Bachelor Thesis

Determinants and Informational Efficiency of Intraday Electricity Price Dynamics in Great Britain.

Authors:

Ivars Kraulis

Artis Kempelis

Supervisor:

Olegs Matvejevs

JEL codes: Q41, G14

March, 2026

Riga

COPYRIGHT DECLARATION AND LICENCE

Names of the authors in full: Artis Kempelis, Ivars Kraulis

Title of the Thesis: Determinants and Informational Efficiency of Intraday Electricity Price Dynamics in Great Britain.

We hereby certify that the above-named thesis is entirely the work of the persons named below, and that all materials, sources and data used in the thesis have been duly referenced. This thesis – in its entirety or in any part thereof – has never been submitted to any other degree commission or published.

In accordance with Section 1 of the Copyright Law of Latvia, the person(s) named below is (are) the authors of this thesis.

Pursuant to Article 40 of the Copyright Law the authors hereby agree and give an explicit licence to SSE Riga to deposit one digital copy of this thesis in the digital catalogue and data base at SSE Riga Library for an unlimited time and without royalty. The licence permits SSE Riga to grant access to the duly deposited thesis to all users of the catalogue and data base without royalty and limitations to downloading, copying and printing of the digital thesis in whole or in part provided we are indicated as the authors of the thesis according to Clause 4 Section 1 Article 14 of Copyright Law. We assert our right to be identified as the authors of this thesis whenever it is reproduced in full or in part.

Signed

______________________ ______________________

Artis Kempelis Ivars Kraulis

Date

……………………………………

List of Abbreviations

Abbr.	Full Term	Description
ACER	Agency for the Cooperation of Energy Regulators	EU agency coordinating national energy regulators and monitoring wholesale energy markets
API	Application Programming Interface	Software interface enabling automated data retrieval from platforms such as Elexon and NESO
AST	Augstsprieguma tīkls	Latvian transmission system operator
BM	Balancing Mechanism	GB mechanism through which NESO procures balancing services to maintain system frequency
BMRS	Balancing Mechanism Reporting Service	Elexon-operated platform publishing data on the GB balancing mechanism
BRELL	Belarus–Russia–Estonia–Latvia–Lithuania	Former synchronous electricity network connecting Baltic states to the Russian grid; desynchronised February 2025
BRP	Balance Responsible Party	Market participant financially accountable for deviations between scheduled and actual position per settlement period
BSC	Balancing and Settlement Code	Governance framework defining rules for electricity balancing and imbalance settlement in GB
BSP	Balancing Service Provider	Entity providing balancing energy or capacity to the TSO
CACM	Capacity Allocation and Congestion Management	EU regulation governing cross-zonal transmission capacity allocation and DA/ID coupling
CEER	Council of European Energy Regulators	Association of European national energy regulatory authorities
CfD	Contract for Difference	Long-term government-backed contract guaranteeing a fixed price for low-carbon electricity generation
CPC	Cumulative Price Change	Cumulative VWAP adjustment following an information event; measures speed of price incorporation
CRT	Cost of Round Trip	Liquidity measure capturing the implicit cost of a buy–sell transaction for the same contract
DA	Day-Ahead	Market timeframe referring to the auction held on D–1 for delivery on day D
DISEBSP	Derived Indicative System Energy Buy/Sell Price	IRIS data item containing indicative imbalance settlement prices published after each SP
DRM	De-Rated Margin	Difference between available capacity (adjusted for reliability) and forecast peak demand

Abbr.	Full Term	Description
DSO	Distribution System Operator	Entity operating the low- and medium-voltage electricity distribution network
EBGL	Electricity Balancing Guideline	EU Commission Regulation 2017/2195 establishing common rules for electricity balancing
ENTSO-E	European Network of Transmission System Operators for Electricity	Association of European TSOs coordinating cross-border transmission and operating the Transparency Platform
EPEX	European Power Exchange	Power exchange operating DA and ID electricity markets across multiple European countries, including GB
EU	European Union	Political and economic union of European member states
FBMC	Flow-Based Market Coupling	Cross-zonal capacity allocation method based on physical power flow modelling
GB	Great Britain	England, Scotland, and Wales; the geographic scope of the market analysed in this thesis
GW	Gigawatt	Unit of power equal to 1,000 MW
GWh	Gigawatt hour	Unit of energy equal to 1,000 MWh
ID	Intraday	Market timeframe covering continuous trading and auctions between the DA auction and gate closure
IDA1	Intraday Auction 1	First GB intraday auction, held at 17:30 on D–1
IDA2	Intraday Auction 2	Second GB intraday auction, held at 08:00 on day D
IRIS	Indicative and Reference Imbalance Settlement	Elexon system publishing indicative imbalance data shortly after each settlement period
ISP	Imbalance Settlement Period	Time interval over which imbalances are calculated and settled; 30 minutes in GB
LoLP	Loss of Load Probability	Probabilistic indicator of the risk that available generation will be insufficient to meet demand
MAPE	Mean Absolute Percentage Error	Forecast accuracy metric; problematic for electricity due to near-zero and negative prices
MARI	Manually Activated Reserves Initiative	European platform for the exchange of manually activated frequency restoration reserves
MASE	Mean Absolute Scaled Error	Forecast accuracy metric normalised by naive benchmark forecast error
mFRR	Manual Frequency Restoration Reserve	Balancing reserve activated manually by the TSO to restore system frequency
MRID	Master Resource Identifier	Unique identifier for a REMIT UMM, used to track revisions of the same outage event

Abbr.	Full Term	Description
MTU	Market Time Unit	Smallest tradeable delivery period; 30 min in GB (transitioning to 15 min for DA from October 2025)
MW	Megawatt	Unit of power; used to express generation capacity and outage magnitudes
MWh	Megawatt hour	Unit of energy; the standard trading unit in European electricity markets
NESO	National Energy System Operator	GB system operator responsible for real-time balancing, forecast publication, and system security
NIV	Net Imbalance Volume	Aggregate system imbalance per SP; positive = system short, negative = system long
OLS	Ordinary Least Squares	Standard linear regression estimation method
P305	BSC Modification P305	Reform implementing single cash-out pricing and marginal imbalance pricing in GB
PAR	Price Average Reference	Volume of most expensive balancing actions used to calculate cash-out price; 1 MWh under P305
REMA	Review of Electricity Market Arrangements	UK government programme reviewing GB electricity market design
REMIT	Regulation on Wholesale Energy Market Integrity and Transparency	EU/UK regulation prohibiting insider trading and market manipulation in wholesale energy markets
RSP	Reserve Scarcity Price	Administered price component added when LoLP indicates scarcity; calculated as LoLP $\times$ VoLL
SBP	System Buy Price	Imbalance price charged to short BRPs; equals SSP under single pricing (post-P305)
SDAC	Single Day-Ahead Coupling	Pan-European mechanism coupling national DA electricity auctions
SIDC	Single Intraday Coupling	Pan-European continuous cross-border ID trading mechanism; GB exited following Brexit
SP	Settlement Period	30-minute delivery window in GB, numbered 1–48 per day
SSP	System Sell Price	Imbalance price paid to long BRPs; equals SBP under single pricing (post-P305)
STOR	Short Term Operating Reserve	Contracted reserve service providing additional generation or demand reduction at short notice
TSO	Transmission System Operator	Entity operating the high-voltage transmission network and maintaining system balance
TWh	Terawatt hour	Unit of energy equal to 1,000 GWh
UMM	Urgent Market Message	REMIT disclosure reporting unplanned outages or changes in available capacity
UTC	Coordinated Universal Time	Time standard used for timestamping electricity market data

Abbr.	Full Term	Description
VoLL	Value of Lost Load	Administered estimate of cost of involuntary disconnection; £6,000/MWh in GB
VWAP	Volume-Weighted Average Price	Price measure: ratio of total traded value to total traded volume over a defined window
WMAE	Weighted Mean Absolute Error	Forecast accuracy metric normalising absolute errors by the mean price over the evaluation period
XBID	Cross-Border Intraday	Trading platform underlying SIDC for continuous cross-border ID trading in Europe

Abstract

This thesis examines the determinants and informational efficiency of intraday electricity prices in Great Britain using 70.2 million transactions from the EPEX continuous trading platform (2021-2025). Adapting event study methods from financial markets, we measure the price response to four categories of publicly available information: NESO wind and solar forecast revisions, derated margin and Loss of Load Probability updates, REMIT generation outage announcements, and indicative imbalance settlement prices and volumes. All four information types significantly affect intraday prices, but the public NESO wind forecast does not reach significance in the pooled sample, which we attribute to proprietary forecasts pre-empting the public signal. Across all variables, price sensitivity concentrates within the final hour before gate closure, confirming that the approaching imbalance settlement amplifies the value of information near delivery. Incorporation speed varies by signal type: imbalance prices are absorbed within seconds with evidence of overreaction, REMIT outages show persistent underreaction over 10 minutes, and solar revisions fall between. The market has become measurably more responsive over the sample period, coinciding with a tenfold increase in trading activity and consistent with growing algorithmic participation. We connect these findings to consumer welfare: following Koch and Hirth (2019), more responsive intraday trading reduces residual imbalances and the balancing costs recovered from consumers. However, the growing efficiency appears driven by well-resourced participants with proprietary data, raising concerns about participant structure. Our results validate GB's public data provision regime, establish empirical baselines for REMIT compliance monitoring, and demonstrate what new entrants need to know to participate in GB intraday trading.

Keywords: intraday electricity market, price discovery, event study, Great Britain, REMIT, imbalance price, renewable forecasts, market efficiency

JEL codes: Q41, G14

1. Introduction

Short-term electricity markets have become increasingly important as power systems rely more heavily on renewable generation. Technologies such as wind and solar reduce carbon emissions, but they also introduce uncertainty because their output depends on weather conditions that change during the day. As forecasts are updated and new information becomes available, market participants must continuously adjust their expectations. This creates large price swings in the hours leading up to delivery and increases the value of markets that operate closer to real time.

In Great Britain, these adjustments take place mainly in the intraday market. The day-ahead auction sets the initial schedule for the next day, but much of the new information that affects system balance arrives afterwards. Wind and solar forecasts are revised several times a day, demand estimates improve as delivery approaches, and generators can experience unexpected outages. Market participants respond to these changes by trading intraday contracts, which reduces their exposure to imbalance prices in the balancing mechanism. Because Great Britain settles imbalances in 30-minute periods and no longer participates in the European SIDC platform, intraday prices are shaped mostly by domestic conditions and short-run system signals.

Although many studies have analysed intraday price formation in continental Europe, the evidence for Great Britain is still limited, especially for the current regime after the EU exit. Most empirical work covers markets that participate in the Single Intraday Coupling (SIDC/XBID), where prices are jointly formed within a common European order book. These findings do not directly apply to Great Britain, which now trades intraday outside SIDC, relies on explicit cross-border allocation, and uses local intraday auctions together with 30-minute settlement periods. Despite these structural differences, no published study has systematically examined GB intraday prices using high-frequency transaction data for the post-2021 period.

The existing literature has also increasingly shifted its methodological focus. As Thakare et al. (2023) document, the dominant trend in intraday electricity research has been toward neural networks, gradient-boosted trees, and other black-box models that optimise predictive accuracy. While these models can achieve strong out-of-sample performance, they do not identify the marginal effect of a specific wind forecast revision, a generation outage, or an imbalance price publication on the intraday price, nor do they reveal whether the market response is immediate or delayed, proportional or asymmetric. For the GB market specifically, where no systematic empirical analysis exists for the post-2021 period, the priority is not yet to forecast prices but to first understand what drives them. The fundamentals we study (renewable forecast revisions, system tightness indicators, REMIT outage disclosures, and imbalance settlement publications) represent clearly identifiable information events with known timing and magnitude. This makes them well-suited for studying the causal relationship between information arrival and price adjustment, analogous to the analysis of earnings announcements or macroeconomic releases in traditional financial markets. This causal, event-driven perspective on GB intraday price formation remains largely absent from the literature. This motivates our research question:

What are the determinants of intraday electricity prices in Great Britain, and how efficiently do prices incorporate new fundamental information as delivery approaches?

Beyond the academic contribution, answering this question is directly relevant for market transparency, monitoring and policy. Understanding which fundamentals drive intraday price dynamics helps regulators identify where information asymmetries may arise and which data streams are most critical for efficient price formation. For the Review of Electricity Market Arrangements (REMA) monitoring teams, clearer knowledge of how prices should respond under normal conditions makes it easier to detect abnormal patterns or behaviour that may signal manipulation. For National ESO and other system operators, identifying which variables market participants rely on most heavily can inform investment decisions in forecasting systems, improve data publication standards and grid infrastructure. A market that reacts efficiently to the right signals supports better operational planning, reduces balancing costs and ultimately improves system reliability. Therefore, analysing how information is incorporated into GB intraday prices contributes not only to the literature but also to ongoing efforts to build a transparent, well-functioning electricity market.

Unlike most financial assets, where only holders of the asset bear exposure to price deviations, electricity price dynamics affect every household and business that consumes power. A shareholder in a publicly listed company bears the risk of that company's stock volatility; a non-shareholder does not. Electricity is fundamentally different. According to the European Commission's Household Budget Survey, housing, water, electricity, gas, and other fuels constituted the single largest category of EU household expenditure in 2020, accounting for 32.9% of total spending across 26 member states (Eurostat, 2022). The Odyssee-Mure project estimates that energy expenditure alone averaged 4% of household budgets across the EU in 2023, rising above 6% in Czechia, Estonia, and Latvia (Enerdata, 2024). This universal exposure distinguishes electricity from conventional financial assets: wholesale market efficiency is not merely a concern for trading firms but a direct determinant of the price paid by every end consumer.

Market efficiency, in turn, depends on the breadth of participation. When fewer participants compete, the remaining players face weaker competitive pressure and can extract rents. Batalla-Bejerano et al. (2022) analyse electricity prices and switching rates across 27 EU member states from 2000 to 2019 and find that retail competition is associated with lower consumer prices, with the effect of full liberalisation being larger than that of partial liberalisation. Ito and Reguant (2022) document the converse: studying U.S. electricity deregulation, they find that where market power persists, wholesale markups dominate efficiency gains, leading to higher consumer prices and lower consumer welfare. The implication is clear: competitive markets deliver lower prices only when participation is broad enough to discipline necessary behaviour.

Latvia illustrates this dynamic very well. The Latvian power market is dominated by Latvenergo, a state-owned utility that held a 51% retail market share in Latvia in 2024 (Latvenergo, 2025). As of 2020, Latvenergo accounted for 57% of national electricity generation, making Latvia one of only three EU member states, alongside Croatia and Lithuania, with a single major generating company (Eliņa, 2022). Latvian household electricity prices, which historically remained well below the EU average, have since risen to approximately 20% above it (Enerdata, 2024). Concentrated markets allow dominant players to set prices with limited competitive discipline, and the welfare cost falls directly on consumers through their electricity bills. One critical channel through which market efficiency translates into lower consumer costs is the relationship between intraday trading and system balancing. In European electricity markets, every generator, supplier, and large consumer is assigned to a balance responsible party (BRP), which is financially accountable for any deviation between its scheduled and actual position in each settlement period (Commission Regulation (EU) 2017/2195, Article 2(8)). When a BRP's portfolio is short or long at delivery, the transmission system operator must activate balancing reserves to close the gap, and the cost of this activation is reflected in the imbalance settlement price charged to the deviating party. Crucially, the system is designed to be financially neutral for the Transmission System Operator (TSO): any positive or negative financial outcome from the settlement process must be passed on to network users in accordance with national rules (Commission Regulation (EU) 2017/2195, Article 44(1)(i)). In practice, this means that the overall cost of procuring balancing services is transferred to end consumers partly through imbalance charges levied on BRPs, which suppliers pass through in retail tariffs, and partly through balancing capacity costs socialised directly in network tariffs (Commission Regulation (EU) 2017/2195, Article 44). The higher the residual system imbalance, the more reserves the TSO must activate, and the more expensive the balancing actions become, costs that ultimately appear on every consumer's electricity bill.

Active intraday trading reduces these costs by allowing market participants to correct their positions before delivery, thereby shrinking the residual imbalance that the TSO must resolve. The mechanism is straightforward: when updated forecasts reveal that the system will be shorter than the day-ahead schedule anticipated, informed participants buy additional power on the intraday market. This trading pulls flexible generation, such as gas-fired plants with short start-up times, into the market. Consider a settlement period where the day-ahead auction cleared at €20/MWh, but revised wind forecasts indicate lower-than-expected output. A gas plant with a marginal cost of €90/MWh offers into the intraday continuous market at €110/MWh, and a BRP that recognises it will be short purchases this power to close its position. The trade is expensive relative to the day-ahead price, but the alternative is worse: if the BRP does not act, the same megawatt of deficit reaches the balancing mechanism, where the TSO procures it from balancing reserves at the imbalance settlement price, which, during tight conditions, can reach multiples of the intraday price. The intraday trade thus serves as a pressure valve: it brings the marginal unit of generation online before delivery, reduces the volume the TSO must balance, and in doing so reduces the total cost passed on to consumers.

Koch and Hirth (2019) provide empirical evidence of this mechanism in Germany. Despite wind and solar generation nearly doubling between 2011 and 2017, both balancing reserve requirements and reserve activation volumes declined by approximately 50%, a phenomenon they term the "German Balancing Paradox." They attribute much of this improvement to increased intraday trading activity, finding that the shift to quarter-hourly trading products alone reduced balancing energy volumes by 17%, and that market participants respond efficiently to imbalance price signals when adjusting their portfolios. Eicke et al. (2021) formalise this relationship further, estimating that in Germany, the demand for imbalance energy declines by 2.2 MW for each €1/MWh increase in the imbalance price, which confirms that price signals in the balancing mechanism incentivise active portfolio management on the intraday market.

When intraday markets are thin or when participants lack the information needed to trade effectively, this self-correcting mechanism breaks down. The consequences can be extreme: during the summer of 2025, the Baltic balancing markets recorded imbalance settlement prices reaching +€9,976/MWh and -€10,007/MWh, which reflects a system in which residual imbalances were large and the marginal cost of resolving them was borne by a small number of balance responsible parties (Baltic Transparency Dashboard, 2026). These costs do not remain confined to the BRPs who incurred them. Under the EU's electricity balancing framework, BRPs that deviate from their schedules are charged the imbalance settlement price, and these charges are incorporated into the wholesale cost base that retail suppliers pass through to end consumers in their tariffs (Commission Regulation (EU) 2017/2195, Article 54-55). As the ACER-CEER (2024) Market Monitoring Report notes, without significant changes to demand patterns and market participation, prices will ultimately increase for all consumers regardless of contract type, as the system costs of managing imbalances and congestion must be recovered through network tariffs and retail pricing (para. 17). In less competitive markets dominated by incumbent suppliers with limited switching, as is the case in several Baltic and Central European member states where fixed-price contracts account for 100% of household contracts (ACER & CEER, 2024, Figure 1), the pass-through is more direct and consumers have fewer alternatives. Either way, the cost of an inefficient balancing market materialises on every consumer's electricity bill. By contrast, a well-functioning intraday market with broad participation and timely access to fundamental information allows participants to trade out their forecast errors before delivery, compressing the residual imbalance volume, reducing the frequency of extreme balancing actions, and lowering the total system cost that must ultimately be recovered from consumers.

However, entering electricity markets requires understanding what drives prices and when information matters. In continental European markets, much of this intelligence is gatekept behind proprietary data and forecasting services. Providers such as Volue Insight, Montel/Energy Quantified, Dexter Energy, Enapsys, and Meteorologica offer fundamental forecasts, intraday price predictions, and algorithmic trading tools as commercial subscriptions. While EU Regulation 543/2013 mandates that transmission system operators publish fundamental data freely through the ENTSO-E Transparency Platform, the raw TSO data alone is insufficient for trading: it requires substantial processing, cleaning, and modelling to translate into actionable signals. Even as of 2025, the Agency for the Cooperation of Energy Regulators continues to recommend improvements to the platform, noting the need for greater consistency in data definitions across legal frameworks and improved clarity in publications related to balancing and cross-zonal capacities (ACER, 2025). A concurrent study documents persistent data gaps and instances of implausible values on the platform, including generation figures exceeding installed capacity, alongside significant performance differences in day-ahead forecasts across bidding zones (Mascarenhas & Kazmi, 2025). The commercial data providers fill these gaps by layering proprietary weather models, machine-learning forecasts, and curated data feeds on top of the freely available but imperfect fundamentals. For a small utility, an independent renewable producer, or a new trading entrant, the combined cost of multiple data subscriptions, forecasting tools, and the technical infrastructure to process them represents a meaningful barrier to participation, particularly in smaller markets like the Baltics where trading revenues may not justify the expenditure.

