

Intraday Power Trading: Towards an Arms Race in Weather Forecasting?

Thomas Kuppelwieser · David Wozabal

Received: date / Accepted: date

Abstract We propose the first weather-based algorithmic trading strategy on a continuous intraday power market. The strategy uses neither production assets nor power demand and generates profits purely based on superior information about aggregate output of weather-dependent renewable production. We use an optimized parametric policy based on state-of-the-art intraday updates of renewable production forecasts and evaluate the resulting decisions out-of-sample for one year of trading based on detailed order book level data for the German market. Our strategies yield significant positive profits, which suggests that intraday power markets are not semi-strong efficient. Furthermore, sizable additional profits could be made using improved weather forecasts, which implies that the quality of forecasts is an important factor for profitable trading strategies. This has the potential to trigger an arms race for more frequent and more accurate forecasts, which would likely lead to increased market efficiency, more reliable price signals, and more liquidity.

Keywords policy optimization · intraday power markets · algorithmic trading · weather based trading · stochastic optimization

1 Introduction

In the last decades, the electricity industry in many countries has seen rapid changes. One driver of these developments was the transition from a highly vertically integrated, state controlled sector of the economy to a largely competitive and decoupled industry [Pollitt \(2019\)](#). Another reason is the climate crisis and the increasing efforts to transition to a carbon neutral society. The electricity sector is the key to sustainable energy systems changing the nature of energy supply by sharply increasing production from *variable renewable energy sources* (VRES) such as wind and photovoltaics.

Thomas Kuppelwieser

School of Management, Technical University of Munich, Arcisstraße 21, 80333 Munich, Germany

E-mail: thomas.kuppelwieser@tum.de

David Wozabal

School of Management, Technical University of Munich, Arcisstraße 21, 80333 Munich, Germany

E-mail: david.wozabal@tum.de

In the majority of industrialized countries electricity is traded on a range of future markets whose products differ in their time to maturity. Recently, the weather-dependent and unpredictable nature of VRES production has increasingly shifted the focus to markets with a high temporal resolution that trade close to delivery when production forecasts are reasonably accurate.

Short-term trading is mostly organized in real-time markets or continuous intraday markets. While the former is the prevailing design in the US (Milligan et al., 2016), the latter is, for example, used in Europe. These volatile markets are attractive for firms that can quickly adapt their demand or production profiles and can thus sell their *flexibility* to other market participants with balancing needs driven by, for example, forecast errors in VRES production. Short-term trading thus provides incentives to invest in flexible energy sources such as gas turbines and storage, which are required to balance the intermittent production from ever growing VRES capacities.

Apart from flexibility providers, short-term markets are increasingly interesting for speculative traders who neither own production assets nor trade their own electricity demand. In this paper, we propose a trading strategy for speculative trading on continuous intraday markets. Our approach is motivated by algorithmic trading strategies in continuous financial markets that are triggered by *signals* indicating a change in the fundamental value of an asset. Since, as discussed above, VRES production is an important driver of short-term electricity trading, we use forecast errors of aggregate VRES production as signals for our strategies. The rationale for this choice is that if forecasts for VRES production are inaccurate, producers have to correct their positions taken on the day-ahead market, which, if the errors are large enough, causes a shift in intraday prices (Kiesel and Paraschiv, 2017; Kremer et al., 2020a,b).

While the literature on asset backed trading on intraday power markets is extensive (see for example Boomsma et al., 2014; Kumbartzky et al., 2017; Séguin et al., 2017; Bertrand and Papavasiliou, 2019; Wozabal and Rameseder, 2020; Rintamäki et al., 2020), there is virtually no research on optimal bidding strategies for speculative traders that have no assets of their own.

In the following, we review those papers that come closest to our trading strategies. Kath and Ziel (2018) introduce a forecast for the volume weighted continuous intraday price for 15-minutes contracts and develop a strategy to choose between trading on the day-ahead auction market and the continuous intraday market. Monteiro et al. (2020) evaluate future trading strategies on the Spanish Mibel market based on long-term electricity futures. Maciejowska et al. (2019) study the problem of a small VRES producer that trades on the day-ahead and the intraday market. Wozabal and Rameseder (2020) study trading strategies for a storage that arbitrages between Spanish day-ahead and intraday markets. Furthermore, Kath and Ziel (2020) explore optimal order execution strategies with the aim to minimize liquidity cost and Glas et al. (2019, 2020) explore optimal VRES trading strategies on the intraday market in an optimal control setting. Bertrand and Papavasiliou (2019) use reinforcement learning to optimize a Markovian strategy for an electricity storage on the German intraday market for power.

