Covid-19 Impact on Electricity*. Accessed: 2025-03-18. 2021. URL: https://www.iea.org/reports/covid-19-impact-on-electricity.

[125] Friderike Kuik, Jakob Feveile Adolfsen, Eliza Magdalena Lis, and Aidan Meyler. *Energy Price Developments in and out of the COVID-19 Pandemic – From Commodity Prices to Consumer Prices*. Accessed: 2025-03-

18. 2022. URL: https://www.ecb.europa.eu/pub/economic-bulletin/articles/2022/html/ecb.ebart202204_01~7b32d31b29.en.html.

[126] Energy Exemplar. *Power Industry Solutions*. Accessed: 2025-03-21. 2025. URL: https://www.energyexemplar.com/industries/power.