Great Britain offers a structurally different transparency regime. National ESO and Elexon publish wind, solar, and demand forecasts, REMIT outage data, imbalance settlement prices and volumes, loss-of-load probabilities, and derated margins freely and at high frequency through open APIs. Unlike the fragmented ENTSO-E ecosystem, where data quality and completeness vary substantially across member states and where definitions remain inconsistent even after a decade of operation (ACER, 2025), the GB system provides a single, well-documented, internally consistent data ecosystem maintained by one TSO and one settlement body. Our research exploits this transparency directly. By identifying which of these publicly available fundamentals drive GB intraday prices, quantifying the magnitude of their effects in £/MWh, and measuring how quickly the market absorbs each type of information shock, we provide an empirically grounded framework that any market participant can replicate without purchasing proprietary data. In effect, this study demonstrates what a new entrant needs to know and what a TSO needs to publish in order to support an efficient, well-informed intraday market.

2. Literature Review

2.1 Overview of electricity markets

Electricity markets differ fundamentally from most commodity markets because electricity cannot be stored economically at scale. This non-storability implies that the generation and the consumption of electricity must be the same and be balanced in real time. This means that the grid requires continuous adjustments by system operators, and that creates price dynamics that are extremely volatile. (Weron, 2014). The physical requirement that generation equals consumption at every moment means that market participants must form expectations about supply and demand ahead of delivery and update these expectations as new information arrives.

Most liberalised systems organise trading into three sequential timeframes, each of which processes information that becomes available before delivery. The Day-Ahead (DA) market is a centralized auction that sets prices for the following day. It aggregates participants' expectations about demand, available generation, cross-border capacity, and expected renewable output. The Intraday (ID) market provides a continuous or auction-based venue for adjusting positions closer to real time. Participants use the ID market to respond to new information such as updated wind or solar forecasts, changes in expected demand, or unexpected generator outages. The Balancing Market (BM) is operated by the system operator to correct any remaining discrepancies between supply and demand at the moment of delivery. The balancing auctions taken in this stage determine the imbalance settlement price.

This sequence of markets reflects the way information becomes available. Day-ahead prices incorporate expectations formed roughly one day before delivery. Intraday prices adjust to new information that arrives within the same day, such as forecast revisions or operational surprises. Balancing prices capture the final stage of system adjustment when deviations must be corrected immediately.

For Great Britain, this sequencing is central to understanding intraday price formation. Recent years have brought rapid growth in wind and solar generation, new interconnectors, and the introduction of local intraday auctions. These developments have made the GB system more sensitive to short-term forecast accuracy and have increased the importance of flexible trading opportunities closer to real time.

2.2 Day-ahead Market

The Day-Ahead (DA) market is the central forward-trading stage in Great Britain and determines scheduled generation and reference prices for delivery on the following day. Before the end of the Brexit transition period, GB was part of the European Single Day-Ahead Coupling (SDAC). This meant that day-ahead prices were jointly cleared with continental Europe through the Euphemia algorithm. Euphemia computes a single welfare-maximising solution for all participating zones and simultaneously determines market clearing prices for each bidding zone and the optimal cross-zonal electricity flows that respect interconnector capacity limits (ENTSO-E, n.d.; European Commission, 2015). SDAC uses implicit allocation. Under implicit allocation, traders submit only buy or sell orders for electricity and the algorithm automatically assigns transmission capacity to arbitrage price differences across borders. According to the CACM Regulation, implicit allocation is designed to maximise social welfare because it ensures that lower-cost generation meets demand across coupled zones and reduces the need for traders to predict cross-border congestion (European Commission, 2015).

Since 1 January 2021, Great Britain is no longer part of SDAC and Euphemia no longer includes GB in the European clearing process. The UK government explains that after EU Exit, GB electricity is no longer traded through the EU market coupling regime and that interconnector capacity between GB and France, Belgium and the Netherlands is allocated through explicit capacity auctions rather than jointly with the energy market (Department for Business, Energy & Industrial Strategy, 2021). Under explicit allocation, energy and transmission rights are purchased separately. Traders must first buy interconnector capacity in a dedicated auction before they can import or export electricity. The post-Brexit analysis of EU and GB transmission system operators notes that explicit allocation leads to less efficient use of transmission lines because capacity decisions are taken without the knowledge of final electricity price outcomes and that this tends to reduce cross-border price convergence (CEPA, 2021).

In the post-SDAC environment, GB day-ahead trading is conducted through multiple independent auctions operated by N2EX (Nord Pool) and EPEX SPOT UK. Nord Pool's main N2EX auction trades hourly products and closes at 09:50 GMT (Nord Pool, 2022). EPEX SPOT operates a parallel 60-minute GB Day-Ahead auction with gate closure at 09:20 GMT, again independent of any continental coupling mechanism (EPEX SPOT SE, n.d.-a). Because these auctions clear separately and do not rely on a shared capacity allocation mechanism, they can produce different prices for the same delivery hours.

Beyond these hourly auctions, GB also has next-day auctions in half-hour products. This reflects its 30-minute imbalance settlement period. EPEX SPOT runs a GB 30-minute Day-Ahead auction at 15:30 GMT, allowing participants to trade all 48 half-hour settlement periods of the following day with more granularity (EPEX SPOT SE, n.d.-b). Nord Pool offers a comparable GB Half-Hour Auction with an identical gate time of 15:30 GMT (Nord Pool, 2022). These half-hour auctions provide additional opportunities for participants to refine their positions with better granularity than the hourly auctions allow.

The absence of SDAC coupling means that none of these auctions incorporate cross-border flows directly in their clearing processes. Instead, the auction outcomes reflect domestic supply and demand conditions, weather expectations, fuel prices and bilateral trading positions, while interconnector flows depend on the outcome of separate explicit capacity auctions. This separation has made GB day-ahead prices more responsive to local fundamentals and more insulated from continental supply shocks. At the same time, it has reduced the automatic arbitrage that previously aligned British and European prices. Today, therefore, GB differs from most SDAC-coupled markets in two key ways: (i) its day-ahead prices are set in national auctions that are not coupled with continental Europe, and (ii) cross-border capacity on GB-EU interconnectors is allocated explicitly at day-ahead, while implicit allocation remains the reference model within the EU (European Commission, 2015; ENTSO-E, 2023). This background is important for our research, since GB intraday prices evolve on top of a day-ahead structure that is not jointly optimised with continental markets and interconnector positions are determined through separate capacity auctions rather than by Euphemia within SDAC. This affects how cross-border fundamentals and forecast revisions can feed into GB intraday price dynamics.

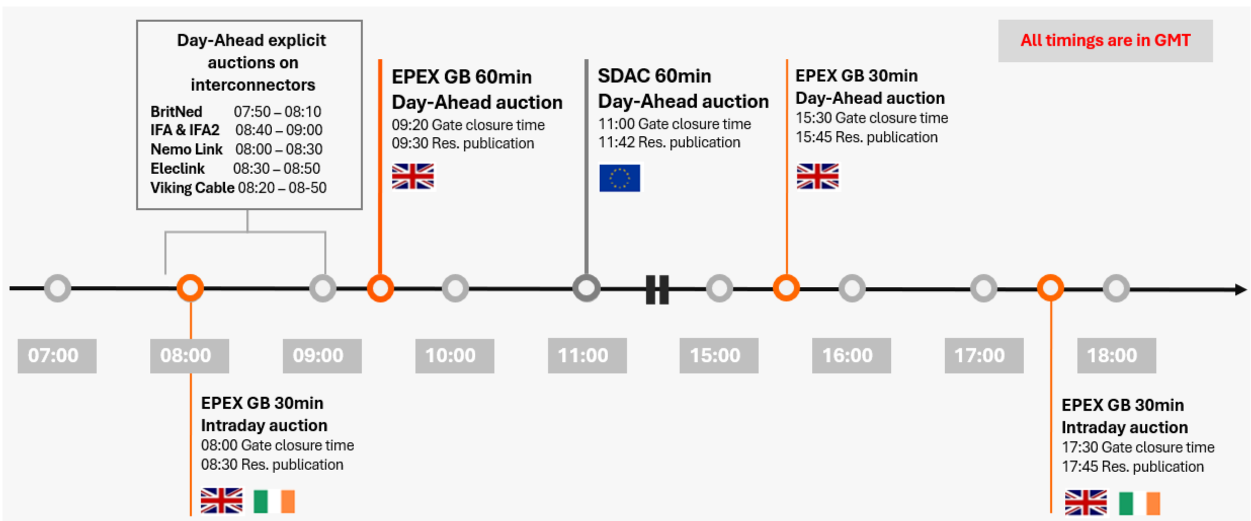

Figure 1. A time-series breakdown of the Day-Ahead and Intraday, and interconnector auctions. Source: EPEX SPOT SE. (n.d.-a)

2.3 Balancing Market

The Balancing Market (BM) is the last stage of the short-term electricity market in Great Britain and is operated by National Grid ESO to maintain real-time physical balance. Whereas the Day-Ahead and Intraday markets allow participants to adjust positions ahead of delivery, the BM resolves any remaining mismatches between contracted positions and actual system conditions. Imbalance settlement in Great Britain is carried out in 30-minute settlement periods, and each half-hour is settled separately. Parties that are long or short relative to their contracted volume in a settlement period are exposed to the Energy Imbalance Price, which is calculated according to the Balancing and Settlement Code (BSC) (Elexon, 2025).

Usually, there is a frequency restoration reserve merit order list, where market participants can submit bids, e.g., an asset is willing to turn on, for example, 14 MW of balancing reserves for a price of 98 Euros per MWh. Many of these assets are then added to a merit-order list, and the imbalance price is usually a function of the marginal/weighted activated balancing bids. This methodology is applied in much of continental Europe, including France, Germany, Belgium, and the Netherlands. However, countries like Belgium can rely on this because energy flows freely and locational constraints are smaller. In Great Britain, which is an island with tight constraints and limited interconnector capacity, the Balancing Mechanism is not merit-order based. ESO dispatches balancing actions according to technical constraints rather than purely on price, and the imbalance price is produced by a post-processing methodology that filters, flags and adjusts actions through a multi-stage algorithm. As a result, GB does not maintain or publish a balancing reserve merit-order list, and the imbalance price does not correspond to the marginal cost of the activated reserve.

2.4 Intraday Market

The intraday market is the second stage of short-term electricity trading in Great Britain. Its purpose is to let participants correct their positions once the day-ahead schedules are set and new information becomes available. Forecasts for wind, solar and demand improve throughout the day, and unexpected outages or system changes often appear with little warning. Because of this, the intraday market has become an essential part of daily trading, especially in a system that relies heavily on renewable generation.

Trading takes place in two main forms: continuous trading and intraday auctions. Continuous trading is offered by both N2EX (Nord Pool) and EPEX SPOT UK. In these markets, participants can buy or sell power for individual 30-minute periods up until shortly before delivery. Prices move quickly as traders react to real-time updates from Elexon, the system operator and interconnectors.

Great Britain also has two local intraday auctions run by EPEX SPOT. The first auction, IDA1, is held at 17:30 GMT on the day before delivery and usually incorporates the latest day-ahead forecast updates. The second, IDA2, takes place at 08:00 GMT on the delivery day and captures the morning revisions to wind, solar and demand. These auctions create short but concentrated windows of liquidity and often reset expectations before the day progresses into continuous trading.

Intraday products in GB trade in 30-minute blocks, which matches the imbalance settlement period used in the balancing mechanism. This link is important, as any position a trader holds after the intraday market closes becomes an imbalance position, so intraday trading is widely used to manage risk and avoid exposure to potentially unfavourable imbalance prices later in the day.

Since leaving the EU, Great Britain no longer participates in the Single Intraday Coupling (SIDC). This means that GB traders do not have access to the common European order book (XBID) and must rely on domestic liquidity. Cross-border capacity is obtained through explicit auctions rather than automatic matching with continental bids and offers. As a result, intraday prices in GB depend mainly on internal conditions such as demand, renewable output, outages and system tightness. Interconnector flows still matter, but mostly when the actual physical flows differ from earlier nominations, since such deviations signal imbalance pressure.

In practice, the intraday market sits between the day-ahead auction and the balancing mechanism. It gives participants the tools to adjust their positions ahead of real-time balancing and to manage uncertainty that arises during the delivery day. These features make the intraday market a key part of how prices evolve in Great Britain, especially in the period after EU Exit.

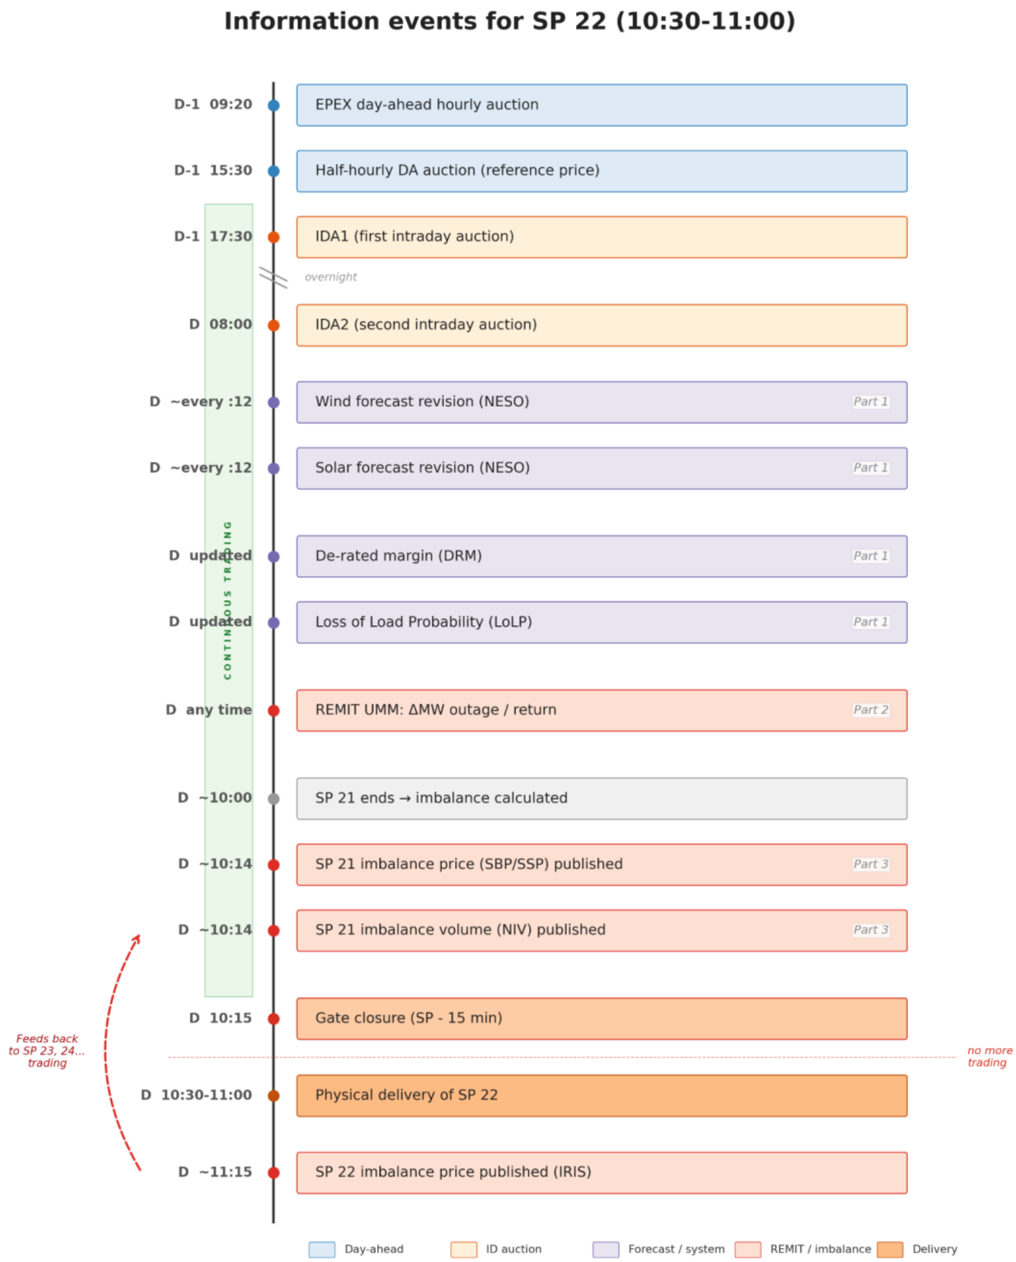

Figure 2. Hypothetical chronological sequence of information events affecting the intraday price of Settlement Period 22 (10:30-11:00).

2.5 Intraday Price Formation

2.5.1 Price Measurement and Methodological Considerations

As discussed previously, the intraday market is used to correct asset production or consumption deviations from forecasts, i.e., from what was submitted in the day-ahead auction. Besides wind or solar deviations, the need to balance in the intraday market can arise due to technical failures in both traditional hard generation plants (coal, gas, nuclear, etc.) as well as renewables. Similarly, the transmission network itself often experiences planned maintenance on larger or smaller scales, as well as technical failures. An intraday price that exceeds the day-ahead price in a specific Market Time Unit (MTU) could mean that for assets that were above the day-ahead price on the merit-order curve, it is now profitable to sell (or buy in the case of BESS) at this level. These and many other factors can theoretically be drivers of the intraday price.

It is important to mention what is meant by the intraday price. Due to lower liquidity when compared to most financial markets, the intraday price is often computed as the volume-weighted average price (VWAP). This can be computed over different time periods until gate closure; for example, 1H VWAP would imply the average price of all transactions, weighted by volume, that took place in the last hour before gate closure. While this is more informative than the open or close prices for a given period, as prices can swing drastically due to large one-off transactions or new market information, many studies just look at the evolution of this VWAP over long timeframes. Other studies often examine the day-ahead and intraday price spread, as electricity has no simple intrinsic price and day-ahead prices can differ drastically across contracts. This helps to normalise the data; however, the nominal deviation is still economically relevant.

2.5.2 Renewable forecasts

Vast empirical evidence is available for the German electricity market. The review by Thakare et al. (2023) on intraday price forecasting studies found that almost all models incorporate wind and solar forecast shocks. Privately owned renewable assets are a main reason that modern short-term electricity markets exist so their impact on the price is quite large. For the German market, Ziel (2017) found forecast errors of wind and solar to be the dominant determinants of intraday price deviations from the day-ahead price. Furthermore, due to this importance, the accuracy of renewable forecasts plays a key role in overall price volatility, as found by Gurtler and Paulsen (2018). In Great Britain, similarly to Germany, wind and solar forecast revisions are published by Elexon and the system operator, which aggregate individual BRPs' forecasts into system-level values. In Germany, the effect of these revisions increases closer to delivery as they form a larger surprise (Kiesel & Paraschiv, 2017). This suggests that the impact of fundamental shocks may differ across contracts depending on time-to-delivery, which we also investigate in the GB setting.

While solar and wind forecast errors can be thought of as causing imbalances and hence needing to be balanced out within the intraday, the balancing market represents the opportunity cost of balancing in the intraday. With a tighter balancing margin, the risk of a higher imbalance price is much larger so the opportunity cost of balancing in the intraday market rises, in turn increasing the intraday price.

2.5.3 System Tightness and Scarcity Pricing

Empirical studies using system tightness or other market-specific proxies have identified it as a relevant determinant of intraday price (Hagemann, 2015). As the full order book of the balancing market cannot be observed by anyone before delivery, often different proxies of balancing capacity are used to estimate the margin. For Great Britain, Elexon publishes a derated margin (DM) measure that can function as a proxy for short-term system tightness. Intuitively, DM approximates the amount of "spare" reliable generation capacity available after accounting for demand and the fact that not all generation is fully dependable at peak times, for example, due to outages or derating. When the margin is small, the system is tight and the risk of not meeting demand increases.