We contribute to the literature in the following ways:

1. While there is a growing literature investigating the impact of VRES production forecast errors on intraday prices (e.g., Garnier and Madlener, 2014; Kiesel and Paraschiv, 2017; Kremer et al., 2020a,b; Kulakov and Ziel, 2019), we are the first to propose a demonstrably profitable trading strategy based on this observation. We take great care to accurately model market mechanisms, the exact clearing algorithm, and the sequence of information. To the best of our knowledge Martin and Otterson (2018);

Bertrand and Papavasiliou (2019); Kuppelwieser and Wozabal (2020) are the only other papers that capture the realities of continuous trading in similar detail. In particular, apart from Bertrand and Papavasiliou (2019), this is the first paper that evaluates a trading strategy based on detailed order book data, which is different from the extant literature that discretizes the trading to 1 minute or 15 minute brackets to be able to deal with the sheer amount of order data (e.g. Glas et al., 2019, 2020; Kath and Ziel, 2020).

The resulting trading problem is characterized by substantial uncertainties about the future state of the continuous market and a high frequency of arrival of new order information, necessitating a large number of decisions which have to be taken at random points in time. Consequently, given the complex information structure of the problem and the number of decisions to be taken, finding *optimal decisions* is clearly computationally intractable (Bertrand and Papavasiliou, 2019). We therefore propose a non-anticipative parametric policy that yields significant positive profits in controlled out-of-sample experiments and uses the forecast errors of renewable production as trading signals.

2. Our results show that intraday power markets are far from efficient. In particular, it is possible to capitalize on information on day-ahead forecast errors of VRES output. This fact suggests that the market disseminates information slowly and in an imperfect manner: While recent results found evidence that intraday electricity markets are weak-form efficient (e.g. Oksuz and Ugurlu, 2019; Narajewski and Ziel, 2020), our results illustrate that they violate the more restrictive semi-strong version of the efficient market hypothesis, which states that it is impossible to consistently generate abnormal returns using publicly available data Malkiel and Fama (1970).
3. Next to demonstrating that strategies based on current state-of-the-art weather forecasting are profitable, we quantify the value of a perfect weather forecast and conclude that there is potential for substantially increased profits from weather-based strategies. This finding suggests that in the future the industry might see an *arms race* in weather forecasting, similar to the arms race for speed observed in the financial markets (e.g., Budish et al., 2015).

In our numerical case study, we consider the German intraday power market. We first examine the insample performance of our policy for 18 months of trading to identify sensible ranges for our parameters and for the timing of trading decisions. We find a trade-off between the quality of the signal that is required to trigger the strategy and the size of the traded position. Generally speaking, profits per trade rise in the quality of the signal. However, if trading is restricted to only those products with high quality signals, trading occurs infrequently reducing overall profits. A similar trade-off can be observed for the size of the position: while profits initially rise with larger positions, the marginal profit per additional traded MWh is diminishing due to liquidity costs that increase in order size.

Furthermore, we find that one of the most important aspect of the trading strategy is how it deals with the lack of liquidity that plagues intraday power markets. In particular, a trader that seeks to capitalize on informational advantages in weather forecasting would ideally want to trade as early as possible on this information. However, since there is usually very little trading activity until 2-3 hours before gate closure, such a strategy is running the risk of being unprofitable due to high transaction costs. We show how *patient* strategies based on a sequence of limit orders can significantly reduce liquidity costs and outperform simpler impatient strategies based on market orders.

In an out-of-sample study, we evaluate our strategies for one year of trading. The results show that the proposed policies yield significant positive profits for both hourly and quarter-hourly products, where the former is characterized by larger volumes, higher profits, and more volatile profits per product, while the latter yields lower profits and also trades less volumes. This differences can mostly be explained by the higher liquidity of hourly products.

We show that the potential additional earnings for a strategy which is based on a perfect intraday forecast of VRES production are significant, increasing profits by one order of magnitude. Hence, there is a strong incentive to invest in better weather forecasts and more frequent updates during the day – a situation which has the potential to trigger an *arms race in short-term weather forecasting*. As opposed to the *arms race for speed* observed in the share market (e.g. [Budish et al., 2015](#)), this development has the potential to increase market liquidity in early hours of intraday trading, the accuracy of price discovery, and therefore ultimately welfare.

The rest of the paper is organized as follows: In Section 2, we describe the relevant features of intraday power markets and discuss liquidity and the impact of VRES. Section 3 is dedicated to our trading policy. Section 4 describes the setting of our case study, while Section 5 discusses its results. Finally, Section 6 concludes the paper and discusses implications as well as avenues for further research.

2 Intraday Markets

In this section, we first describe the typical market design of continuous intraday power markets in Section 2.1, focusing on the German continuous intraday market as one of the most liquid markets. Secondly, we discuss the influence of renewable generation on prices in Section 2.2. Finally, we investigate market liquidity and its dependency on time to delivery in Section 2.3.

2.1 Market Design

Most spot markets for power consist of a day-ahead market that allows market participants to trade electricity one day ahead of delivery and a short-term market, which gives participants the possibility to adjust their positions until shortly before physical delivery. Short-term markets are usually either organized as real-time markets or as intraday markets. Prominent examples for the former include most US power markets, while European short-term markets are examples of the latter category.

In Europe, there are currently two competing types of intraday trading systems: auction markets and continuous intraday trading. In 2015, the EU decided on the long-term goal to couple all European intraday markets in a large continuous market in order to facilitate a secure energy supply, competitiveness, and fair prices ([European Commission, 2015](#)). While most European countries already transitioned to continuous intraday markets that are compatible with the joint European design, some countries such as Italy, Spain, and Portugal still use auction markets. In this paper, we are interested in continuous intraday markets and for the ease of exposition focus on the European market design and its implementation in Germany hosted by the EPEX, the largest power exchange in Europe (see [Viehmann, 2017](#), for a detailed description). However, we note that other markets are very similar in the features crucial for the analysis in this paper.

With the build up of capacities in intermittent and unpredictable production, short-term trading on intraday markets is increasingly gaining traction (EPEX, 2020b). As a result, liquidity of the German intraday market has been improving in the last years with growing trading volumes, but also an increased prevalence of automated trading EPEX (2020b). In particular, due to the short-term nature of the continuous intraday market, marketing of flexible power sources and electricity storage as well as position closing is often relegated to trading algorithms.

On the German intraday market power can be traded on a national market until 30 minutes before physical delivery and until 5 minutes before physical delivery within the four control areas. The market opens shortly after the clearing of the day-ahead market and allows to trade hourly, half-hourly, and quarter-hourly products. Market participants submit orders to the limit order book which are cleared continually. If for a market participant the combined orders from spot and future markets deviate from the actual physical production or consumption at gate closure of the intraday market, the residual quantities are settled on the balancing market. The price charged or paid for these deviations is the so-called symmetric *reBAP* (Bundesnetzagentur, 2012).

Each buy and sell order on the intraday market for a given product contains basic information about quantity, limit price, and validity time. A *market order* is cleared immediately against the best available order in the limit order book (LOB), while a limit order is only executed with matching orders on the other side of the market up to a certain price (the limit). If this is not possible, the order is kept in the limit order book until its *end validity date* to be cleared with future orders. If the quantities of two matched orders do not agree, the order with the higher order quantity is only partially cleared and remains in the order book with a correspondingly reduced quantity.

Market participants can add the usual order qualifiers such as *all-or-nothing*, *immediate-or-cancel*, or *fill-or-kill* (EPEX, 2020a). Additionally, *iceberg* orders are allowed for which only a fraction of the order quantity is visible to other market participants. As soon as the visible quantity is cleared, the next part of the order is automatically placed in the limit order book.

The state of the LOB changes with the placement of a new order, with the modification of an order, and at the end-validity-time of an active order. The limit price of the order with the lowest sell price is called *best-ask*, while the order with the highest buy price defines the *best-bid*, and the difference between the two prices is the *bid-ask-spread*.

2.2 The Influence of Renewable Generation

Because electricity is bought by most consumers for a price that is only infrequently updated, short-term consumption is inelastic. Furthermore, due to limited storage, supply and demand have to be matched instantaneously. Consequently, supply and demand shocks can lead to massive shifts in short-term prices (Weron, 2014).

One frequent source of supply shocks is the deviation of produced wind and solar power from its forecast levels. Typically, owners of VRES sell electricity on the day-ahead market one day before delivery based on forecasts of wind speeds and solar irradiation. If those forecasts turn out to be incorrect, the residual quantities have to be traded on the intraday market or resolved on the balancing markets. Since

the latter is typically more expensive, VRES producers have an incentive to balance forecast errors on the intraday market as best as they can.

In particular, if a trader sold too much energy on the day-ahead market she will try to buy back *missing energy* on the continuous intraday market as soon as more accurate forecasts become available and the error becomes apparent thereby increasing demand. An analogous situation occurs if too little energy was sold, which induces an increased supply leading to downward pressure on the intraday prices. Due to the rapid expansion of VRES capacities in many countries and the high correlation of forecast errors for VRES production within a market zone, large unexpected aggregate deviations from production forecast are frequently observed and significantly influence the intraday price (Kiesel and Paraschiv, 2017; Kulakov and Ziel, 2019).