To translate tightness into an economic scarcity signal, the system operator computes a Loss of Load Probability (LoLP) for each settlement period. LoLP is the probability that available generation will be insufficient to meet demand, meaning the probability that some demand would have to be curtailed if the system becomes short. In parallel, the Value of Lost Load (VoLL) represents the economic cost to consumers of involuntary disconnection, in other words, how costly it is, in welfare terms, to not serve electricity demand. In GB, LoLP and VoLL are combined into the Reserve Scarcity Price (RSP) via:

$$RSP_t = LoLP_t \times VoLL$$

Crucially, RSP does not affect the imbalance price as a small proportional adjustment. Instead, it can enter the imbalance calculation discretely. Eligible balancing actions can be assessed against the scarcity value, and if an action's utilisation price is below RSP, it can be repriced up to the RSP level before the system price is formed. A key example is Short Term Operating Reserve (STOR) actions. STOR is a contracted reserve service where providers agree to make additional generation available, or reduce demand, at short notice when instructed by the system operator. The “certain STOR actions” in this context refer to STOR energy actions that are identified as STOR and delivered within the provider’s STOR availability window, meaning they are treated as reserve deployment rather than ordinary balancing offers. When scarcity is high, repricing these STOR actions up to the RSP level ensures that cash out prices reflect not only the utilisation costs of balancing actions but also the scarcity value of reserves.

Economically, this matters for intraday price formation close to delivery because the expected imbalance price is the outside option for any participant that remains imbalanced. If a trader does not close their position in the intraday market, the position is ultimately settled at the imbalance price. Therefore, changes in DM and LoLP, and hence RSP, shift the expected cost of being out of balance, which should be priced into intraday contracts approaching delivery through traders’ incentives to adjust positions when scarcity risk increases.

2.5.4 Unplanned Outages and REMIT Information Disclosure

Unplanned outages have an impact on the intraday price as they push up the risk of a larger imbalance and hence a larger balancing price (Hagemann, 2015). In GB, these are collected under REMIT messages. Under UK REMIT, market participants must publicly disclose inside information (Article 4). "Inside information" is defined as information that is (i) precise, (ii) not publicly available, (iii) relates (directly or indirectly) to wholesale energy products, and (iv) would be likely to significantly affect prices if made public (Regulation (EU) No. 1227/2011, 2011, Art. 2(1)). Ofgem's enforcement approach assesses conduct against this qualitative definition, meaning whether information qualifies depends on market conditions and the likelihood of significant price effects, rather than any universal fixed MW threshold (Ofgem, 2014a, p.3).

An unplanned generation outage reduces available capacity in specific settlement periods. This increases the expected scarcity of balancing actions and can raise imbalance prices. Forward-looking market participants may then bid intraday prices upward for the affected half-hours to avoid being short and exposed to high expected imbalance prices. Because UMMs have precise publication timestamps and period-specific asset impacts, they provide well-timed information events for an intraday event study, while still requiring standard checks for pre-trends and potential information leakage.

The empirical literature on the price impact of outage disclosures in European electricity markets is limited but growing. Hagemann (2015) identifies unplanned outages as a significant determinant of German intraday prices using hourly VWAP data. Valitov and Maier (2020) provide the closest precedent to our study, examining how private and public information about unplanned outages affects the day-ahead to intraday price spread on the German EPEX market. They find that both private and public information about unplanned outages have a significant positive effect on intraday electricity prices in Germany. Notably, they document evidence of asymmetric information effects, where prices respond to private information before public disclosure, suggesting potential information leakage. Both studies work with relatively coarse price measures: Hagemann uses hourly VWAP indices, while Valitov and Maier examine the day-ahead to intraday spread rather than high-frequency intraday price changes around the exact publication timestamp. No published study has examined the short-run price response to REMIT publications in the GB market at the second-level frequency that our transaction data permits.

2.5.5 Imbalance Price Signals

The imbalance settlement price published by Elexon after each settlement period provides a high-frequency information signal that directly affects intraday trading incentives. Because GB settles imbalances every 30 minutes, a new indicative system price appears approximately every half hour, each representing the realised cost of balancing the system during the preceding period. For any participant still holding an open position in the intraday market, the imbalance price defines the opportunity cost of inaction: failing to close a position before gate closure means settling at whatever imbalance price materialises. When the published imbalance price for the previous settlement period is unexpectedly high, it signals that the system was tighter than anticipated, which should increase the willingness to pay for intraday contracts covering the next settlement period. Conversely, a low imbalance price suggests the system was long, reducing urgency to trade.

Eicke et al. (2021) provide empirical evidence of this mechanism in Germany, estimating cross-market equilibria between intraday and imbalance markets and finding that a shock to the imbalance price triggers a subsequent adjustment of the intraday price. They argue that each publication prompts position rebalancing, as traders update their expectation of the next imbalance price and trade intraday to avoid settlement at an unfavourable level. The net imbalance volume, published alongside the price, provides a complementary signal: large positive volumes indicate aggregate system shortage, while large negative volumes indicate surplus. Both variables represent clearly timed information events with known publication timestamps, making them well-suited for the event study framework we apply in Part 3 of our empirical analysis.

The relationship between intraday prices and the imbalance settlement has also been studied from a theoretical and strategic perspective. Aïd, Cosso, and Pham (2022) formulate an equilibrium model of intraday trading in which agents minimise the sum of trading costs and imbalance penalties. In their framework, the imbalance penalty directly shapes each agent's optimal trading rate: an agent whose forecasted marginal cost exceeds the market price sells, and vice versa, with the intensity of trading increasing as delivery approaches and the imbalance penalty becomes more salient. The model demonstrates that heterogeneity across agents is a necessary condition for the Samuelson effect (increasing price volatility near delivery) to hold, and that the equilibrium price dynamics are consistent with the empirical patterns observed on EPEX. While the model does not directly predict how the published imbalance price feeds back into intraday trading, it formalises the mechanism through which the imbalance penalty structures the incentives that drive intraday price formation.

Koch (2021) takes this relationship to its practical conclusion, analysing a trading strategy in Germany that explicitly exploits the spread between intraday prices and expected imbalance prices. Using a logistic regression model to predict the direction of the overall system balance, Koch demonstrates that a profitable strategy can be constructed by taking positions in the intraday market based on imbalance price expectations. The study documents that the average spread between intraday and imbalance prices is asymmetric depending on the direction of the system balance, and that biased imbalance price incentives can lead to positions that are not system-supportive. This work highlights that the intraday-imbalance price spread is not merely a residual but contains exploitable information, and that the design of the imbalance pricing mechanism has direct consequences for trading behaviour and system stability. For our study, Koch's findings motivate the possibility that the difference between a contract's intraday VWAP and the subsequently realised imbalance price could serve as an alternative measure of information surprise, a specification we discuss in Section 5.1.6 but do not pursue in the present analysis.

2.6 Event Studies in Traditional Financial Markets

Although this study is motivated by electricity-market dynamics, event-study methodologies originate in traditional financial markets and provide a useful benchmark for analysing price responses to discrete information shocks. A large literature examines how asset prices, volatility, and liquidity respond to earnings announcements and macroeconomic releases using high-frequency data. A key similarity with intraday electricity markets is that many such events occur during periods of reduced liquidity, particularly in after-hours trading, where transaction prices are sparse, and price discovery may occur through quotes rather than trades (Grégoire & Martineau, 2022).

Empirical evidence from financial markets consistently shows that prices adjust rapidly to new public information. Smales (2013), for example, documents that interest rate futures incorporate macroeconomic news within seconds, with sharp but short-lived increases in volatility, trading activity, and bid-ask spreads around announcement time. While the frequency of these announcements is far lower than the frequency of fundamental data updates, the same conceptual challenges arise. In particular, isolating the announcement signal from background price noise is a common concern across event-study designs. With even lower and more uneven liquidity than traditional after-hours financial markets, this noise can pose additional challenges for the analysis of GB intraday electricity prices, as individual trades may reflect transient imbalances rather than information-driven price discovery.

2.7 Research Gap

While these studies empirically have highlighted the main drivers, there are multiple caveats. Firstly, to our knowledge, no study has been conducted on UK intraday price formation using the recent period and the combination of fundamentals described above. While Germany has been extensively studied and many other regions like the Nordics have been examined, none are easily comparable. Availability of data differs in various ways, including live data, aggregated data and other metrics that are collected and shared differently in each country. Secondly, GB is not in the SIDC zone anymore, so the intraday prices observable there are primarily driven by internal conditions and bilateral arrangements, rather than a fully coupled European order book. This is quite different from Europe, where large countries like Germany can affect the price elsewhere significantly. Lastly, the focus in the forecasting literature has mainly shifted away from econometric models to more machine learning “black box” models (Thakare et al., 2023). While these models help to improve forecast accuracy, they make it harder to interpret the marginal effect of individual fundamentals on prices.

In this thesis, we therefore focus on relatively transparent econometric models that link changes in intraday prices to a clearly defined set of fundamentals: changes in renewable forecasts, derated margin, and outage events. This allows us to build on the existing empirical literature while filling the gap for Great Britain and keeping the results interpretable. Based on this, our first research question asks how these medium-term and short-term fundamentals jointly explain intraday price movements in GB, and how their effects differ across contracts as delivery approaches.

At the same time, much less is known about how intraday prices behave in the immediate vicinity of large shocks and about the short-run volatility response, especially at higher frequency and in the GB context. High-frequency intraday data have only recently become widely accessible, and many studies either aggregate to relatively coarse time intervals or focus mainly on point forecasting performance. As a result, there is limited evidence on how quickly intraday prices adjust after forecast surprises, changes in system tightness or large deviations in load and interconnector flows, and on whether these periods are systematically associated with higher short-term volatility.

3. Methodology

3.1 Data

3.1.1 Trade Data

As our main data resource, we will be using historic transaction data on the GB EPEX market. The data is collected on a per-transaction basis, which allows to construct a high-frequency analysis. The data is available spanning back to 2015, but for our purposes of studying the GB intraday as a market now disconnected from SIDC, the window of study will be limited from 01.01.2021 till 30.09.2025, when the 15-minute MTU was adopted for the day-ahead auctions. This range should provide a vast sample that would allow us to separate the data across certain half-hour contracts or look at changes over time if the data suggests such patterns.

The dataset includes all intraday transactions executed on the continuous trading platform. Each transaction record contains the contract delivery period (MTU), execution timestamp (to the 1/1000th of a second), traded volume (MWh), and price (£/MWh). Transaction data also allows us to choose our aggregation window, as opposed to older studies such as Hagemann (2015) that have looked at broad one-hour VWAP windows. Our granularity allows us to construct volume-weighted average prices (VWAP) at varying frequencies and to analyse price dynamics at high resolution. For running an event-study based methodology, we aggregate the pre and post VWAP values around time t_i for event *i*:

$$\mathrm{VWAP}^{\mathrm{pre}}_{c,t_i} = \frac{\sum\limits_{t\in[t_i-\tau,\, t_i)} P_{c,t}\cdot V_{c,t}}{\sum\limits_{t\in[t_i-\tau,\, t_i)} V_{c,t}}, \qquad \mathrm{VWAP}^{\mathrm{post}}_{c,t_i} = \frac{\sum\limits_{t\in[t_i,\, t_i+\tau)} P_{c,t}\cdot V_{c,t}}{\sum\limits_{t\in[t_i,\, t_i+\tau)} V_{c,t}}$$

where $P_{c,t}$ is the transaction price, $V_{c,t}$ is the volume, and the sum runs over all transactions which are executed around the period τ for event time t_i . This definition of VWAP is used in all subsequent sections.

3.1.2 Forecast Revision Data

Wind and solar generation forecasts are obtained from the NESO Data Portal, which publishes updated forecasts for each settlement period once every hour. We extract the historical forecast revisions for wind and solar generation for each settlement period. The structure of the data is that NESO publishes their solar and wind forecast on the 12th minute of every hour for every single settlement period. For example, for the settlement period 12:30-13:00, there would be several significant changes in forecasts at 12:12; 11:12;10:12 etc., with the final forecast coming in 18 minutes prior to the settlement period starting. For settlement periods starting at the turn of the hour, for example, at 14:00, the changes in forecasts would be available at 13:12; 12:12;11:12, etc., with the last forecast being available 48 minutes prior to the settlement period starting.

Each forecast revision is matched directly to the delivery contract it references, as the publication explicitly identifies the settlement period it covers. As these forecasts are updated each hour and their historical values are stored, we can match how these forecasts have evolved for each settlement period which corresponds to the respective contract.

If NESO's renewable forecasts significantly affect intraday prices, which they should, since they directly influence imbalance size and drive participants to rebalance, we would expect :00 settlement periods to be less price-efficient than :30 periods. This is because :00 periods have a 48-minute forecast-to-delivery gap versus just 18 minutes for :30 periods, meaning participants trade on less accurate information, potentially causing larger price deviations and slower adjustment to new information.

3.1.3 System Tightness and Pricing Data

The imbalance price and net imbalance volume for each settlement period are sourced from Elexon's IRIS (Indicative and Reference Imbalance Settlement) archive. IRIS publishes indicative system prices shortly after each settlement period ends, before the final D+1 settlement run. We obtain the raw IRIS DISEBSP (Derived Indicative System Energy Buy/Sell Price) files from the Elexon archive, which contain the indicative system buy price (SBP), system sell price (SSP), and net imbalance volume (NIV) for each settlement period, together with a field that records the exact UTC timestamp at which each publication was made available. This timestamp is critical for our event study design: it allows us to identify precisely when the market learned the imbalance outcome, rather than relying on assumptions about publication timing.

From the archive, we extract both the initial indicative publication and any subsequent revisions. For our purposes, only the first publication matters, as it represents the information shock to the market. Across our sample, the median lag between the end of a settlement period and the initial IRIS publication is approximately 15 minutes. For example, SP 21 (10:00-10:30) ends at 10:30, and its indicative imbalance price typically appears on IRIS around 10:45. This creates a clearly identifiable information event: intraday traders learn ex post how tight the system actually was during the previous settlement period, and can update their expectations for upcoming periods accordingly. Since each imbalance publication refers to an already-settled period, we match it to the next tradeable contract - the earliest half-hour contract whose gate closure has not yet passed at the publication timestamp. For example, if the SP 21 imbalance price is published at 10:45, the matched contract is SP 23 (11:00–11:30), whose gate closure at 10:45 has not yet elapsed.

As discussed in Section 2.5, the imbalance price serves as the opportunity cost of intraday trading. A participant who fails to balance their position in the intraday market faces the imbalance price in settlement. When the published imbalance price for the previous period is unexpectedly high, it signals that the system was tighter than anticipated, which should increase the willingness to pay for intraday contracts covering the next settlement period. Conversely, a low imbalance price suggests the system was long, reducing urgency to trade. The net imbalance volume (NIV) provides a complementary signal: positive NIV indicates aggregate system shortage, while negative NIV indicates surplus. Both variables enter our Part 3 specification as information shocks with known publication timing, allowing us to measure the causal price response in intraday trading for subsequent settlement periods.

We additionally collect two forward-looking system tightness indicators from Elexon's Balancing Mechanism Reporting Service (BMRS): the Loss of Load Probability (LoLP) forecast and the derated margin (DM) forecast. These are used by National ESO to compute the Reserve Scarcity Price and are published as rolling forecasts that update 4 times throughout the day, e.g., 8h, 4h, 2h and 1h before the respective SP. Hence, the matching contracts are already provided in the dataset for all historical revisions.

Unlike the imbalance price, which reflects realised system conditions ex post, LoLP and DRM represent the system operator's forward-looking assessment of supply adequacy for upcoming settlement periods. A rising LoLP or falling DRM signals expected tightness, which should be reflected in higher intraday prices for the affected contracts. These variables enter our Part 1 specification alongside the renewable and demand forecast revisions.

3.1.4 REMIT Urgent Market Messages

We collect REMIT (Regulation on Wholesale Energy Market Integrity and Transparency) Urgent Market Messages from the BMRS platform. Each REMIT message contains the following key fields: publication timestamp (when the information became public), fuel type (e.g., Fossil Gas, Wind Offshore), available and unavailable capacity (MW), event start and end times (which settlement periods are affected), and asset identifiers (participant ID, asset ID, unit name). To determine which settlement periods are impacted by each message, we expand the event start and end times into individual half-hourly settlement periods that fall within the outage window, restricted to those still tradeable when the message was published. This allows us to assign the MW capacity impact to specific half-hourly contracts.

However, REMIT messages can be updated multiple times for the same event as conditions evolve. To avoid double-counting capacity impacts, we group messages by their unique event identifier (MRID). We then sort by revision number and calculate the change in unavailable capacity relative to the previous revision. This captures only the new information revealed with each update, rather than the cumulative outage size. We then transform the dataset, so that we have a specific ΔMW variable for every event-settlement period pair. This enables us to transform the incoming dataset into a specific MW change on the total system imbalance volume. This is going to have an effect on the imbalance price. Since the imbalance price represents the opportunity cost of not closing one's intraday position before gate closure, a supply shortfall that increases the expected imbalance should have a direct impact on the intraday price for the affected contracts.

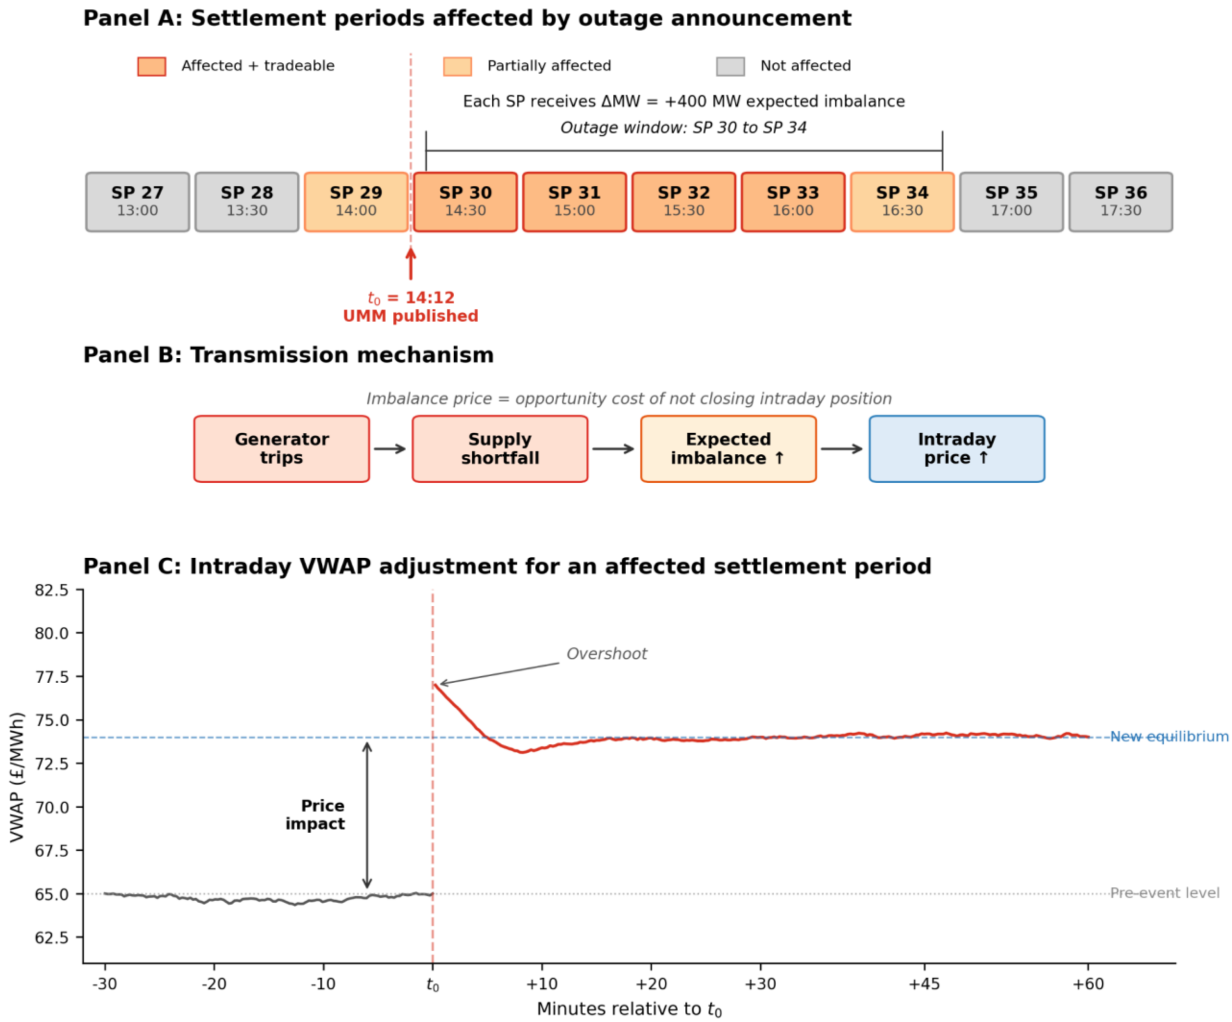

Figure 3. Hypothetical example of intraday price adjustment following a REMIT outage announcement.

3.2 Methodological framework

3.2.1 Dependent Variable

The choice of how to express price movements is not trivial in electricity markets and requires careful consideration. In traditional financial event studies, the standard approach is to work with logarithmic returns:

$$r_{c,i} = \ln\left(\frac{\mathrm{VWAP}^{\mathrm{post}}_{c,t_i}}{\mathrm{VWAP}^{\mathrm{pre}}_{c,t_i}}\right)$$

which are symmetric, time-additive, and approximately normally distributed for most asset classes. However, this convention breaks down fundamentally in electricity markets. GB intraday prices regularly take zero or negative values during periods of high renewable output or low demand. In our sample, approximately 3-4% of settlement periods exhibit negative VWAP levels at some point during the trading window. Since the logarithm is undefined for non-positive arguments, log returns cannot be computed for these observations. Excluding them would systematically remove precisely the low-price, high-renewable periods that are among the most economically interesting for studying forecast revision impacts.

A natural alternative is percentage returns, defined as:

$$r_{c,i} = \frac{\mathrm{VWAP}^{\mathrm{post}}_{c,t_i} - \mathrm{VWAP}^{\mathrm{pre}}_{c,t_i}}{\left|\mathrm{VWAP}^{\mathrm{pre}}_{c,t_i}\right|}$$

This resolves the domain problem by using the absolute value of the base price. However, it introduces a different issue: when the base price is close to zero, the denominator shrinks toward zero and percentage returns become arbitrarily large, regardless of the actual magnitude of the price movement. A price moving from £0.50 to £5.50 registers as a 1000% return, while a move from £50 to £55 registers as 10%, even though both represent the same £5/MWh shift. In our dataset, VWAP values between -£2 and £2 are not uncommon during overnight and early morning hours with high wind output. Any percentage-based measure would assign extreme weight to these observations, distorting the regression estimates.

We considered several approaches to mitigate this. One option was to winsorise or truncate observations where the absolute value of VWAP falls below a threshold, for example, 2 GBP which would eliminate the near-zero explosion at the cost of discarding observations. An alternative would be to normalise by a rolling average VWAP over a longer lookback window rather than the instantaneous pre-event price, which could smooth through transient near-zero prices. However, both approaches introduce arbitrary cutoff choices that could bias the sample toward higher-price periods and away from the renewable-dominated hours where forecast revisions are most relevant.

Weron (2014, Section 3.3), in his comprehensive review of electricity price forecasting methods, documents precisely this problem in the context of forecast evaluation metrics. He notes that the mean absolute percentage error (MAPE), the most widely used percentage-based accuracy measure, becomes misleading for electricity prices: it produces very large values when prices are near zero regardless of the actual forecast error, it yields artificially small values during price spikes regardless of the absolute miss, and it becomes negative and uninterpretable when prices are negative. He recommends normalising by the mean price over the evaluation period (the weighted mean absolute error, WMAE) or by the error of a naive forecast (the mean absolute scaled error, MASE), rather than by individual price observations. While his discussion concerns forecast evaluation rather than regression specification, the underlying statistical problem is identical: dividing by individual electricity prices produces unstable and misleading ratios.

Given these considerations, we follow the established approach in the intraday electricity literature and specify our dependent variable as the absolute price change in £/MWh: We can then define variables as nominal price changes:

$$\Delta P_{c,i} = \mathrm{VWAP}^{\mathrm{post}}_{c,t_i} - \mathrm{VWAP}^{\mathrm{pre}}_{c,t_i}$$

This is the specification adopted by Kiesel and Paraschiv (2017) in their econometric analysis of 15-minute intraday electricity prices on EPEX Germany, and by Kremer, Kiesel, and Paraschiv (2021) in their econometric model of German intraday continuous trading. Both studies work with absolute price changes in €/MWh as the dependent variable in regressions of price movements on fundamental drivers, for the same reasons outlined above. Hagemann (2015) similarly analyses German intraday price dynamics in levels rather than returns.

The specification has a direct economic interpretation: a regression coefficient represents the expected price adjustment in £/MWh per unit change in the explanatory variable (for example, per MW of wind forecast revision or per MW of generation outage). This is the unit in which traders evaluate positions and the unit in which imbalance exposure is priced. Differencing also removes much of the intraday seasonality and serial correlation that characterise electricity price levels (Kiesel & Paraschiv, 2017), which then gives a dependent variable that is closer to stationary and better suited to standard inference. As a robustness check, we re-estimate our main specifications using price changes normalised by the rolling one-hour VWAP average, excluding observations where this average falls below £2/MWh in absolute value, and confirm that the qualitative results are unchanged.

3.2.2 Event study framework

Unlike predictive models that focus on forecast accuracy, the event-study approach allows us to isolate the causal effect of clearly timed information arrivals on prices. This is particularly appropriate in intraday electricity markets, where fundamental information arrives at discrete, observable timestamps and directly affects imbalance risk. To address the research question, our approach is guided by Kremer et al. (2021), who built on the work of Kiesel & Paraschiv

(2017) of studying intraday price determinants in Germany. Their approach regressed the nominal price change between two consecutive 1 minute VWAP within a contract on their lagged values, the latest solar and wind forecast changes and later added a merit-order-curve proxy. While their approach did show the significance of wind and solar forecast updates, it was more geared towards identifying patterns of when this renewable forecast impact was the largest. As we are interested and have access to a wider range of fundamental variables which are posted by the market operator, we see a better fit for an event study analysis, that is, looking at each forecast revision, REMIT message and imbalance price/volume as a significant event, from which affected intraday contract prices adjust.

This is akin to studies of earnings releases and macro announcements in traditional financial markets, for example, Smales (2013), which looked at the response of interest rate futures in Australia following macroeconomic announcements. From this, we adopt the use of bins to split the transactions into more representative chunks. Similar to illiquid after-hours trading studied by Smales, the GB electricity market also isn't as liquid as traditional markets, so binning reduces the amount of inefficient or mispriced transactions and looks at an average.

Drawing from the approaches discussed, to study the determinants of nominal intraday price changes, we use the following event-level general regression:

$$\Delta P_{c,i} = \alpha + \beta\, S_i + \varepsilon_{c,i}$$

where S_i is the surprise (information shock) associated with event *i*. For jointly published variables where multicollinearity is not a risk, this becomes:

$$\Delta P_{c,i} = \alpha + \beta_1\, S_{1,i} + \beta_2\, S_{2,i} + \varepsilon_{c,i}$$

The surprise is defined as:

$$S_i = f_c^{(k)} - f_c^{(k-1)}$$

Where ${f_c}^{(k)}$ is the k-th published value of the forecast for the delivery period affected by event *i*, and ${f_c}^{(k-1)}$ is the immediately preceding publication. For imbalance prices and volumes, we test both the revision difference as well as the level ${f_c}^{(k)}$.

3.2.3 Subsample analysis

To examine heterogeneity in the price response across different market conditions, we estimate the baseline regression across several dimensions. First, we estimate the regression separately by calendar year to assess whether the sensitivity of intraday prices to fundamental information has evolved as the market has matured. The GB continuous intraday market has experienced substantial growth in trading volumes and participant diversity over our sample period.

Furthermore, with increases in renewable production capacity, we would expect the price sensitivity to wind and solar forecast revisions to strengthen over time, as renewables constitute a larger share of the generation mix and forecast errors translate into larger absolute imbalance volumes.

Second, we partition events by the time remaining until gate closure at the moment of publication. For forecast revisions (LOLP, derated margin, wind, and solar), we define four buckets: less than 1 hour, 1-2 hours, 2-4 hours, and 4 or more hours before gate closure. For imbalance prices, where the contract matching produces events concentrated near gate closure, we use a simpler split of either less than 5 or more than 5 minutes before gate closure. This decomposition tests whether proximity to gate closure amplifies the price response, as traders face increasing urgency to close positions and have fewer remaining opportunities to offset new information through subsequent trades.

Third, we vary the post-event observation window τ^{post} while holding the pre-event window fixed at 180 seconds. We estimate the regression for $\tau^{post} \in \{10,30,60,180,300,600\}$ seconds. This horizon analysis adapting Smales (2013) approach serves two purposes: it reveals how quickly information is incorporated into prices, and whether the initial price adjustment persists or reverses at longer horizons. A coefficient that is large at short horizons and stable or growing at longer horizons is consistent with efficient information incorporation. A coefficient that peaks and then declines suggests an initial overreaction followed by partial reversal, while a coefficient that grows with the horizon indicates slow price discovery.

Lastly, for REMIT events, we examine a broader set of subdivisions given the richer metadata available. We separate events by shock direction (capacity deficit versus surplus), by whether the outage was classified as planned or unplanned, by the magnitude of the capacity change (using size buckets of 0-50, 50-100, 100-200, 200-500, 500-1000, and 1000+ MW), and by whether the message represents the initial announcement or a subsequent revision. These cuts allow us to test whether the market responds asymmetrically to supply losses versus returns to service, whether unplanned outages produce larger price adjustments than scheduled maintenance, and whether the magnitude of the response scales with the size of the capacity shock.

3.2.4 Impulse Response Analysis

To visualise the dynamics of price adjustment at high frequency, we construct event-time impulse response functions. Rather than estimating a single pre-to-post price change, we track the evolution of the VWAP relative to a stable baseline across fine time bins around each event.

The baseline VWAP is computed over the interval $[t_i - 360s, t_i - 60s]$, providing a 300-second reference that ends one minute before the event to avoid contamination from anticipatory trading or timestamp misalignment. We then construct 5-second VWAP bins covering 60 seconds before and 300 seconds after the event. For each bin b, the price displacement is:

$$\Delta P_{c,i}^{b} = \mathrm{VWAP}_{c,i}^{b} - \mathrm{VWAP}_{c,i}^{\mathrm{base}}$$

The pre-event bins serve as a placebo test: if the matching and timestamps are correct, no systematic price movement should be visible before the event. To avoid cancellation between positive and negative shocks, we condition on the magnitude of the surprise, plotting the impulse response separately for the top and bottom 5% of the shock distribution with 95% confidence intervals. For forecast revisions and REMIT events, we restrict the sample to the latest revision per contract, as this is the most informationally relevant update for each delivery period. For imbalance prices, we restrict to events where the next tradeable contract is within 5 minutes of gate closure.

4. Analysis of Results

4.1 Descriptive statistics

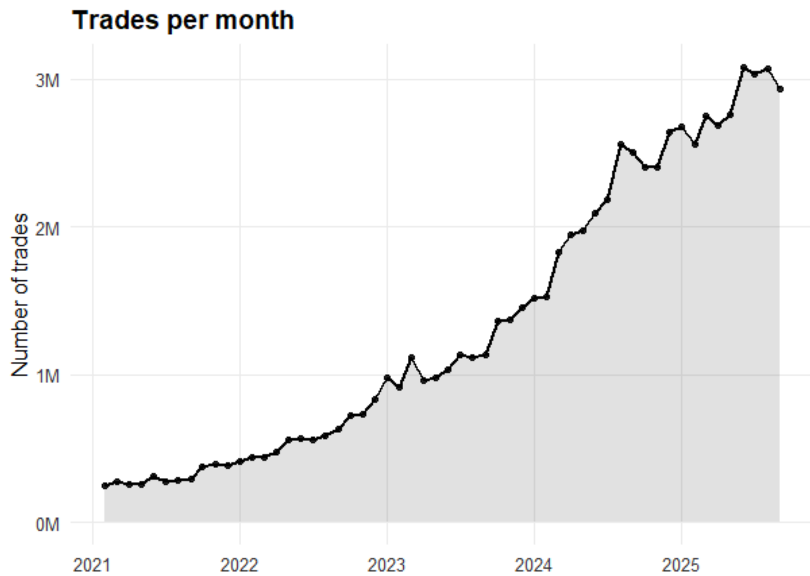

Figure 4. Trades per month, through years 2021-2025.

The number of monthly transactions on the EPEX GB intraday continuous market has grown roughly tenfold over the sample period. It has increased from around 200-300 thousand trades per month in early 2021 to consistently around 3 million by mid-2025. This trend has a direct implication for both research questions. For RQ1, rising liquidity means that the price discovery process has likely become more efficient over time. A thicker order book and more continuous trading allow fundamental information to be incorporated faster. This could imply that the sensitivity of intraday prices to forecast revisions and system tightness indicators has strengthened over the sample period, or alternatively that the adjustment window has narrowed.

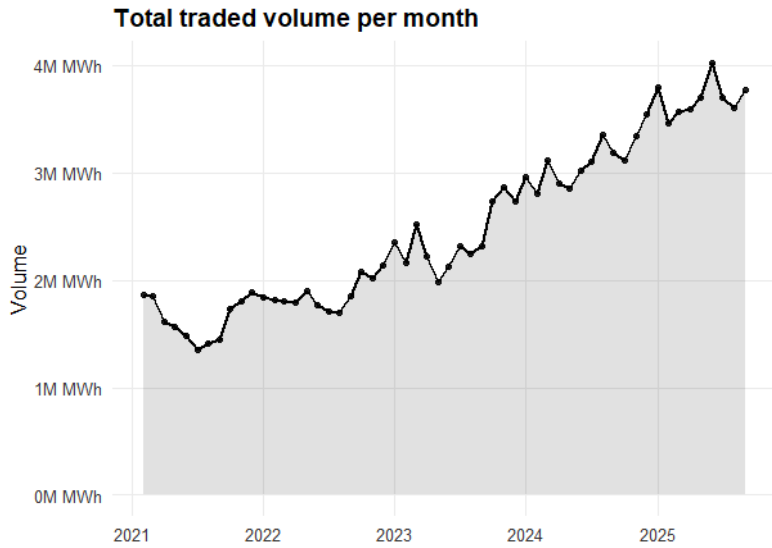

Figure 5. Traded Volume per month, through years 2021-2025.

Monthly traded volume follows a similar upward trajectory, growing from roughly 1.5-2M MWh in 2021 to nearly 4M MWh by mid-2025. The growth in volume is somewhat less dramatic than in trade count, as volume roughly doubled while trade count increased tenfold. This implies that average trade size has fallen substantially over the period. This is consistent with increasing algorithmic participation, where strategies tend to slice orders into smaller, more frequent transactions. Seasonal patterns are also visible, with slight dips during summer months when demand, residual load, and hence price volatility are lower, consistent with the merit-order effect described by Cludius et al. (2014) and Ketterer (2014).

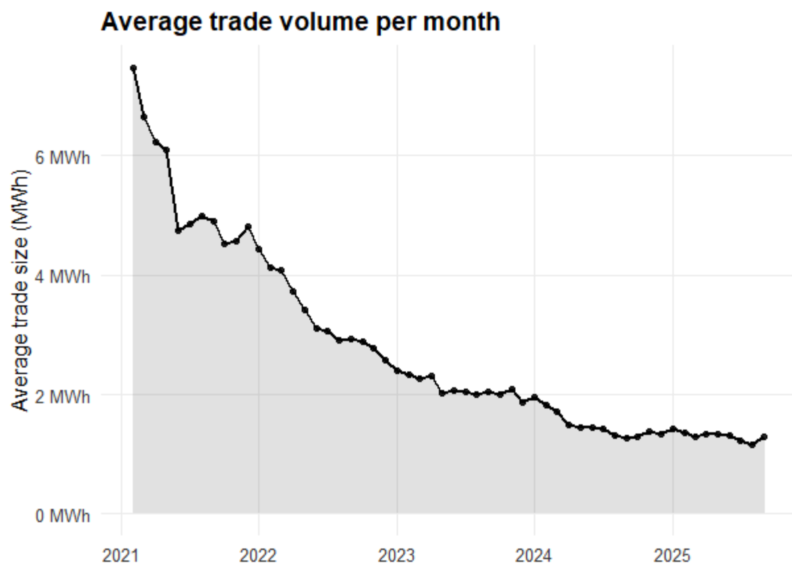

Figure 6. Average Trade Size (MWh) through years 2021-2025.

For our methodology, the declining average trade size reinforces the choice of volume-weighted average prices (VWAP) rather than last-trade or simple average prices. As Kremer et al. (2021) note, VWAP is the appropriate measure in low-liquidity or fragmented markets because it prevents small outlier trades from distorting the price signal. The growing volume also suggests that the market is increasingly being used for its intended purpose, which is correcting positions as forecast revisions and system information arrive.

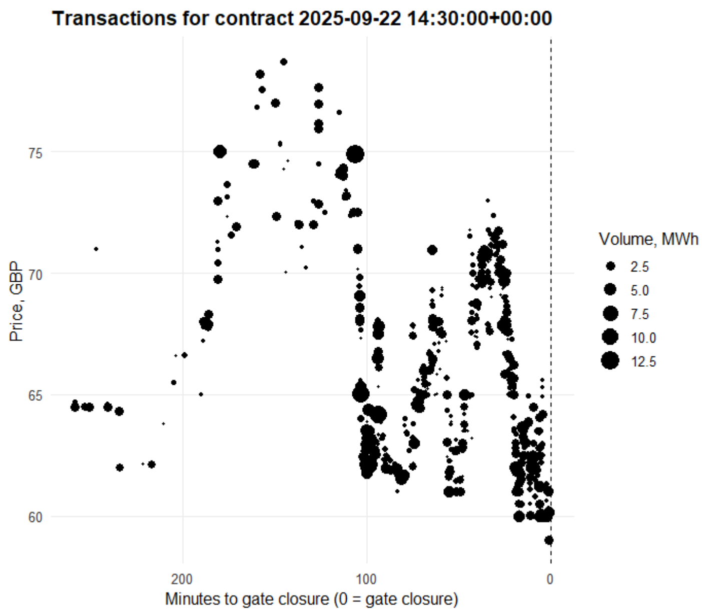

Figure 7. Example of a transaction scatter for a contract.

This scatter plot shows all individual transactions for a single half-hour delivery contract, with price (GBP/MWh) on the y-axis and minutes to gate closure on the x-axis, bubble size representing volume. As can be seen, little trading takes place beyond 200 minutes before gate closure. Liquidity increases dramatically inside 100 minutes to delivery, where a dense cluster of transactions forms. In this particular example, there is a clear downward price drift, which suggests that updated forecasts, potentially higher-than-expected wind or solar output, pushed the market down as delivery approached. This example illustrates the core mechanism behind RQ1. As NESO publishes updated wind and solar forecasts on the 12th minute of each hour, each revision introduces a medium-term information shock that shifts expectations for the remaining settlement periods. The price path visible here is consistent with the hypothesis drawn from Kiesel and Paraschiv (2017) that the impact of forecast revisions increases closer to delivery, as the revisions represent larger surprises relative to remaining uncertainty. The sparse, high-variance trading early in the session versus the concentrated, convergent trading near gate closure also motivates our methodological choice to examine how regression coefficients vary across contracts at different times to delivery, as the same variable (e.g., a wind forecast revision) may have a negligible effect 4 hours out but a highly significant one within the last hour.

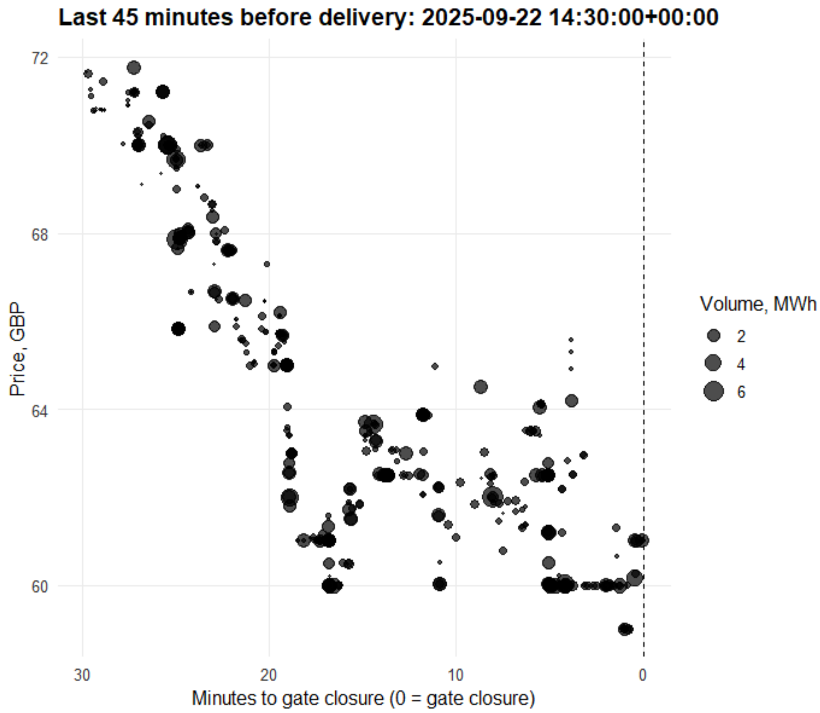

Figure 8. Transactions during the last 45 minutes before delivery.