Traditionally weather forecasts are based on large computationally expensive models that depend on satellite images and high altitude measurements of planes and weather balloons, which are only collected every couple of hours. These forecasts are therefore updated too infrequently to be used as inputs for algorithmic trading strategies on the intraday market.

However, recently, several providers specialized in combining these traditional global weather forecasts with real-time production data and local weather models to offer frequent updates of forecasts for renewable production of single plants. Currently, there are many providers such as Enfor, ConWX, Meteologica, Gnarum, enercast, weathernews, or windsim that compete to provide more accurate VRES power production forecasts and more frequent updates.

2.3 The Role of Liquidity

Liquid markets are necessary for the successful implementation of the trading strategies considered in this paper. The observations in this section therefore informs the discussions in the later sections. For a more comprehensive treatment of the liquidity of the German intraday market, we refer to Kuppelwieser and Wozabal (2020).

Liquid markets allow trading for fair prices at low transaction costs and with little scope for price manipulation by dominant players. While traded volumes on the German continuous intraday market have been continuously increasing in the last years, the liquidity of the market is still rather limited at times. Most orders are placed shortly before the market closes and consequently, liquidity is typically low at the beginning of the trading session, increases towards physical delivery, and decreases again shortly before the market closes.

As can be seen by comparing panel 1 with panel 2 and 3 of Figure 1, the liquidity of the intraday power market is significantly worse than that of financial markets. The comparison reveals that, relative to the price, the bid-ask-spread for a share of a large company is roughly 50 times smaller than the bid-ask spread of the continuous power market during its most liquid period. Inspecting the lower two plots depicting bid and ask prices on the German intraday market for a typical trading session of an hourly product, we recognize the characteristic *L-shape* in the bid-ask spread with large differences between the two prices which suddenly falls to a low value close to delivery as also observed by Balardy (2018). We note that the market for half-hourly and quarter-hourly products is even thinner than that for hourly products (e.g. Narajewski and Ziel, 2020). The comparison of the two plots in panel 2 and 3 reveals

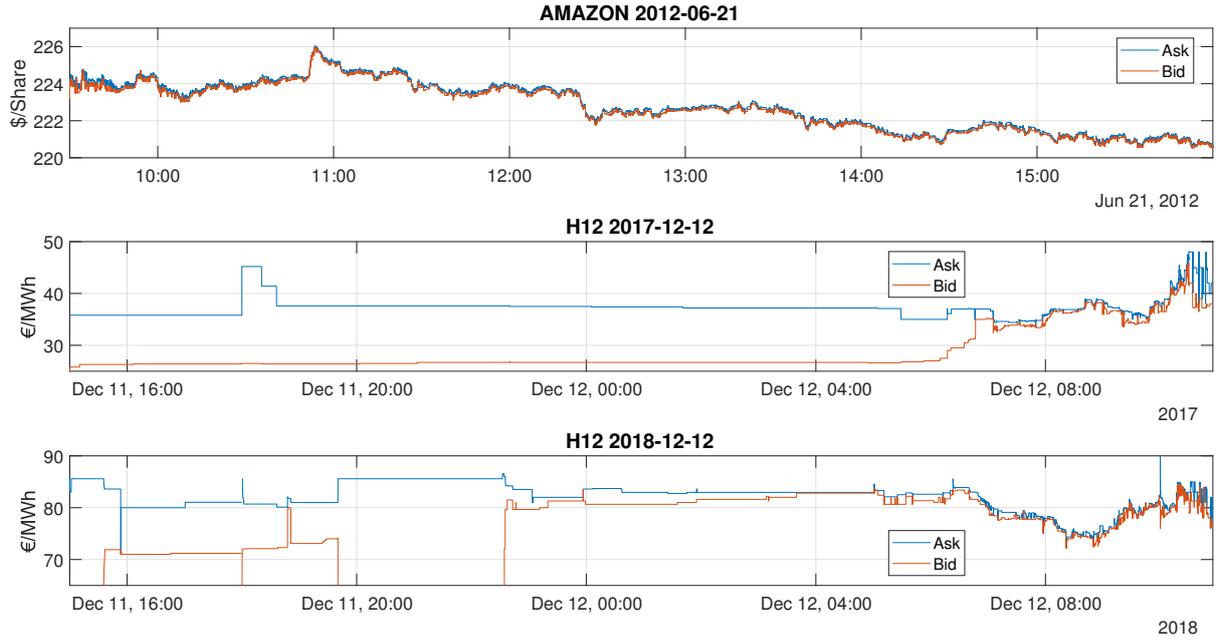


Fig. 1 Financial markets vs EPEX SPOT: The three plots show the best-bid and the best-ask of one trading session. The upper plot shows the Amazon share (AMZN) traded on Nasdaq, the middle plot shows prices for the product H12 which delivers power from 11:00 to 12:00 on the 12.12.2018 as traded on EPEX and the lower plot shows the same product one year after to highlight the increase of trading activity. The data on the Amazon share has been obtained from LobsterData (<https://lobsterdata.com/>).

evidence for an increase in liquidity between the years 2017 and 2018. Finally, the high volatility of the intraday price during the trading session, makes the market attractive for speculative trading.