This data shows the most liquid and informationally rich part of trading, In the graph we can see that the drop in price is not smooth, but moves in discrete steps, which is in line with new information shocks arriving at identifiable timestamps. This figure directly motivates our research question, as the discrete price steps suggest that the market does not adjust smoothly but reacts in bursts to specific information events, which is the type of short-run dynamics our event study methodology is designed to capture.

Descriptive statistics for intraday trades (GB Half-Hour Power)	
Statistic	Value
Number of trades	70 208 147
Number of trade IDs	70 208 147
Number of contracts	83 957
Mean price (GBP/MWh)	93.82
Std. dev. price	66.82
Minimum price	-200.00
1st percentile	-15.00
5th percentile	8.28
Median price	85.20
95th percentile	200.00
99th percentile	359.00
Maximum price	5 803.00
Mean trade size (MWh)	1.31
Total volume (MWh)	91 705 695
First trade	2020-12-31 18:48:53
Last trade	2025-10-17 21:14:54

Figure 9. Trade data summary

The sample covers the full history of the EPEX GB continuous intraday market from 31 December 2020 through 17 October 2025, totalling 70.2 million transactions across 83,957 delivery contracts. The total traded volume is 91.7 TWh with a mean trade size of 1.31 MWh. Prices range from -£200/MWh to £5,803/MWh around a mean of £93.82 and standard deviation of £66.82. For the analysis we will be using trades until 30th september.

Mean trades in symmetric X-second window around random times, by year						
Year	10s	30s	60s	180s	300s	600s
2021	11.20	14.60	18.30	28.80	39.30	63.10
2022	13.70	19.40	26.60	48.50	67.50	104.00
2023	12.20	18.40	26.80	54.00	81.60	137.00
2024	13.10	22.00	30.90	66.40	102.00	183.00
2025	14.10	25.30	40.10	95.60	145.00	271.00
All	12.90	19.60	27.40	55.00	81.00	140.00

Figure 10. Mean number of transactions observed in symmetric windows of varying width. 100 randomly sampled timestamps, by year.

A key decision in our empirical design is the width of the VWAP aggregation window used to construct the dependent variable. The window must be wide enough to contain a sufficient number of trades for a meaningful volume-weighted average, yet narrow enough to isolate the price response to a specific information event without contamination from subsequent shocks. To inform this choice, we conduct a liquidity analysis by sampling 100 random timestamps uniformly across the trading day and counting the number of transactions executed for the same settlement period contract within symmetric windows of varying width around each timestamp. At very short windows of ± 10 or ± 30 seconds, the median contract has fewer than 20 trades, which makes the resulting VWAP noisy and sensitive to individual transactions. At ±600 seconds (10 minutes), the window captures over 140 trades on average across the full sample, but for an event study this width risks blending the response to the event of interest with subsequent information arrivals, particularly given that NESO forecast revisions and IRIS imbalance publications can arrive within minutes of each other. The ±180-second (3-minute) base window offers a practical balance: it contains an average of 55 trades across the full sample and 96 trades in 2025, which is sufficient to construct a stable VWAP, while remaining narrow enough to attribute the observed price movement to a single identifiable event.

Descriptive statistics for forecast variables

Variable	Obs	Contracts	Mean	Std. dev.	Min	Max
Solar	2 467 718	83 089	1 534.00	2 326.00	0.00	13 850.00
Wind	2 467 718	83 089	1 778.00	1 110.00	121.00	5 945.00
Derated Margin	412 560	83 068	14 287.00	6 043.00	-636.00	34 948.00
LOLP	412 560	83 068	0.0002	0.0062	0.00	0.80

Figure 11. Forecast variable data summary

The forecast variables (Embedded Wind and Solar) are observed at much higher frequency than the system variables because NESO publishes updated forecasts every hour, and each forecast update is matched to all contracts still open for trading at that time. This produces 2.47 million forecast-contract observations. Wind forecasts average 1,778 MW and solar 1,534 MW, but solar has roughly double the standard deviation (2,326 vs 1,110 MW). This is of course due to the fact that solar production is 0 during the night, whereas wind production is more evenly distributed across the day.

The system-level variables tell a different story. Derated Margin averages 14,287 MW across 412,560 observations, which indicates that GB generally operates with substantial reserve capacity. However, the range extends to -636 MW, which does confirm that genuine capacity shortfalls do occur. The fact that DM has roughly more observations than Imbalance Price (412,560 vs 84,709) despite covering a similar number of contracts (83,068 vs 84,709) reflects that DM is published at multiple intervals (8, 4, 2, and 1 hours ahead of each settlement period), which gives traders progressively updated, granular information about system conditions throughout the day.

Descriptive statistics for imbalance price and volume	
Statistic	Value
Observations	37 223
Contracts	37 188
First publish	2024-01-31
Last publish	2026-03-26
Mean price (GBP/MWh)	76.60
Std. dev. price	58.10
Min price	-96.20
Max price	2 900.00
Mean volume (MWh)	-43.80
Std. dev. volume	312.00
Min volume	-2 834.00
Max volume	2 545.00

Figure 12. Imbalance price/volume data summary

Figure 12 reports the distributional properties of the imbalance settlement data that forms the basis for Part 3 of our empirical analysis. Each observation represents the first indicative system price published by IRIS for a given settlement period. The minor discrepancy between the total number of publications (37,223) and unique contracts (37,188) arises from a small number of settlement periods for which IRIS released near-simultaneous initial publications; in these cases, we retain only the earliest record.

The mean imbalance price of £76.60/MWh is broadly consistent with the average intraday VWAP over the same period, which aligns with the theoretical expectation that intraday prices should converge toward the expected imbalance price as delivery approaches (Kremer et al., 2021). However, the distribution exhibits pronounced right skewness: while typical settlement periods clear at moderate levels, the maximum reaches £2,900/MWh, which shows that the extreme costs that balance responsible parties face during periods of acute system tightness. These tail realisations are economically significant. As discussed in our relevance analysis, it is precisely during such episodes that active intraday trading is most valuable, because even a modest reduction in the residual imbalance volume can avoid the activation of the most expensive balancing reserves, generating disproportionate savings that ultimately flow through to consumers (Koch & Hirth, 2019).

The negative mean NIV of -43.80 MWh indicates that the GB system was, on average, slightly long over this period, meaning that aggregate physical generation marginally exceeded consumption. This is consistent with the incentive structure created by the single cash-out price reform (Ofgem, 2014b), which was fully implemented in November 2018 and made the imbalance price fully marginal by reducing the Price Average Reference volume to 1 MWh (Ofgem, 2014b). Under single pricing, a BRP that is short when the system is short faces the full marginal cost of the most expensive balancing action, whereas a BRP that is long when the system is short is paid the same price. This asymmetric risk encourages participants to maintain a slightly long position as a buffer, which manifests in the negative mean NIV observed in our data. Eicke et al. (2021) document an analogous response in Germany, where market participants systematically adjust their portfolio positions in response to the expected imbalance price, confirming that imbalance settlement design directly shapes intraday trading behaviour.

Descriptive statistics for REMIT outage events

Statistic	Value
Events	504 459
Unique MRIDs	34 018
Contracts	79 794
First publish	2021-01-01
Last publish	2025-09-30
Mean delta (MW)	144.00
Std. dev. delta	225.00
Min delta	-1 400.00
Max delta	2 000.00

Figure 13. Descriptive statistics for REMIT outage events

This figure summarises the REMIT outage dataset used in Part 2 of the analysis. The raw data is sourced from the Elexon BMRS REMIT feed, which publishes all Urgent Market Messages submitted by GB generators in compliance with UK REMIT. We process the raw messages by collapsing multi-row revisions within each MRID and computing the incremental change in unavailable capacity (delta MW) between successive publications, as described in Section 3.1.4. Each event is then expanded into individual settlement period observations covering the affected delivery window, yielding 504,459 event-SP pairs across 79,794 unique contracts, derived from 34,018 unique outage events.

The mean delta is 144 MW, consistent with the capacity of a mid-sized CCGT. The standard deviation of 225 MW indicates substantial heterogeneity in event magnitude. The maximum delta of 2,000 MW corresponds to the loss of an interconnector, while the minimum of −1,400 MW represents a return to service. A single outage typically spans multiple settlement periods: a gas plant tripping at 14:00 and returning at 18:00 affects eight consecutive half-hourly contracts, each of which constitutes a separate observation in our Part 2 specification.

4.2 Part 1 | Forecast Revisions and System Tightness

4.2.1 Wind and Solar Forecast Revisions

Wind and solar: joint regression results

		ΔWind		ΔSolar	
Group	n	Coef.	p-value	Coef.	p-value
Baseline					
All	143 759	-5.2304e-04	1.29e-01	-3.2727e-04**	3.34e-02
By year					
2021	11 564	1.8057e-04	9.32e-01	-1.0724e-03	2.63e-01
2022	16 006	-2.9448e-03	1.33e-01	-2.2737e-04	8.11e-01
2023	28 962	-6.7809e-04	3.27e-01	-5.8875e-05	8.63e-01
2024	49 726	4.7681e-04	1.92e-01	-4.7435e-04***	1.07e-03
2025	37 501	-2.0951e-04	5.59e-01	-2.1947e-04	1.75e-01
By GC bucket					
<1h	54 591	-1.5324e-03**	4.74e-02	-5.3280e-04*	7.23e-02
1–2h	46 803	-3.1074e-04	5.37e-01	1.4957e-04	5.35e-01
2–4h	36 204	3.7183e-04	3.52e-01	-4.4314e-04**	3.35e-02
≥4h	6 161	9.3887e-04	4.24e-01	-3.8947e-04	4.73e-01
By horizon					
10s	43 763	-2.7781e-04	5.22e-01	-1.4846e-04	4.25e-01
30s	77 169	-9.7063e-05	7.87e-01	-2.4849e-04	1.11e-01
60s	103 391	-2.6675e-04	4.23e-01	-4.6764e-04***	1.45e-03
180s	143 759	-5.2304e-04	1.29e-01	-3.2727e-04**	3.34e-02
300s	158 924	-4.5730e-04	2.14e-01	-4.1661e-04**	1.18e-02
600s	172 021	-7.6269e-04*	7.27e-02	-3.1029e-04	1.05e-01

Figure 14. Significance of joint wind/solar forecast adjustments on intraday price change.

Figure 14 reports the results of regressing the intraday price change on the revisions to NESO's wind and solar generation forecasts, estimated jointly since both variables are published simultaneously on the 12th minute of each hour. Across all 143,759 matched forecast-contract events, the wind revision coefficient does not reach statistical significance. The solar revision coefficient is significant at the 5% level. Both coefficients carry a negative sign, which is directionally correct: an upward revision in expected renewable generation increases anticipated supply for the affected settlement period, reducing residual load and pushing the intraday price downward.

When the forecast revision arrives less than one hour before gate closure, wind becomes significant at 5% and solar is marginally significant at 10%. At longer horizons (1–2h, ≥4h), neither variable is significant. The strongest results for both variables are concentrated in the <1h bucket. Across post-event horizons, solar first reaches significance at 60 seconds and remains significant through 300 seconds, while wind is insignificant at all horizons except 600 seconds.

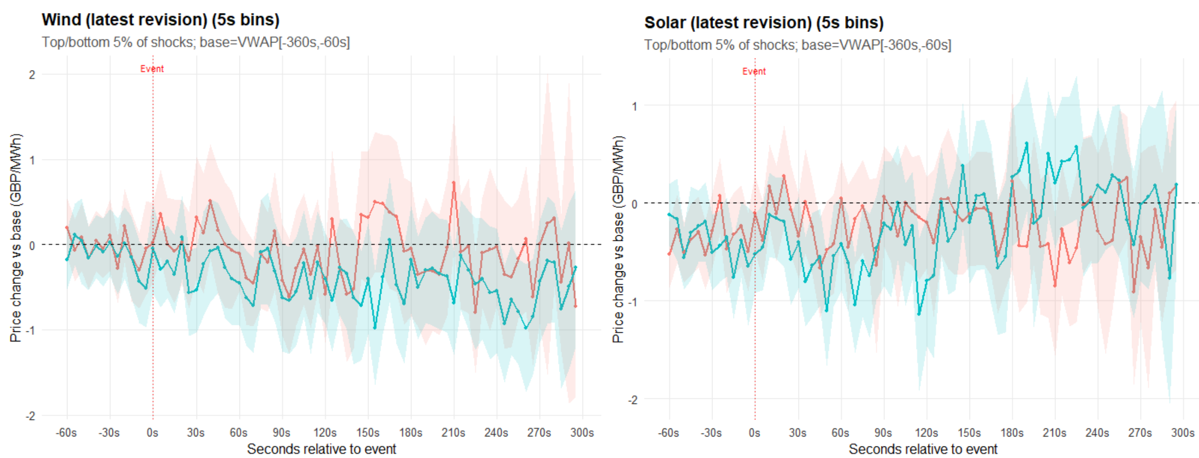

Figure 15. Impulse response functions on top/bottom 5% of solar/wind forecast revisions by 5s bins (red - bottom 5% / blue - top 5%).

These graphs show the impulse response (IR) for the top and bottom 5% of wind and solar revisions. Neither variable displays a sharp post-event adjustment; the responses drift gradually with wide confidence bands that overlap through most of the post-event window. In testing, restricting the sample to contracts closing within 3 minutes of the forecast release produces no clearer pattern. As the revisions happen on an hourly basis, the mean revision is relatively small, so even in the extremes there might not be enough variation to merit a response from the market.

4.2.2 Loss of Load Probability and Derated Margin

LOLP and derated margin: joint regression results					
		ΔLOLP		ΔDM	
Group	n	Coef.	p-value	Coef.	p-value
Baseline					
All	118 998	-3.9902e+01***	3.17e-17	-8.1770e-05**	1.60e-02
By year					
2021	9 195	-2.0167e+02***	1.41e-04	2.2271e-04	4.45e-01
2022	15 088	-3.9325e+01***	1.36e-03	-1.6415e-04	3.72e-01
2023	26 559	3.0890e+01***	6.23e-03	-3.4189e-05	5.49e-01
2024	38 775	-4.0700e+01***	1.82e-31	-9.6338e-05***	1.59e-04
2025	29 381	-2.0596e+02***	4.75e-30	-1.0937e-04***	3.29e-03
By GC bucket					
<1h	64 378	-2.4170e+01***	5.35e-03	-1.2347e-04**	4.64e-02
1–2h	44 024	-7.1253e+01***	5.00e-49	-4.1456e-05	2.48e-01
2–4h	10 508	3.1015e+00	6.44e-01	-5.5182e-05	2.47e-01
≥4h	88	-7.5608e+00	9.25e-01	4.0517e-04	3.99e-01
By horizon					
10s	36 987	8.0891e+00	1.92e-01	-1.1469e-04***	2.75e-03
30s	68 123	-1.4655e+01***	6.08e-03	-1.4554e-04***	1.80e-05
60s	90 214	-2.6365e+01***	2.56e-07	-1.1043e-04***	9.92e-04
180s	118 998	-3.9902e+01***	3.17e-17	-8.1770e-05**	1.60e-02
300s	128 580	-2.5271e+01***	6.35e-08	-9.0213e-05**	1.22e-02
600s	137 004	-1.9818e+01***	1.44e-04	-1.3264e-04***	8.30e-04

Figure 16. Significance of joint LOLP/DM forecast adjustments on intraday price change.

The table reports the results of regressing the intraday price change on the joint revisions to LOLP and derated margin. Both coefficients are statistically significant. The DM coefficient is negative, meaning that an increase in forecasted spare capacity is associated with a decrease in the intraday price. DM is significant across all individual years, significant in the <1h and 1–2h gate closure buckets, and significant at all post-event horizons from 30 seconds onward. This makes it the most robust fundamental driver in our Part 1 analysis.

The LOLP coefficient is also negative, which is counterintuitive: a higher LOLP indicates a greater probability that available generation will be insufficient to meet demand, which should push prices upward rather than downward. However, as reported in Section 4.1, the mean LOLP in our sample is 0.0002 with the vast majority of observations at or near zero. Given that LOLP revisions are effectively zero for most events and only take non-trivial values during rare scarcity episodes, the coefficient is driven by a small number of extreme observations and should not be interpreted as a reliable marginal effect.

4.2.3 Derated Margin Impulse Response

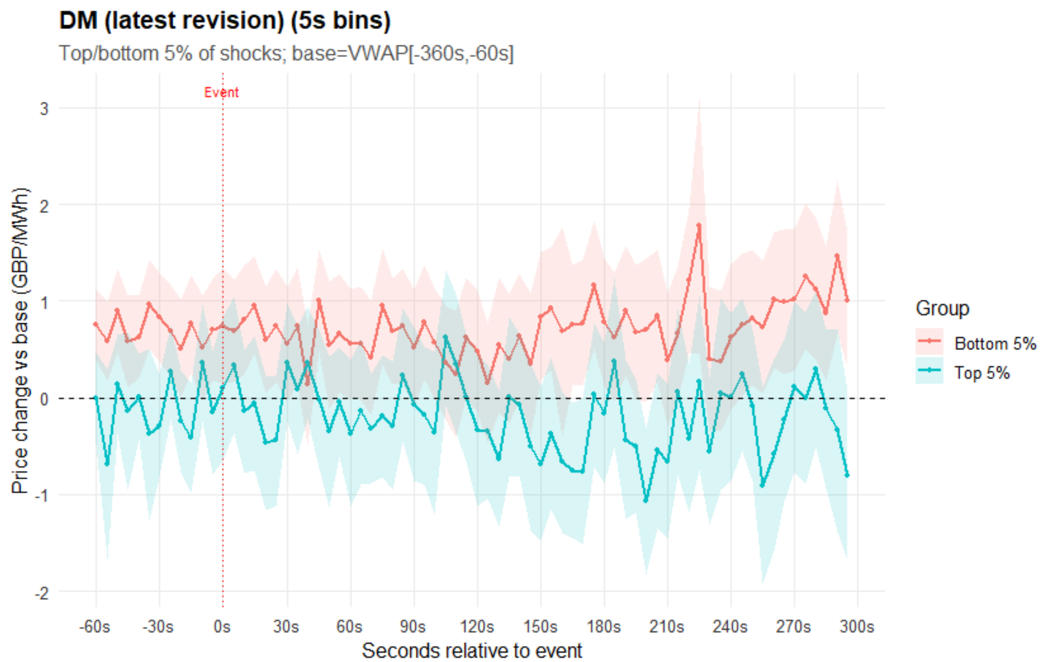

Figure 17. Impulse response functions on top/bottom 5% of DM forecast revisions by 5s bins.

The graph shows the impulse response for DM revisions. The bottom 5% of DM shocks (tightening system) are associated with an upward price drift of approximately +1.0 £/MWh by 180 seconds. The top 5% (loosening system) produce a roughly symmetric downward drift. The pre-event period shows no systematic trend for either group, confirming that the event timing is correctly identified. However, the bottom 5% of revisions show an already elevated price level relative to the baseline even before the publication timestamp. This is consistent with tight system conditions being partially observable through other channels, such as earlier DM publications at longer lead times, or correlated signals from low wind forecasts and high demand, before the specific revision captured in our event window. The post-event drift then represents the additional adjustment attributable to the new DM publication itself.