3 Trading Strategy

Our trading strategy rests on the assumption that a large number of VRES plants sell their forecast production on the day-ahead market and use the intraday market to re-balance their positions so as to take into account updated production forecasts on the day of delivery. The idea behind the strategies discussed in this section is to capitalize on early intraday updates of aggregate VRES production forecasts for the whole of Germany by anticipating the direction of the correction in prices.

To get an accurate measurement of profits, we evaluate the proposed strategy based on detailed limit order book data. In particular, we do not merely rely on tick data or discretized version of the market as for example in Glas et al. (2019, 2020); Kath and Ziel (2020), but take into account the exact rules of continuous intraday market clearing as well as detailed data on orders by other market participants to calculate the price at which we buy and sell electricity.

We are interested in trading strategies that work without physical assets or electricity demand, implying that every product has to be traded separately and positions have to be closed before gate closure. We base our algorithms for the product that delivers electricity in period t on the updates in the forecast of renewable production s hours before delivery

$$\varepsilon_t^s = f_t^{DA} - f_t^s, \quad (1)$$

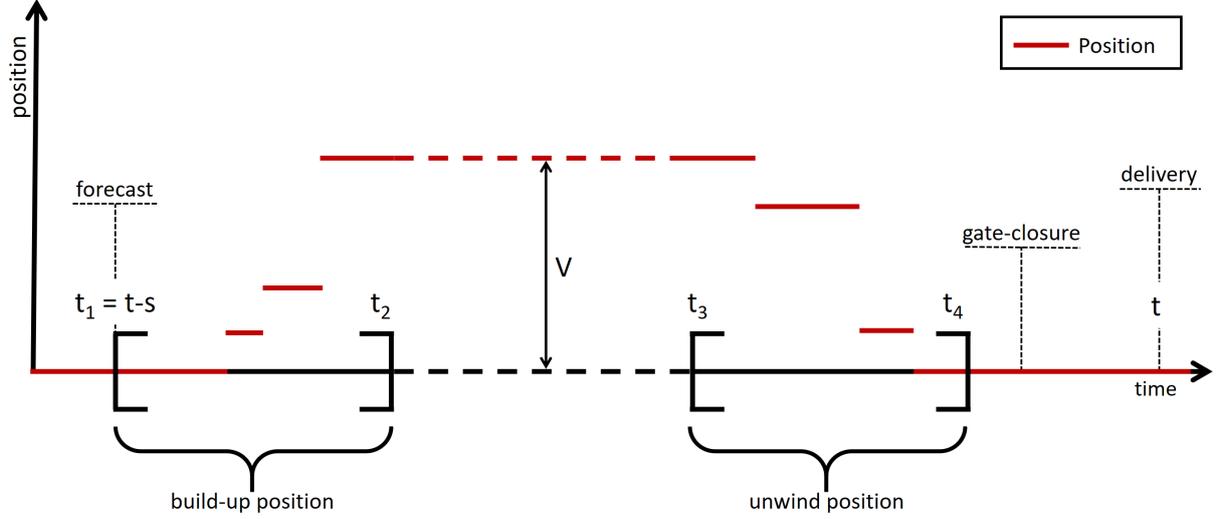


Fig. 2 Schematic depiction of the sequence of events of the proposed trading strategy for the case where energy is bought.

where f_t^{DA} is the day-ahead forecast of renewable production in t while f_t^s is the updated forecast at time $t - s$. The quantity ε_t^s is thus the best estimate of the forecast error in aggregate VRES production at time t which is available at time $t - s$. We adopt the convention that f_t^0 is the actual production, making ε_t^0 the true forecast error of the day-ahead forecast.

Our algorithm takes the form of a classic algorithmic trading strategy on financial markets and uses ε_t^s as a signal that can be used to infer a change in the fundamental value of the product, i.e., electricity to be delivered in period t . This is based on the assumption that traders that first become aware of the errors in forecasts can capitalize on this knowledge by trading accordingly. For example, as a result of a positive ε_t^s , a trader would buy electricity on the intraday market anticipating a rise in prices once the rest of the market becomes aware of the shortage.

However, unlike signals in financial markets like earning announcements or prices of other assets, which can be regarded as public information as soon as they are revealed, information on VRES forecast errors is gradually improved as increasingly better forecasts become available. In particular, the notion of a trader *reacting first* makes much less sense than for signals typically used for high frequency trading on shares markets, since orders cannot be placed *as soon as information arrives* and the decision when to act on updated forecasts becomes important. Traders thus face a trade-off between the reliability of the signal and the speed of the reaction.

To define our strategy, we specify a traded quantity, a price for which we place orders, as well as the timing of orders. We depict the sequence of events in Figure 2. The strategy is triggered by the arrival of a new forecast for VRES production at time t_1 , which is a pre-defined length of time s before delivery of a product t , i.e., $t_1 = t - s$. If the forecast error ε_t^s is large enough, we build up a position in the time interval $[t_1, t_2]$. Subsequently, we hold the position until $t_3 > t_2$ and finally unwind the position in the time interval $[t_3, t_4]$, where t_4 is close to gate closure. Note that since we assume that the trader does not have a physical asset, we require the position to be closed at the end of trading to avoid open positions on the balancing market.

More specifically, we open a position of size $V^\pm > 0$ if the signal ε_t^s observed at time t_1 exceeds a threshold Δ^\pm depending on the sign of the deviation. We thus define the traded quantity at time t_1 as

$$x_{t_1} = \begin{cases} V^+, & \text{if } \varepsilon_t^s > \Delta^+ \\ -V^-, & \text{if } \varepsilon_t^s < -\Delta^- \\ 0, & \text{otherwise,} \end{cases} \quad (2)$$

where positive quantities correspond to buying of electricity, i.e., we buy V^+ MWh of electricity if forecasts are corrected downward by more than a threshold Δ^+ .

Apart from the traded quantity V^\pm , we also need to specify a price to place an order. We investigate two strategies: an *impatient strategy* using market orders and a *patient strategy* based on limit orders. If market orders are used, the price is set to the $\pm 9999\text{€}/\text{MWh}$ which is the maximum/minimum price the trading system allows, i.e., the quantity x_{t_1} is always immediately cleared at time t_1 regardless of the price, provided the order book on the opposing side of the market is not too small to cover the full quantity x_{t_1} . If a market order cannot be (fully) cleared due to a lack of market depth, it is removed from the order book and the second trading phase operates with the correspondingly smaller position. Similarly, at time t_4 the position is closed using market orders. Choosing this impatient strategy thus makes sure that a position is opened as soon as possible and closed at the last possible moment. The downside is that if market depth is insufficient, trading might happen at unfavorable prices.

In contrast, the *patient strategy* places limit orders and accepts a delay in order execution in exchange for potentially more favorable prices. The strategy places an order that outbids the other orders in the system by a small margin $\delta > 0$. For example, if $\varepsilon_t^s > \Delta^+$, i.e., we are seeking to buy, we set the price to be the best bid plus $\delta\text{€}$. If an order with a higher price is added to the order stack at time t' with $t_1 < t' < t_2$ by another party, we update the price of our order to ensure that we outbid the best bid by $\delta\text{€}$. We continue in this fashion until either the whole quantity is traded or time $t_2 > t_1$ comes at which point we remove the order from the system.

We start closing the position at t_3 by again setting the price such that the order is on top of the respective side of the order book and update prices as new orders arrive. Finally, if the position is not closed at time $t_4 > t_3$, we place a market order to close the position. If the order cannot be fully cleared against orders in the LOB at t_4 , the rest of the order is cancelled and the residual quantity is cleared on the balancing market.

Note that opposed to the patient strategy the impatient strategy incurs the full bid-ask spread. For example, if the intention is to buy, then an order on the *ask side* of the market is accepted instead of placing orders on the *bid side* as it is done when using limit orders. Similarly, when closing the position with a market order an existing bid is accepted instead of placing an ask order in the system. Hence, loosely speaking the patient strategy avoids the bid-ask spread for the price of delayed order execution.

In order to calculate the resulting profit, we denote by \mathcal{T}_1 the set of time points at which the LOB changes in the period $[t_1, t_2]$, by \mathcal{T}_2 the set of time points when the LOB changes after t_3 until the end of trading of the product at t_4 , and by V_τ as the quantity traded as consequence of order stack changes at times $\tau \in \mathcal{T} := \mathcal{T}_1 \cup \mathcal{T}_2$. Further, for $\tau \in \mathcal{T}$, we denote by P_τ as the volume weighted average per MWh price for which the quantity at time τ is traded.