4.3 Part 2 | REMIT Outage Events

REMIT outage events: regression results

Group	n	Coef.	p-value	Group	n	Coef.	p-value
Baseline				*By year*			
All	69 869	1.3093e-03***	4.36e-24	2021	7 659	3.0225e-03***	8.65e-04
By shock direction				2022	7 832	2.9690e-03***	1.54e-06
Deficit	62 080	1.5021e-03***	1.19e-21	2023	18 533	1.0655e-03***	4.14e-10
Surplus	7 789	4.3682e-04	1.66e-01	2024	19 580	5.6829e-04***	1.71e-11
By planned/unplanned				2025	16 265	1.0156e-03***	4.74e-22
Planned	25 700	2.1176e-03***	1.64e-11	*By GC bucket*			
Unplanned	44 169	9.9981e-04***	5.91e-15	<1h	37 793	1.7189e-03***	4.65e-18
By initial/revision				1–2h	19 350	1.1118e-03***	2.53e-06
Revision	13 498	1.1359e-03***	2.85e-07	2–4h	11 401	4.6398e-04***	5.98e-04
Initial	56 371	1.5195e-03***	6.33e-20	≥4h	1 325	-1.2366e-03**	2.82e-02
By size bucket (MW)				*By horizon*			
0–50	29 334	-4.6031e-04	8.65e-01	10s	21 815	6.6534e-04***	1.19e-06
50–100	9 661	-4.9230e-05	9.78e-01	30s	39 902	9.9947e-04***	1.98e-12
100–200	8 186	2.7709e-04	8.38e-01	60s	52 890	1.1474e-03***	4.10e-20
200–500	17 649	1.2732e-03***	7.83e-09	180s	69 869	1.3093e-03***	4.36e-24
500–1000	4 883	2.4988e-04	3.86e-01	300s	75 013	1.4290e-03***	1.27e-24
1000+	156	5.5156e-04	7.57e-01	600s	79 530	1.8249e-03***	8.56e-26

Figure 18. Significance of REMIT UMM size and other factors on intraday price change.

Across 69,869 event-contract pairs, the coefficient is positive and highly significant. A generation outage that removes capacity from the system raises the intraday price for the affected settlement periods. Capacity deficits (62,080 events) produce a positive and highly significant coefficient. Capacity surpluses, i.e. returns to service (7,789 events), produce a positive but statistically insignificant coefficient. The market responds detectably to supply losses but not to supply returns. Both planned (25,700 events) and unplanned (44,169 events) outages produce significant positive coefficients. The planned outage coefficient is roughly double the unplanned coefficient in magnitude. Outages below 200 MW do not produce significant coefficients. The 0-50 MW bucket coefficient is negative but insignificant. The 200-500 MW bucket (17,649 events) is the only size group that is significant. The 500-1000 and 1000+ MW buckets are positive but insignificant, with small sample sizes of 4,883 and 156 events respectively.

The coefficient is positive and significant in the <1h, 1-2h, and 2-4h buckets, with the largest magnitude in the <1h bucket and declining through each subsequent bucket. The ≥4h bucket (1,325 events) produces a negative and significant coefficient at the 5% level, meaning outages announced more than 4 hours before gate closure are associated with a price decrease rather than an increase for the affected contracts.

The coefficient is positive and significant at all horizons from 10 seconds through 600 seconds. The magnitude increases monotonically with horizon length: it is smallest at 10 seconds and roughly triples by 600 seconds. The price adjustment continues to accumulate for at least 10 minutes after the UMM is published, rather than completing in a single jump.

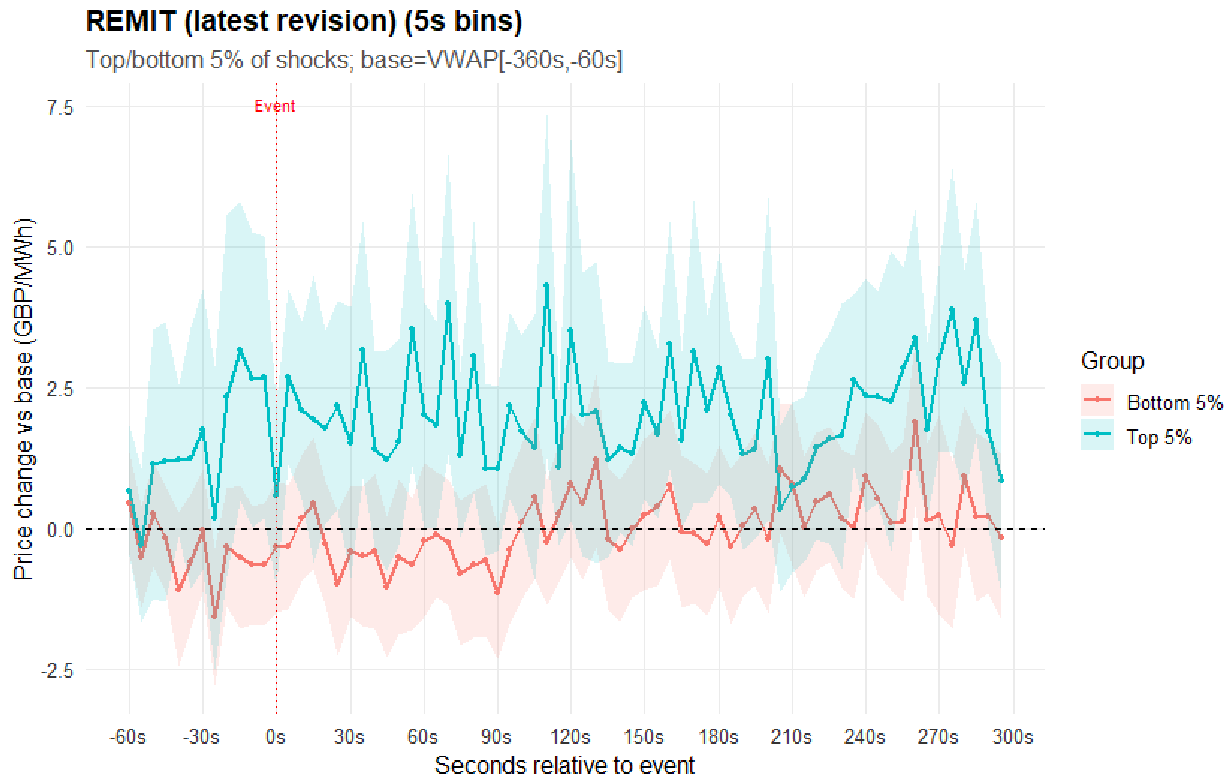

Figure 19. Impulse response functions on top/bottom 5% of REMIT announcements by size.

The figure shows the impulse response for the top and bottom 5% of REMIT shocks (latest revision per contract), using 5-second bins. The top 5% of shocks (large capacity losses) produce a clear upward drift, rising from near zero at the event time to approximately +3 to +5 £/MWh by 300 seconds. The bottom 5% (large capacity returns) show a much flatter response, hovering near zero throughout the post-event window. The pre-event period is elevated for the top 5%, which is indicative of information being incorporated into the price prior to the announcement.

4.4 Part 3 | Imbalance Price and Imbalance Volume

4.4.1 Delta Specification

Imbalance price and volume: delta regressions

		ΔIMB price		ΔIMB volume	
Group	n	Coef.	p-value	Coef.	p-value
Baseline (univariate)					
All	28 760	1.4195e-02***	2.34e-136	2.5517e-04*	6.50e-02
Baseline (joint)					
All	28 760	1.5638e-02***	2.15e-146	-1.0236e-03***	1.91e-12
By year					
2024	15 218	1.9120e-03**	3.02e-02	-1.0932e-04	4.07e-01
2025	13 542	1.7134e-02***	1.32e-100	6.2351e-04**	1.22e-02
Near-GC split					
$<$5 min	5 033	2.6114e-03	1.22e-01	-2.5550e-04	2.32e-01
$\geq$5 min	23 727	1.4785e-02***	1.09e-125	3.6356e-04**	2.43e-02
By horizon					
10s	15 532	2.7039e-02***	7.42e-253	3.0717e-04*	8.89e-02
30s	23 572	1.6566e-02***	3.88e-180	1.1989e-04	3.58e-01
60s	27 332	1.6450e-02***	3.51e-227	2.0290e-04	1.07e-01
180s	28 760	1.4195e-02***	2.34e-136	2.5517e-04*	6.50e-02
300s	28 786	1.4464e-02***	4.18e-106	2.7779e-04*	8.22e-02
600s	28 787	2.1177e-02***	9.15e-128	5.1958e-04**	1.47e-02

Figure 20. Significance of univariate imbalance price/volume delta on intraday price change.

In the univariate regressions across 28,760 events, both ΔIMB price and ΔIMB volume are positive. ΔIMB price is highly significant; ΔIMB volume is marginally significant at the 10% level. Both signs are directionally correct: an increase in the previous settlement period's imbalance price or a shift toward system shortage (more positive volume) is associated with a higher intraday price for the next tradeable contract.

In the joint specification, the ΔIMB price coefficient remains positive and highly significant, while the ΔIMB volume coefficient flips to negative and becomes highly significant. This sign change reflects the different information each variable carries once estimated jointly. The ΔIMB price coefficient captures the marginal cost of balancing: a higher imbalance price signals scarcity and raises the expected outside option for intraday traders. The ΔIMB volume coefficient, conditional on the price change, captures a different signal. If the system was very short (large positive volume change) but the imbalance price did not spike proportionally, this indicates that the TSO resolved a large imbalance at a moderate cost, suggesting sufficient balancing reserve depth. Conditional on the same price movement, a larger volume therefore signals more available balancing capacity rather than more scarcity, which reduces concern about the next settlement period and pushes the intraday price down. The univariate specification gives the unconditional relationship that reflects the directional exposure traders face, while the joint specification decomposes the signal into a scarcity component (price) and a depth component (volume). The subsample analyses below use the joint specification.

The ΔIMB price coefficient is positive and highly significant at all horizons from 10 seconds through 600 seconds. The magnitude is largest at 10 seconds and declines through 30s and 60s before stabilising from 180s onward. ΔIMB volume is marginally significant at 10 seconds, insignificant at 30s and 60s, marginally significant at 180s, and significant at the 5% level at 300s and 600s. The pattern for ΔIMB price indicates that the initial price reaction is strongest immediately after publication and partially attenuates at longer horizons.

4.4.2 Level Specification

Imbalance price and volume: level regressions

		IMB price		IMB volume	
Group	n	Coef.	p-value	Coef.	p-value
Baseline (univariate)					
All	28 760	3.7793e-03***	3.21e-17	-2.4379e-04***	6.55e-03
Baseline (joint)					
All	28 760	5.6777e-03***	9.78e-29	-7.9006e-04***	1.00e-14
By year					
2024	15 218	-2.8389e-03***	3.22e-06	-4.4573e-04***	2.25e-07
2025	13 542	5.9363e-03***	1.70e-19	-4.2240e-05	7.92e-01
Near-GC split					
$<$5 min	5 033	-5.0248e-03***	4.65e-05	-4.3009e-04***	2.94e-03
$\geq$5 min	23 727	4.4685e-03***	1.01e-19	-2.0664e-04**	4.65e-02
By horizon					
10s	15 532	1.6620e-02***	2.35e-164	7.9585e-05	4.94e-01
30s	23 572	6.7995e-03***	2.65e-51	-5.2405e-05	5.36e-01
60s	27 332	8.8994e-03***	1.67e-100	-8.9416e-05	2.74e-01
180s	28 760	3.7793e-03***	3.21e-17	-2.4379e-04***	6.55e-03
300s	28 786	-3.5671e-05	9.45e-01	-2.7277e-04***	8.48e-03
600s	28 787	-5.9182e-04	3.92e-01	-2.8197e-04**	4.13e-02

Figure 21. Significance of univariate imbalance price/volume levels on intraday price change.

In the univariate regressions, both IMB price and IMB volume are highly significant. IMB price is positive: a higher absolute imbalance price in the previous settlement period is associated with a higher intraday price. IMB volume is negative: a more positive (system-short) imbalance volume is associated with a lower intraday price. In the joint specification, both remain significant with the same signs, though magnitudes shift. The negative IMB volume sign in the univariate level specification differs from the positive ΔIMB volume sign in the univariate delta specification. The two specifications capture different dynamics. The delta captures momentum: the system getting shorter than the period before signals deteriorating conditions, which pushes prices up. The level captures mean-reversion: when the previous period's absolute imbalance was large, market participants and the TSO respond by correcting positions and activating reserves, so the market expects the next period to be less extreme, which pushes prices down.

4.4.3 Imbalance Price and Volume Impulse Response

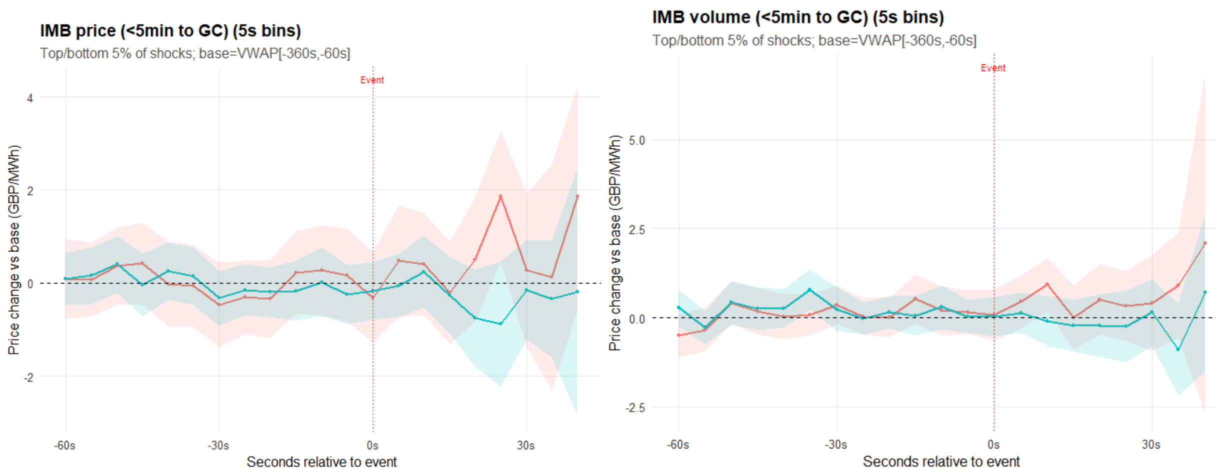

Figure 22. Impulse response functions on top/bottom 5% of imbalance price/volume levels by 5s bins (only announcements just before GC).

Figure 22 shows the impulse response for the top and bottom 5% of imbalance price and volume shocks, restricted to events where the next tradeable contract is within 5 minutes of gate closure. The pre-event period is relatively flat. After the event, the bottom 5% of price shocks (previous period settled at an unexpectedly low price) are associated with a rise in the intraday VWAP of approximately +2 £/MWh within 30 seconds. The top 5% show a slight downward drift with wide confidence bands. However, these impulse responses should be interpreted with caution. Because IRIS publications typically arrive around the 14th minute after each settlement period, the remaining trading window before gate closure of the next contract is extremely short. This means that the later 5-second bins in the IRF are populated by very few events where trading actually continued long enough to fill the full post-event window. The large price movements visible at the tail end of the IRF are therefore more likely to reflect thin liquidity and elevated volatility near gate closure than a genuine sustained price response to the imbalance publication.

5. Discussion

This section interprets the empirical findings presented in Section 4 in relation to the research question, the theoretical framework developed in the literature review, and the broader context of GB electricity market design. Our research question asks: What are the determinants of intraday electricity prices in Great Britain, and how efficiently do prices incorporate new fundamental information as delivery approaches? The results across all three parts of the analysis provide a detailed answer to both components of this question.

5.1 Determinants of Intraday Prices

5.1.1 The forecast data problem

The most important distinction between our study and the German literature concerns the nature of the forecast data used. Kiesel and Paraschiv (2017) and Kremer et al. (2021) both used a proprietary dataset of intradaily updated wind and solar forecasts provided by EWE TRADING GmbH, a commercial energy trading firm. Kremer et al. (2021) describe this as "a unique data set" that had, at the time, "solely been analysed by Kiesel and Paraschiv." These are not TSO forecasts. These are the kind of forecasts that an active trading desk would subscribe to from providers like Volue Insight, Meteorologica, Dexter Energy, or similar commercial services, forecasts that update at sub-hourly frequency, incorporate proprietary weather models, and are designed specifically to give traders an informational edge.

Our study, by contrast, uses NESO's publicly available wind and solar forecasts, which update once per hour on the 12th minute. This is the baseline information available to every market participant. Any trading firm with a commercial forecast subscription has access to higher-frequency, potentially more accurate forecasts, and will have already acted on them before NESO publishes. What our regression measures, therefore, is not the total information content of renewable forecast revisions for intraday prices, but rather the residual information content of the public TSO forecast after the market has already partially incorporated private forecast signals.

This distinction is essential for interpreting our results. Kiesel and Paraschiv (2017) found wind and solar forecast revisions to be significant at the 1% level, with strong explanatory power during peak hours. We find wind insignificant in the pooled sample and solar significant only at the 5% level. This does not mean that renewable forecasts matter less in Great Britain than in Germany. It means that the public hourly NESO forecast contains less marginal information than a proprietary trading-grade forecast. The fact that we find any significance at all for solar, and that wind becomes significant within the last hour before gate closure, indicates that the NESO publication does still contain some information that the market has not fully anticipated, but it is a smaller and noisier signal than what the German studies captured.

This has a direct practical implication. A new entrant to the GB intraday market relying solely on NESO's public forecasts would be trading on information that is partially stale by the time it is published. The significance of solar close to gate closure suggests that the public forecast retains some value in the final hour, when remaining uncertainty is highest and even small informational edges translate into trading decisions. But for wind, where the public forecast is insignificant at all but the longest horizon, the implication is that the market has already priced in the expected wind revision before NESO publishes it, and a participant without access to proprietary forecasts is systematically disadvantaged in this dimension.

This finding directly supports the argument made in the introduction that commercial data providers fill gaps left by public TSO data. We noted that these providers layer proprietary weather models and machine-learning forecasts on top of freely available fundamentals, and that the cost of these subscriptions represents a meaningful barrier to entry for smaller participants. Our results provide an indication of the information asymmetry this potentially creates: the public NESO wind forecast is effectively priced in before publication, meaning that participants without private forecasts cannot trade profitably on this signal. This is a concrete channel through which information barriers reduce competitive participation, consistent with the broader argument that market efficiency depends on the structure of how much of the market participants are informed (Batalla-Bejerano et al., 2022).

5.1.2 Gate closure proximity

Across all three parts of the analysis, the most robust and consistent finding is that the price response to fundamental information strengthens as gate closure approaches. Wind and solar become significant only within the last hour. DM is significant only in the <1h bucket. REMIT coefficients decline with distance to gate closure, and even flip sign for the ≥4h bucket. Imbalance price effects are largest for contracts more than 5 minutes from gate closure, where the signal still represents new information.

The universality of this pattern across fundamentally different information types, renewable forecasts, system tightness indicators, outage announcements, and imbalance settlement data, indicates that it reflects a structural feature of the continuous intraday market rather than a property of any individual variable. The mechanism is the approaching imbalance settlement price. As delivery nears, a trader who has not closed their position faces the imbalance price as the binding outside option. Any information that updates the expected imbalance price therefore has a higher marginal value closer to gate closure, because the cost of ignoring it is no longer hypothetical but imminent.

Kiesel and Paraschiv (2017) indicated this pattern, arguing that forecast revisions should have a larger impact closer to delivery because they represent a larger proportion of remaining uncertainty. Our results confirm this prediction for the GB market and extend it beyond renewable forecasts to system tightness indicators, outage events, and imbalance publications. The fact that the pattern holds across all information types strengthens the conclusion: it is the time structure of the market, not the nature of the information, that determines when fundamentals matter most.

5.1.3 System tightness and scarcity pricing

The derated margin is a significant negative determinant of intraday prices: more spare capacity is associated with lower prices. This confirms the mechanism described in Section 2.5.3, where DM affects intraday prices through its relationship to the expected Reserve Scarcity Price and hence the expected imbalance settlement price. The emergence of DM significance only in 2024 and 2025 has a straightforward interpretation: the tenfold growth in trading activity documented in Section 4.1 means that by 2024, sufficient participants were monitoring BMRS publications and actively trading on system tightness signals for these to be reflected in the VWAP. In 2021-2023, the market was simply too thin for the tightness signal to propagate reliably into traded prices.

LOLP, despite being statistically significant, carries a counterintuitive negative sign in the joint regression. As discussed in Section 4.2.2, this reflects the variable's distributional properties: LOLP is effectively zero for over 99% of observations and only spikes during rare extreme scarcity events. A linear regression coefficient has no meaningful marginal interpretation for a variable with this distribution. The derated margin, which is continuously distributed and mechanically related to LOLP, captures the economically relevant tightness signal.

5.1.4 REMIT outages and information asymmetry

REMIT outage events produce a significant positive price response, confirming that the GB intraday market prices supply losses. The asymmetry between deficits (significant) and surpluses (insignificant) reflects two distinct dynamics. First, supply losses are unambiguous shocks: a generator trips and capacity disappears from the system immediately. Returns to service, by contrast, are often anticipated (a scheduled maintenance window ending on time) and may already be priced in. Second, the merit order curve is convex: removing capacity from a tight system moves the marginal price along a steep portion of the supply stack, while adding capacity to a long system moves it along a flat portion. The same MW therefore produces an asymmetric price response depending on the direction of the shock.

The finding that planned outages carry a larger coefficient than unplanned outages is initially counterintuitive but has a clear explanation. Planned outages in the REMIT data represent large scheduled maintenance events that remove major units (nuclear, large CCGT) for extended periods and affect many consecutive settlement periods. Their average MW is typically larger than an unplanned trip, which often involves a partial reduction in output rather than a full unit loss. More importantly, planned outage announcements contain precise timing information (exact start and end hours) that resolves residual uncertainty and triggers immediate portfolio rebalancing, whereas unplanned trips may be announced with imprecise or evolving duration estimates that require subsequent revisions.

The pre-event price pattern in our REMIT impulse response is worth noting. The top 5% of outage shocks show an elevated price level before the publication timestamp. This could be a potential indication of information leakage, as UK REMIT (Article 4) requires market participants to disclose inside information as soon as possible. However, there are alternative explanations. Large outages tend to coincide with already tight system conditions that are observable through other public signals. Our framework does not isolate these channels, and testing directly for insider trading falls outside the scope of this thesis.

5.1.5 Imbalance prices as the dominant signal

The imbalance price publication produces the strongest and most immediate price response in our analysis, with significance appearing within 10 seconds in the delta specification. This confirms the theoretical mechanism described by Eicke et al. (2021): the imbalance settlement price is the opportunity cost of remaining imbalanced, and its publication directly updates traders' expectations of this outside option. Eicke et al. (2021) estimate cross-market equilibria between intraday and imbalance markets in Germany, finding that a shock to the imbalance price triggers a subsequent adjustment of the intraday price, which is precisely the dynamic our delta specification captures for the GB market.

The finding that the delta specification (change from the previous period's imbalance price) produces a stronger and more persistent signal than the level specification (absolute imbalance price) indicates that traders respond primarily to new information about the trajectory of system conditions. The level of the previous period's imbalance price may already be reflected in the prevailing intraday VWAP through prior trading, so the marginal information content lies in the direction of change rather than the absolute value. This is consistent with the level specification's IMB price coefficient becoming insignificant at the 300-second and 600-second horizons, while the delta specification remains significant throughout: the level produces a transient flash reaction, while the change carries persistent directional information.

The contrasting signs of the imbalance volume coefficient across specifications reveal that volume carries layered information. In the univariate delta specification, the coefficient is positive: the system getting shorter than the previous period signals deteriorating conditions, and the next-period intraday price rises. This is the unconditional momentum signal that a trader monitoring IRIS publications would observe directly. In the joint specification, conditional on the price change, the volume coefficient flips negative. This conditional sign makes economic sense: if the system was very short (large positive volume change) but the imbalance price did not spike proportionally, this indicates that the TSO resolved a large imbalance at moderate cost, signalling available balancing reserve depth rather than escalating scarcity. Holding the price movement constant, more volume therefore indicates more capacity headroom, which reduces scarcity concerns and pushes the intraday price down.

In the univariate level specification, the volume coefficient is negative, capturing a different dynamic: mean-reversion. When the absolute imbalance volume in the previous period was large, the market expects corrective action. BRPs who observed the large imbalance adjust their positions, and the TSO's concurrent balancing actions reduce the likelihood that the next period will be equally extreme. The market prices this expected correction, so a high absolute volume level is associated with a lower intraday price for the subsequent contract.

These three results (delta univariate positive, delta joint negative, level univariate negative) are internally consistent and collectively demonstrate that imbalance volume is not a simple directional indicator. It contains information about both the state of the system and the system's capacity to absorb shocks, and which signal dominates depends on whether volume is considered unconditionally or conditional on the price outcome.

5.1.6 The intraday-imbalance spread as an alternative measure

Our specifications measure the imbalance surprise either as the period-on-period change or the absolute level. An alternative approach would define the surprise as the difference between the contract's own closing VWAP and the subsequently realised imbalance price for the same settlement period. This spread would directly measure how well the intraday market anticipated the imbalance outcome: a narrow spread implies the market correctly priced the expected imbalance, while a wide spread implies a genuine pricing error.

This concept has empirical precedent. Koch (2021) explicitly analyses trading strategies in Germany that exploit the spread between intraday prices and expected imbalance prices, documenting asymmetric average spreads depending on the direction of the system balance. He demonstrates that a logistic regression model predicting the direction of the system balance can generate profitable strategies based on this spread, confirming that it contains exploitable information. Aïd, Cosso, and Pham (2022) provide theoretical foundations in an equilibrium model of intraday trading where agents minimise the sum of trading costs and imbalance penalties. In their framework, the imbalance cost directly shapes trading incentives and the equilibrium price dynamics, implying that the spread between the intraday price and the expected terminal settlement value should narrow as agents trade more actively closer to delivery.

In our context, one could define the spread for settlement period h as the difference between the final VWAP of contract h and the realised imbalance price for the same period, published by IRIS approximately 15 minutes after delivery. Regressing this spread on the fundamental variables from Parts 1 and 2, such as the last forecast revision before gate closure or the presence of a REMIT outage during the delivery period, would test whether specific information events systematically predict the market's pricing error relative to the realised imbalance outcome. This would provide a more direct test of informational efficiency than our current framework, which measures short-run price reactions around event timestamps but does not evaluate whether those reactions are ultimately correct in magnitude.

We do not pursue this specification for two reasons. First, the dependent variable (the spread) is only observable after delivery, which introduces a look-ahead element incompatible with our event study design, where identification relies on the precise timing of information arrival. Second, the spread conflates the market's pricing error with the unpredictable component of the balancing mechanism outcome, including TSO dispatch decisions, STOR activations, and Reserve Scarcity Price adjustments, none of which are known to intraday traders before gate closure. Decomposing the spread into a predictable component (driven by observable fundamentals) and an unpredictable component (driven by balancing mechanics) would require a structural model of the settlement process that is beyond the scope of this thesis. Nevertheless, the VWAP-imbalance spread represents a natural extension of our work and would provide a direct link between the intraday price discovery process we document and the ultimate settlement outcome.

5.2 Informational Efficiency

We do not test semi-strong efficiency in its strict form, which would require a benchmark for the 'correct' price level given all public information. Instead, our event study provides descriptive evidence on the speed and persistence of price responses to clearly identifiable information shocks. Where adjustment is fast and stable, we interpret this as consistent with efficient incorporation; where it drifts (REMIT) or overshoots (imbalance prices), we interpret this as evidence of underreaction or overreaction relative to that informal benchmark.

The second component of our research question asks how efficiently prices incorporate new fundamental information. Our methodological framework, adapted from Smales (2013), was designed to answer precisely this: by varying the post-event VWAP horizon and constructing second-by-second impulse response functions, we can measure both the speed and completeness of price adjustment to each information type.

The benchmark from traditional financial markets sets a high bar. Smales (2013) documents that Australian interest rate futures incorporate macroeconomic announcements within seconds, with sharp but short-lived spikes in volatility, trading activity, and bid-ask spreads. The adjustment is essentially complete by the time the first minute of post-announcement trading has elapsed. This rapid incorporation reflects the microstructure of deep, liquid, centrally cleared futures markets with dedicated market makers, tight spreads, and continuous two-sided quoting.

The GB intraday electricity market operates in a fundamentally different environment. There are no designated market makers. The order book is continuous but thin, particularly outside the final hour before delivery. The average trade size has fallen to 1.31 MWh, and even in 2025, a 3-minute window around a random timestamp captures an average of only 96 trades. Individual transactions often reflect asset-based position management, where a generator or supplier adjusts their physical exposure, rather than speculative information-driven trading. A wind farm operator selling intraday power to cover a production shortfall does not respond to REMIT publications or IRIS data in the same way that a proprietary futures trader responds to a CPI release. The majority of volume in the GB intraday market is driven by participants managing physical delivery obligations, not by speculative traders seeking to profit from information asymmetries. This means that the price discovery process relies on a relatively small subset of informationally motivated participants operating within a market dominated by flow-driven activity.

These structural differences explain why we do not observe the clean, instantaneous price jumps that characterise financial market event studies. Instead, we observe three distinct adjustment patterns that vary systematically with the nature of the information signal.

Imbalance prices produce the fastest response, with evidence of initial overreaction. The ΔIMB price coefficient is largest at the 10-second horizon and declines at longer horizons, indicating that the first trades after an IRIS publication overshoot and are partially corrected by subsequent activity. This pattern is consistent with a small number of algorithmic participants parsing the standardised, regularly timed IRIS publication and trading on it within seconds, pushing prices beyond the informationally efficient level before the broader market catches up. The level specification reinforces this interpretation: the IMB price level coefficient becomes insignificant by 300 seconds, confirming that the absolute level is a transient signal rather than a persistent driver. The overreaction-then-correction pattern is the closest analogue to the Smales (2013) benchmark in our data, but even here the adjustment unfolds over minutes rather than seconds.

REMIT outages produce the slowest response, with evidence of persistent underreaction. The regression coefficient increases monotonically from 10 seconds to 600 seconds, and the impulse response shows continuous drift over the full 5-minute post-event window. Under the semi-strong efficiency benchmark, all publicly available information should be reflected in a single price jump at the moment of publication with no predictable subsequent drift. The persistent drift we observe constitutes underreaction. The explanation lies in the complexity of the signal. A REMIT UMM is not a single number: it requires the participant to identify the affected asset, assess the MW impact, determine which settlement periods fall within the outage window, evaluate the implications for system tightness given current conditions, and decide whether and how to adjust their position. In a market without dedicated market makers, this multi-step assessment translates into sequential trades over several minutes rather than a single instantaneous price adjustment. The practical implication is that participants who can automate the parsing of BMRS REMIT feeds and pre-compute price impact estimates may find short-lived opportunities in the minutes following an outage announcement.

Renewable forecast revisions fall between these extremes. Solar significance emerges at the 60-second horizon and persists through 300 seconds, indicating an adjustment speed that is slower than the imbalance price response but faster than the REMIT drift. Wind, as discussed in Section 5.1.1, barely registers at all, reflecting the limited residual information content of the public NESO forecast rather than slow market processing. The intermediate speed for solar is consistent with the nature of the signal: a forecast revision is a single number with clear directional implications (more solar = lower residual load = lower price), but translating it into a specific £/MWh adjustment requires some assessment of its magnitude relative to total generation and current market conditions.

This hierarchy, fast for simple numerical signals, slow for complex asset-specific events, intermediate for standardised forecast revisions, follows the same qualitative logic documented in financial markets: information processing speed depends on signal complexity, the number of participants monitoring the source, and the liquidity available to trade on it. The GB intraday market processes information at a fundamentally slower pace than the futures markets studied by Smales (2013), but the ordering across signal types is consistent with theoretical expectations.

One important caveat applies to all our impulse response results. As noted in Section 4.4.3, the imbalance price and volume IRFs are constructed from events where the next tradeable contract is within 5 minutes of gate closure. The later bins in these IRFs are populated by very few events where trading continued long enough to fill the full post-event window. The large price movements visible at the tail end may therefore reflect thin liquidity and elevated volatility near gate closure rather than a genuine sustained price response. This caveat does not affect the regression-based horizon analysis, which uses all events in the sample and is not subject to the same survivorship issue, but it does mean that the visual IRF patterns for imbalance variables should be interpreted with more caution than those for REMIT and forecast revisions, where the post-event trading window is longer.

5.3 Market Maturation and Consumer Welfare

A recurring pattern across all three parts of the analysis is that the responsiveness of intraday prices to fundamental information has strengthened over the sample period, particularly in 2024 and 2025. Solar forecast revisions are significant only in 2024. The derated margin becomes significant in 2024 and 2025 after being insignificant in all earlier years. The ΔIMB price coefficient roughly quintuples from 2024 to 2025. These temporal patterns coincide with the tenfold increase in monthly trade counts documented in Section 4.1, from approximately 300,000 in early 2021 to over 3 million by mid-2025, alongside a halving of average trade size from over 6 MWh to approximately 1 MWh, consistent with growing algorithmic participation.

This co-evolution of liquidity and informational responsiveness is not coincidental. A thicker order book means that when a NESO forecast revision, a REMIT outage, or an IRIS imbalance publication arrives, there are more participants present in the market who are capable of interpreting the signal and willing to trade on it. Each additional informed participant who submits a buy or sell order in response to new information pushes the VWAP closer to the informationally efficient level. In 2021, with 300,000 trades per month, the number of actively trading participants responding to any given information event was small, and a single forecast revision might produce no detectable VWAP movement simply because too few counterparties were present. By 2025, the same revision is observed and acted upon by a larger and more diverse set of participants, producing a measurable and statistically significant price response.

The declining magnitude of the REMIT coefficient from 2021 to 2024 provides complementary evidence from the opposite direction. While forecast and imbalance coefficients have strengthened (the market has become more responsive to these signals), the REMIT price impact per MW has fallen. This is not a contradiction. It reflects the standard microstructure prediction that deeper liquidity reduces the price impact of order flow. When an outage announcement triggers buying pressure from affected BRPs seeking to close their positions, a thicker order book absorbs this flow with a smaller price displacement. The same MW outage that moved the market by £3.02/MWh per MW in 2021 moved it by only £0.57/MWh per MW in 2024. The market is not less responsive to outages in any fundamental sense; rather, it is better able to absorb them without extreme price dislocations.

This evolution has direct consequences for consumer welfare, connecting to the central argument developed in the introduction. We established, drawing on the EU electricity balancing framework (Commission Regulation (EU) 2017/2195), that the cost of procuring balancing services is ultimately transferred to end consumers through two channels: imbalance charges levied on BRPs, which suppliers pass through in retail tariffs, and balancing capacity costs socialised in network tariffs. The higher the residual system imbalance that NESO must resolve through the Balancing Mechanism, the more expensive the balancing actions become, and the greater the cost that appears on every consumer's electricity bill.

Koch and Hirth (2019) provide the most direct empirical evidence of the mechanism linking intraday trading to balancing costs. Despite wind and solar generation nearly doubling in

Germany between 2011 and 2017, both balancing reserve requirements and reserve activation volumes declined by approximately 50%. They attribute much of this to increased intraday trading activity, finding that the shift to quarter-hourly trading products alone reduced balancing energy volumes by 17%. The mechanism is the one described in our introduction: when a BRP recognises from updated forecasts or system data that it will be short at delivery, purchasing power on the intraday market before gate closure is expensive relative to the day-ahead price, but cheaper than the alternative of facing the imbalance settlement price, which during tight conditions can reach multiples of the intraday price. Each intraday trade that closes an open position removes one megawatt of imbalance that NESO would otherwise have to resolve by activating balancing reserves.

Our results suggest that this self-correcting mechanism is becoming increasingly effective in Great Britain. The growing responsiveness of intraday prices to fundamental signals means that participants are reacting to forecast revisions, system tightness indicators, and imbalance publications with greater speed and accuracy. Positions are corrected more efficiently before delivery, shrinking the residual imbalance volume. The declining REMIT price impact confirms that the market absorbs supply shocks with smaller dislocations, reducing the frequency and severity of extreme price events.

The contrast with markets where this mechanism is underdeveloped illustrates the stakes. We noted in the introduction that during the summer of 2025, the Baltic balancing markets recorded imbalance settlement prices reaching +€9,976/MWh and -€10,007/MWh (Baltic Transparency Dashboard, 2026). These extreme prices reflect a system where the intraday market was too thin for BRPs to correct their positions before delivery, leaving large residual imbalances for the TSO to resolve at extraordinary cost. Under the EU balancing framework, these costs do not remain with the BRPs who incurred them: they flow through to all consumers via imbalance charges and network tariffs. As the ACER-CEER (2024) Market Monitoring Report notes, without significant changes to demand patterns and market participation, prices will ultimately increase for all consumers regardless of contract type. In less competitive markets dominated by incumbent suppliers with limited switching, the pass-through is even more direct.

The GB market's trajectory from 2021 to 2025 represents the opposite of this scenario. The tenfold growth in trading activity, the increasing responsiveness to publicly available fundamental data, and the declining price impact of supply shocks together indicate a market that is progressively better at performing its intended function: allowing participants to trade out their forecast errors before delivery, compressing the residual imbalance, and reducing the total system cost that must ultimately be recovered from consumers.

However, this positive trajectory comes with a qualification that follows directly from Section 5.1.1. The growing efficiency of the GB intraday market appears to be driven substantially by algorithmic and proprietary-data-equipped participants. The insignificance of the public NESO wind forecast in our regressions, combined with the significance of commercially sourced forecast data documented in the German literature (Kiesel & Paraschiv, 2017; Kremer et al., 2021), suggests that much of the informational edge in intraday trading resides with participants who can afford commercial forecast subscriptions, automated trading infrastructure, and the technical expertise to process high-frequency data feeds. If market efficiency depends on the participation of these sophisticated actors, then the barrier to entry for smaller utilities, independent renewable producers, and new trading entrants becomes a concern for long-run market health. A market that is efficient because a small number of well-resourced participants dominate price discovery may be fragile: it depends on those participants continuing to trade, and it may not deliver the broad competitive discipline that Batalla-Bejerano et al. (2022) identify as necessary for translating wholesale efficiency into lower retail prices. The introduction noted that the combined cost of data subscriptions, forecasting tools, and trading infrastructure represents a meaningful barrier to participation, particularly in smaller markets. Our results suggest that this barrier is not merely hypothetical but has measurable consequences for which information signals are priced and which are not.

This tension between deepening efficiency and narrowing participation is not unique to electricity markets. It parallels the evolution of equity markets following the growth of high-frequency trading, where market quality metrics (spreads, volatility, price efficiency) improved while concerns about access, fairness, and fragility intensified. For GB electricity market policy, the implication is that initiatives to improve public data quality, increase forecast publication frequency, and reduce the technical barriers to accessing BMRS and NESO APIs may be as important for long-run market health as the structural reforms being considered under REMA. A market where the public data infrastructure is good enough that smaller participants can compete on at least some dimensions, even without proprietary forecasts, is likely to be more resilient and more competitive than one where efficiency depends entirely on a handful of algorithmic traders.

5.4 Implications for Transparency and Market Design

The introduction argued that Great Britain offers a structurally different transparency regime from continental Europe. While the EU mandates fundamental data publication through the ENTSO-E Transparency Platform under Regulation (EU) No 543/2013, the platform suffers from persistent quality issues. ACER (2025) continues to recommend improvements, noting inconsistent data definitions across legal frameworks and insufficient clarity in publications related to balancing and cross-zonal capacities. Mascarenhas and Kazmi (2025) document concrete instances of implausible values, including generation figures exceeding installed capacity, alongside significant performance differences in day-ahead forecasts across bidding zones. In this environment, commercial data providers such as Volue Insight, Montel/Energy Quantified, and Dexter Energy fill the gap by layering proprietary weather models and curated data feeds on top of the freely available but imperfect fundamentals, at a cost that represents a meaningful barrier to entry for smaller participants.

The GB system operates differently. NESO and Elexon publish wind, solar, and demand forecasts, REMIT outage data, imbalance settlement prices and volumes, LoLP, and derated margins freely and at high frequency through open APIs, maintained by a single TSO and a single settlement body within a coherent, well-documented data ecosystem. Our results demonstrate that this data is not merely published but actively used by market participants for price discovery. Every major public data stream we examine produces a statistically significant price response, at least within the last hour before delivery: solar forecast revisions, derated margin updates, REMIT outage announcements, and imbalance price and volume publications all move intraday prices in the directions predicted by economic theory.

This finding validates the design choice of comprehensive, free, high-frequency public data provision. It also provides empirical support for the direction of ACER's (2025) recommendations: improving the quality and consistency of publicly available fundamental data is not an abstract governance objective but a prerequisite for efficient price discovery in intraday markets.

However, our results also reveal the limits of what public data alone can achieve. The insignificance of the NESO wind forecast in our regressions, combined with the strong significance of proprietary forecasts documented by Kiesel and Paraschiv (2017) and Kremer et al. (2021) in Germany, demonstrates that the public TSO forecast is not the binding constraint on wind-related price discovery. By the time NESO publishes its hourly wind update, participants with commercial subscriptions have already acted on higher-frequency private forecasts. The public forecast retains marginal information value for solar, particularly close to gate closure, but for wind it is largely redundant for the informationally active participants who drive price discovery. This creates a two-tier market: sophisticated participants trade on proprietary data, while smaller participants relying solely on public sources face a systematic informational disadvantage, at least for wind-related fundamentals.