The profit and loss of the strategy in period t can thus be calculated as follows

$$\Pi_t = \sum_{\tau \in \mathcal{T}} V_\tau P_\tau + R_t \sum_{\tau \in \mathcal{T}} V_\tau - F \sum_{\tau \in \mathcal{T}} |V_\tau|, \quad (3)$$

where R_t is the symmetric balancing market price for period t and F is the per MWh trading fee. Note that fees on the EPEX are exclusively payable for cleared volumes while modifications of limit orders are not charged. However, we note that the number of modifications is limited to avoid an overload of the trading system. For this purpose, the *order-to-trade ratio* (OTR), defined by the number of order changes divided by the number of placed orders, is limited to 100 by the EPEX.

4 Case Study: Setup & Data

In this section, we discuss the LOB data and the weather reports that we use in the case study in Section 4.1 and Section 4.2, respectively. In Section 4.3, we discuss how we use the data to calibrate the parameters of our strategy.

4.1 Limit Order Book Data

We use German LOB-data for the years 2017 and 2018 as input for the clearing algorithm. The data consists of all submitted orders including information on order changes with timestamps in milliseconds resolution. To test our strategies, we implement the exact EPEX clearing algorithm in *JAVA*. To enable a concise discussion of results, we limit our attention to hourly and quarter-hourly products and do not consider half-hourly products.

Since intraday markets in Europe are increasingly interconnected, some orders in our observation period are cleared against orders from neighboring countries at times when transmission capacities permit cross-border trading. We use the same idea as [Martin and Otterson \(2018\)](#) to deal with this issue by reconstructing the corresponding *foreign* orders using the clearing logs included with the limit order book data. In particular, we check for a counterpart for each executed order in the German LOB. If such a counterpart cannot be found, we add an order with the corresponding price and quantity to the German order book as described in [Martin and Otterson \(2018\)](#), making sure that we can reconstruct published prices with our clearing algorithm. In the considered period there are 47 000 560 orders for hourly products, 1 405 055 (2.9%) of which were cleared against *foreign* orders. For quarter-hourly products there are 139 169 564 orders with 1 495 763 (1.06%) of orders cleared against orders from other markets.

We identify orders for which order quantities are updated immediately after the volume was fully cleared as iceberg orders. These orders are treated as iceberg in our algorithm with the overall quantity that is seen in cleared trades.

The algorithm calculates a clearing at each modification of the limit order book, i.e., if a new order is added, an active order is updated, or an order reaches its end-validity-time. If multiple orders with the same price arrive simultaneously, orders with lower ids are cleared first.

Similar to the results in [Martin and Otterson \(2018\)](#), the prices and cleared quantities computed by our clearing algorithm show a good match with the historical transaction data published by the EPEX.