The practical implication is that the value of public data provision lies not primarily in giving active traders information they do not already possess. Sophisticated participants will always maintain an informational edge through commercial services. Rather, the public data ecosystem serves three functions that our results help to quantify. First, it establishes a transparent information baseline that enables regulatory monitoring. Our impulse response functions, which show no systematic pre-event price drift before REMIT publications, provide evidence that the disclosure framework is functioning as intended. Deviations from the established response patterns, such as abnormal pre-announcement price movements or unusually muted responses to large forecast revisions, could serve as indicators of market dysfunction or manipulation for REMA monitoring teams. Second, it provides the minimum information set that enables participation by less resourced market participants, even if their informational position is weaker than that of proprietary-data-equipped traders. The significance of solar and DM in the public data confirms that at least some dimensions of the market can be navigated using freely available information. Third, it provides the raw inputs for the commercial forecasting layer. The proprietary forecasts from Volue, Meteorologica, and others are built on top of TSO and ENTSO-E data; without reliable public fundamentals, the commercial layer could not exist.

For REMA specifically, our results establish empirical baselines for how GB intraday prices respond to each category of public information under normal market conditions. We document the typical coefficient magnitudes, the speed of adjustment across different horizons, and the shape of the impulse response for each data stream. These baselines have operational value. If a period of abnormally weak price response to REMIT outages were detected, it could indicate reduced participation, degraded data quality in the BMRS feed, or strategic withholding of trading activity. If the pre-event impulse response for a specific asset class began showing systematic drift before the publication timestamp, it would warrant investigation into potential information leakage. The flat pre-event baselines we establish provide the reference against which such anomalies can be measured.

A further implication concerns data publication frequency. Our horizon analysis reveals that different data streams serve different functions in the price discovery process depending on their timing characteristics. Imbalance prices, which arrive approximately every 30 minutes at a predictable timestamp, produce the fastest price response and the strongest statistical significance. Forecast revisions, published hourly, produce a slower and weaker response. REMIT messages, which arrive irregularly and require multi-step interpretation, produce the slowest adjustment. This ordering suggests that increasing the publication frequency of currently hourly data streams, such as the NESO wind and solar forecasts, could improve price discovery by providing more frequent and timelier information to the market. If NESO were to publish forecasts every 15 or 30 minutes rather than hourly, the gap between the public signal and the private proprietary forecasts would narrow, potentially reducing the informational disadvantage faced by participants without commercial subscriptions. Whether the cost of more frequent forecast computation is justified by the welfare gain from improved price discovery is an empirical question that our results can help frame but cannot answer directly.

5.5 Limitations

Several limitations of our analysis should be acknowledged. Our regressions assume a constant linear marginal effect per MW of forecast revision or outage. In practice, the price sensitivity to a given shock depends on the slope of the merit order curve at the prevailing operating point: a 500 MW outage during a period when the system is operating on the steep portion of the supply stack should produce a larger price response than the same outage during a flat, uncongested period (Sensfuß et al., 2008). Kremer et al. (2021) address this with a threshold regression that allows coefficients to switch between flat and steep merit order regimes, estimated using EPEX's published day-ahead supply curves as a proxy. Hirsch and Ziel (2024) compute the slope directly from the anonymous day-ahead bid and offer curves, under the assumption that the supply stack does not change significantly between the day-ahead and intraday timeframes. An alternative approach would use live intraday order book snapshots at the moment of each event, estimating the actual depth of liquidity available to absorb the shock, though this would need to account for the presence of iceberg orders. We do not pursue any of these approaches, as the GB EPEX aggregated day-ahead curves require additional processing and validation that falls outside our scope, and we do not have access to live order book snapshot data. Incorporating state-dependent price impacts is an important avenue for future research and would allow estimation of how the same information shock produces different price responses under different system conditions.

We use NESO's publicly available hourly forecasts rather than the proprietary trading-grade forecasts used by Kiesel and Paraschiv (2017) and Kremer et al. (2021), who employed a unique dataset from EWE TRADING GmbH. This means our wind and solar coefficients measure the residual information content of the public TSO forecast after the market has already partially incorporated private signals. Our estimates represent a lower bound on the total influence of renewable forecast information on intraday prices. A study using proprietary GB forecasts from a commercial provider would likely find stronger, more pervasive, and earlier-emerging effects, particularly for wind.

Our analysis covers only the EPEX continuous trading platform. Bilateral OTC trades, activity on the N2EX (Nord Pool) platform, and the outcomes of the IDA1 and IDA2 intraday auctions are not captured. EPEX continuous trading represents a substantial and growing share of GB intraday activity, but omitting other venues means our results reflect price discovery on one platform rather than the aggregate market. If informed trading is concentrated on EPEX while OTC activity is more flow-driven, our estimates may overstate the efficiency of the aggregate market. Conversely, if significant information-driven trading occurs on N2EX or in bilateral channels, our estimates may understate the total price response to fundamental information.

The period from 2021 to 2025 coincides with extraordinary volatility in European energy markets, driven by the post-COVID recovery, the energy crisis triggered by the Russian invasion of Ukraine, and the subsequent normalisation. The coefficient magnitudes we estimate, particularly the large REMIT coefficients in 2021-2022, may not generalise to calmer market conditions. The declining REMIT coefficient over the sample period already suggests that the market's price sensitivity is evolving, and our estimates should be understood as conditional on the specific market environment of this period rather than as permanent structural parameters.

The LOLP coefficient carries a counterintuitive negative sign in the joint regression with DM. As discussed in Section 4.2.2, this reflects the extreme skewness of the LOLP distribution: the mean is 0.0002, the variable is effectively zero for over 99% of observations, and it only takes non-trivial values during rare scarcity episodes. A linear regression coefficient has no reliable marginal interpretation for a variable with this distribution. The derated margin, which is continuously distributed and mechanically related to LOLP, captures the economically relevant tightness signal. Future work could address this by estimating a non-linear specification, such as a threshold model that activates a separate LOLP coefficient only when the variable exceeds a critical level, or by using the Reserve Scarcity Price (RSP = LOLP × VoLL) directly as the regressor, which would translate the probability into a £/MWh scarcity premium with a more interpretable coefficient.

Section 5.3 draws on Koch and Hirth (2019) and the EU balancing framework to argue that increased intraday trading activity reduces balancing costs and benefits consumers. While this argument is supported by our evidence of growing market responsiveness, we do not directly estimate the causal effect of intraday trading on balancing volumes or consumer prices in the GB market. Establishing this causal link would require a different empirical strategy, potentially exploiting exogenous variation in intraday market access or liquidity, and is beyond the scope of this thesis.

6. Conclusion

This thesis examines the determinants and informational efficiency of intraday electricity prices in Great Britain using 70.2 million transactions from the EPEX continuous trading platform (2021-2025). Adapting event study methods from financial markets, we measure the price response to four categories of publicly available information: NESO wind and solar forecast revisions, derated margin and Loss of Load Probability updates, REMIT generation outage announcements, and indicative imbalance settlement prices and volumes. All four information types significantly affect intraday prices, but the public NESO wind forecast does not reach significance in the pooled sample, which we attribute to proprietary forecasts pre-empting the public signal. Across all variables, price sensitivity concentrates within the final hour before gate closure, confirming that the approaching imbalance settlement amplifies the value of information near delivery. Incorporation speed varies by signal type: imbalance prices are absorbed within seconds with evidence of overreaction, REMIT outages show persistent underreaction over 10 minutes, and solar revisions fall between. The market has become measurably more responsive over the sample period, coinciding with a tenfold increase in trading activity and consistent with growing algorithmic participation. We connect these findings to consumer welfare: following Koch and Hirth (2019), more responsive intraday trading reduces residual imbalances and the balancing costs recovered from consumers. However, the growing efficiency appears driven by well-resourced participants with proprietary data, raising concerns about participation breadth. Our results validate GB's public data provision regime, establish empirical baselines for REMIT compliance monitoring, and demonstrate what new entrants need to know to participate in GB intraday trading.

7. References

ACER. (2025). Opinion No 08/2025 on the fourth update of the manual of procedures for ENTSO-E's central information transparency platform. Agency for the Cooperation of Energy Regulators.

https://www.acer.europa.eu/sites/default/files/documents/Official_documents/Acts_of_the_Agency/Opinions/Opinions/ACER-Opinion-08-2025-ENTSOE-Transparency-Platform.pdf

ACER & CEER. (2024). 2024 Market Monitoring Report on energy retail and consumer protection. Agency for the Cooperation of Energy Regulators. https://www.ceer.eu/wp-content/uploads/2024/09/ACER-CEER_2024_MMR_Retail-1.pdf

Aïd, R., Cosso, A., & Pham, H. (2022). Equilibrium price in intraday electricity markets. Mathematical Finance, 32(2), 517–554. https://doi.org/10.1111/mafi.12340

Baltic Transparency Dashboard. (2026). Imbalance settlement prices. https://baltic.transparency-dashboard.eu/node/41

Batalla-Bejerano, J., Trujillo-Baute, E., & Villa-Arrieta, M. (2022). Is more competition better? Retail electricity prices and switching rates in the European Union. Utilities Policy, 79, 101436. https://www.sciencedirect.com/science/article/abs/pii/S0301421520303475

CEPA. (2021). Cost benefit analysis of multi-region loose volume coupling (MRLVC): Summary report. Retrieved from https://consultations.entsoe.eu/markets/cost-benefit-analysis-of-multi-region-loose-volume/supporting_documents/MRLVC_CBA_analytical%20results_April_2021_final_publication.pdf

Cludius, J., Hermann, H., Matthes, F. C., & Graichen, V. (2014). The merit order effect of wind and photovoltaic electricity generation in Germany 2008–2016: Estimation and distributional implications. Energy Economics, 44, 302–313. Retrieved from https://www.sciencedirect.com/science/article/abs/pii/S0140988314001042

Commission Regulation (EU) 2017/2195 of 23 November 2017 establishing a guideline on electricity balancing. Official Journal of the European Union, L 312, 6–53. https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32017R2195

Department for Business, Energy & Industrial Strategy. (2021). Trading electricity with the EU. Retrieved from https://www.gov.uk/government/publications/trading-electricity-with-the-eu

Do, P., Jamasb, T., & Meier, H. (2021). Renewable energy intermittency and electricity market integration: The impact of wind and solar power on electricity prices. Applied Energy, 283, 116315. Retrieved from https://www.sciencedirect.com/science/article/pii/S0306261920316846

Eicke, A., Ruhnau, O., & Hirth, L. (2021). Electricity balancing as a market equilibrium: An instrument-based estimation of supply and demand for imbalance energy. Energy Economics, 104, 105640. https://doi.org/10.1016/j.eneco.2021.105640

Eliņa, L. (2022, March 16). Latvia compared to EU peers in electricity and gas companies. Latvian Public Broadcasting (LSM). https://eng.lsm.lv/article/economy/economy/latvia-compared-to-eu-peers-in-electricity-and-gas-companies.a448147/

Elexon. (2024, February 26). Indicative Settlement Price data now available on the Insights Solution. Retrieved from https://www.elexon.co.uk/bsc/article/indicative-settlement-price-data-now-available-on-the-insights-solution

Elexon. (2025). Imbalance Pricing Guidance (Version 16.0; effective from 14 July 2025). Retrieved from https://www.elexon.co.uk/bsc/documents/training-guidance/bsc-guidance-notes/imbalance-pricing

Elexon. (n.d.-a). BSC data (Insights Solution). Retrieved from https://www.elexon.co.uk/bsc/data

Elexon. (n.d.-b). Detailed system prices (Insights Solution). Retrieved from https://bmrs.elexon.co.uk/detailed-system-prices

Elexon. (n.d.-c). System sell buy prices (Insights Solution). Retrieved from https://bmrs.elexon.co.uk/system-prices

Enerdata. (2024). Households energy expenditures. Odyssee-Mure. https://www.odyssee-mure.eu/publications/efficiency-by-sector/households/energy-expenditures-of-households.html

ENTSO-E. (n.d.). Single day-ahead coupling (SDAC). Retrieved from https://www.entsoe.eu/network_codes/cacm/implementation/sdac/

EPEX SPOT SE. (2020). Single Intraday Coupling: Making borders flexible. Retrieved from https://www.epexspot.com/sites/default/files/sites/annual-report-2020/innovation/single-intraday-coupling-making-borders-flexible/

EPEX SPOT SE. (n.d.-a). Focus GB. Retrieved from https://www.epexspot.com/en/focus-gb

EPEX SPOT SE. (n.d.-b). Trading products. Retrieved from https://www.epexspot.com/en/tradingproducts

European Commission. (2015). Commission Regulation (EU) 2015/1222 establishing a guideline on capacity allocation and congestion management. Retrieved from https://eur-lex.europa.eu/eli/reg/2015/1222/oj

European Parliament, & Council of the European Union. (2011). Regulation (EU) No 1227/2011 of the European Parliament and of the Council of 25 October 2011 on wholesale energy market integrity and transparency (Text with EEA relevance). Official Journal of the European Union. Retrieved from https://eur-lex.europa.eu/eli/reg/2011/1227/oj/eng

Eurostat. (2022). Household budget survey: Statistics on consumption expenditure. European Commission. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Household_budget_survey_-_statistics_on_consumption_expenditure

Figueiredo, N. C., & da Silva, P. P. (2019). The impact of wind and solar generation on the Portuguese electricity market. Renewable and Sustainable Energy Reviews, 109, 131–142. Retrieved from https://www.sciencedirect.com/science/article/abs/pii/S1364032118307950

Grégoire, V., & Martineau, C. (2022). How is earnings news transmitted to stock prices? Journal of Accounting Research, 60, 261–297. https://doi.org/10.1111/1475-679X.12394

Gurtler, M., & Paulsen, T. (2018). The effect of wind and solar power forecasts on day-ahead and intraday electricity prices in Germany. Energy Economics, 75, 150–162. Retrieved from https://www.sciencedirect.com/science/article/pii/S0140988318302512?via%3Dihub

Hagemann, S. (2015). Price determinants in the German intraday market for electricity: An empirical analysis. Journal of Energy Markets, 8(3), 21–45. https://doi.org/10.21314/JEM.2015.128

Hirsch, S., & Ziel, F. (2024). Simulation-based forecasting for intraday power markets: Modelling fundamental drivers for location, shape and scale of the price distribution. The Energy Journal, 45(3), 55–82. https://doi.org/10.5547/01956574.45.3.shir

Hu, X., Jaraitė, J., & Kažukauskas, A. (2021). The effects of wind power on electricity markets: A case study of the Swedish intraday market. Energy Economics, 96, 105159. https://doi.org/10.1016/j.eneco.2021.105159

Ito, K., & Reguant, M. (2022). Deregulation, market power, and prices: Evidence from the electricity sector (Working Paper No. 2022-008). MIT Center for Energy and Environmental Policy Research. https://ceepr.mit.edu/wp-content/uploads/2022/04/2022-008.pdf

Karanfil, F., & Li, Y. (2017). The role of continuous intraday electricity markets: The integration of large-share wind power generation in Denmark. The Energy Journal, 38(2), 107–130. https://doi.org/10.5547/01956574.38.2.fkar

Ketterer, J. C. (2014). The impact of wind power generation on the electricity price in Germany. Energy Economics, 44, 270–280. Retrieved from https://www.sciencedirect.com/science/article/abs/pii/S0140988314000875

Kiesel, R., & Paraschiv, F. (2017). Econometric analysis of 15-minute intraday electricity prices. Energy Economics, 64, 77–90. https://doi.org/10.1016/j.eneco.2017.03.002

Koch, C. (2021). Intraday imbalance optimization: Incentives and impact of strategic intraday bidding behavior. Energy Systems, 13, 409–435. https://doi.org/10.1007/s12667-021-00445-9

Koch, C., & Hirth, L. (2019). Short-term electricity trading for system balancing: An empirical analysis of the role of intraday trading in balancing Germany's electricity system. Renewable and Sustainable Energy Reviews, 113, 109275. https://doi.org/10.1016/j.rser.2019.109275

Kremer, M., Kiesel, R., & Paraschiv, F. (2021). An econometric model for intraday electricity trading. Philosophical Transactions of the Royal Society A, 379(2202), 20190624. https://doi.org/10.1098/rsta.2019.0624

Latvenergo. (2025). Half of the electricity sold by Latvenergo Group in Lithuania and Estonia [Press release]. https://latvenergo.lv/en/jaunumi/preses-relizes/relize/half-electricity-sold-latvenergo-group-lithuania-and-estonia

Mascarenhas, M. M., & Kazmi, H. (2025). Quis custodiet ipsos custodes: Bringing more transparency to the ENTSO-E TP data and forecasts on transmission system operations. In 2025 21st International Conference on the European Energy Market (EEM) (pp. 1–6). IEEE. https://doi.org/10.1109/EEM64765.2025.11050146

NEMO Committee. (2020, November 6). Information note about the exiting GB parties in SDAC [PDF]. Retrieved from https://www.nemo-committee.eu/assets/files/information-note-about-the-exiting-gb-parties-in-sdac.pdf

Nord Pool. (2022). The N2EX Day-Ahead Auction in the UK. Retrieved from https://support.nordpoolgroup.com/support/solutions/articles/8000088463-about-the-n2ex-day-ahead-auction

Office of Gas and Electricity Markets. (2014a). Inside information disclosure under the Regulation on wholesale energy market integrity and transparency (REMIT) [Open letter]. Retrieved from https://www.ofgem.gov.uk/sites/default/files/docs/2014/07/ofgem_open_letter_on_remit_inside_information_0.pdf

Office of Gas and Electricity Markets. (2014b). Electricity Balancing Significant Code Review: Final policy decision. Office of Gas and Electricity Markets. https://www.ofgem.gov.uk/publications/electricity-balancing-significant-code-review-final-policy-decision

Regulation (EU) No 543/2013 of 14 June 2013 on submission and publication of data in electricity markets. Official Journal of the European Union, L 163, 1–12. https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32013R0543

Sensfuß, F., Ragwitz, M., & Genoese, M. (2008). The merit-order effect: A detailed analysis of the price effect of renewable electricity generation on spot market prices in Germany. Energy Policy, 36(8), 3086–3094. https://doi.org/10.1016/j.enpol.2008.03.035

Smales, L. A. (2013). Impact of macroeconomic announcements on interest rate futures: High-frequency evidence from Australia. Journal of Financial Research, 36(3), 371–388. https://doi.org/10.1111/j.1475-6803.2013.12015.x

Thakare, S., Bokde, N. D., & Feijóo-Lorenzo, A. E. (2023). Forecasting different dimensions of liquidity in the intraday electricity markets: A review. AIMS Energy, 11(5), 918–959. https://doi.org/10.3934/energy.2023044

Trebbien, S. (2023). The impact of renewable forecast errors on cross-border electricity prices and flows. arXiv preprint arXiv:2310.03339. Retrieved from https://arxiv.org/pdf/2310.03339

Valitov, N., & Maier, A. (2020). Asymmetric information in the German intraday electricity market. Energy Economics, 89, 104785. Retrieved from https://www.sciencedirect.com/science/article/pii/S0140988320301250

Weron, R. (2014). Electricity price forecasting: A review of the state-of-the-art with a look into the future. International Journal of Forecasting, 30(4), 1030–1081. Retrieved from https://www.sciencedirect.com/science/article/pii/S0169207014001083

Ziel, F. (2017). Modeling the impact of wind and solar power forecasting errors on intraday electricity prices. In 2017 14th International Conference on the European Energy Market (EEM) (pp. 1–5). IEEE. https://doi.org/10.1109/EEM.2017.7981900

Acknowledgement

We would like to thank our supervisor, Olegs Matvejevs, for his guidance throughout the research process.

Use of Artificial Intelligence Tools

AI-based tools, specifically Anthropic's Claude, were used in three capacities during the preparation of this thesis. First, Claude assisted with writing and debugging SQL and Python code for data extraction, cleaning, and transformation, particularly for querying the Elexon BMRS API, processing REMIT Urgent Market Messages, and constructing the event-matched datasets used in the empirical analysis. Second, Claude was used as a research aid to locate, summarise, and cross-reference academic literature, regulatory documents, and technical specifications related to GB electricity market design, the EU balancing framework, and intraday price formation. Third, Claude assisted with structuring and drafting sections of the thesis text, including the discussion of regression outputs and the interpretation of findings in relation to existing literature. In all cases, the authors independently verified the accuracy of code outputs, confirmed that all referenced sources exist and are correctly cited, and exercised full editorial judgement over the final text. All empirical results, analytical conclusions, and economic interpretations are the authors' own.