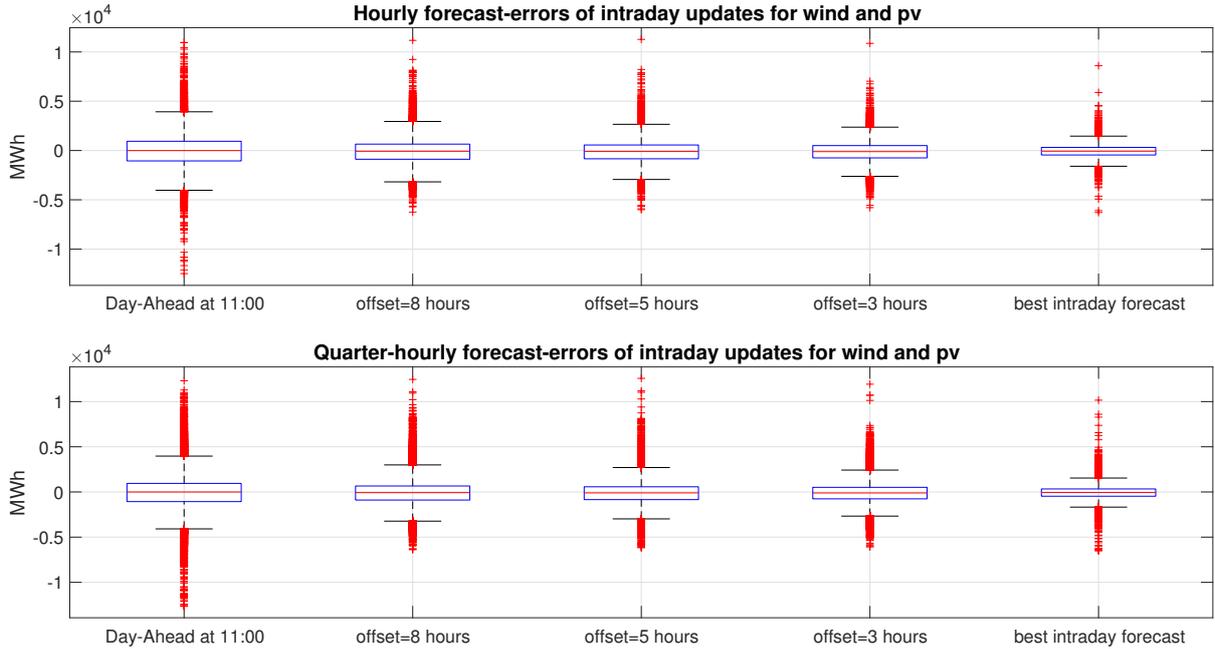


Fig. 3 Forecast errors of intraday forecasts for hourly and quarter-hourly products traded on the German intraday power market between July 2017 and December 2018. The best forecast refers to the last forecast before delivery whose exact timing slightly varies with the product.

We thus are able to accurately evaluate how the market would have cleared additional orders added to the LOB by our trading strategies, which enables us to conduct a historical backtesting.

4.2 Weather Forecasts

In order to execute our strategies, we require the signals ε_t^s defined in (1), which are defined based on aggregated historical forecasts of solar and wind power production in Germany kindly provided by Meteologica¹. Our data consists of day-ahead forecasts available at 11 a.m. the day before delivery, the latest available intraday forecast before gate closure, and intraday forecasts with an offset of 8, 5, and 3 hours before the delivery of a product from July 2017 until December 2018.

To assess the forecast errors, we use data on realized production of solar plants and wind parks for the four German control areas as provided by ENTSOE.² Box plots of the forecast errors are provided in Figure 3. We observe an increasing accuracy with smaller offsets as better weather forecasts and measurements of realized production become available.

Our strategy is based on the expectation that errors in day-ahead forecasts are predominantly traded on the intraday market and therefore have the potential to change intraday prices for power, i.e., can be used as valid signal for changes in the true fundamental value of the product. Consequently, for our strategy, the most important aspect of weather forecasts is whether the sign of the error of the day-ahead VRES forecast can be predicted from the updated intraday forecasts.

¹ <http://www.meteologica.com/>

² <https://transparency.entsoe.eu/>

$ \varepsilon_t^s $	Hourly Contracts						Quarter Hourly Contracts					
	$s = 8$		$s = 5$		$s = 3$		$s = 8$		$s = 5$		$s = 3$	
	%	Hits	%	Hits	%	Hits	%	Hits	%	Hits	%	Hits
>0	100.0	71.2	100.0	74.5	100.0	77.9	100.0	71.1	100.0	74.2	100.0	77.4
>100	87.6	73.8	90.3	76.8	91.5	80.0	87.6	73.6	90.3	76.4	91.7	79.7
>200	76.4	76.3	80.8	79.0	83.4	82.1	76.3	75.8	80.8	78.7	83.5	81.7
>300	66.3	77.9	72.0	81.0	75.5	83.8	66.5	77.6	72.2	80.6	75.8	83.4
>400	58.3	79.6	64.0	82.7	68.5	85.4	58.1	79.3	64.2	82.3	68.5	85.1
>500	50.4	81.2	56.8	84.2	61.9	87.0	50.7	80.7	57.2	83.8	62.0	86.5
>1000	25.7	89.0	31.6	90.2	36.1	92.4	26.0	88.5	31.7	89.9	36.4	92.0
>1500	13.6	93.9	17.5	94.3	20.8	96.1	13.8	93.1	17.7	94.1	21.0	95.5
>2000	7.3	97.2	9.7	97.1	12.1	97.8	7.5	96.9	9.9	96.4	12.3	97.6

Table 1 Distribution of the size of absolute forecast errors (in MWh) in intervals (%) and fraction of correct predictions (hits) of the sign of the forecast error ε_t^0 based on the magnitude of the signals ε_t^s .

We investigate this aspect in Table 1, which displays how often the sign of the forecast error ε_t^0 is correctly predicted by ε_t^s depending on the magnitude of the signal, i.e., $|\varepsilon_t^s|$. In line with expectations and the results in Figure 3, the precision of the forecast increases as the data is restricted to products with higher absolute values of ε_t^s for all s and both types of products. It can also be observed that shorter time to gate closure yields a consistently higher hit rate. However, the increase in accuracy is only moderate. Hence, it seems that earlier signals are not much worse while at the same time give the trader more time to react to the signal. Finally, comparing hourly with quarter-hourly products, we observe that the latter yield worse forecasts of the sign of ε_t^0 in most cases, but the differences are minute.

4.3 Calibration and Evaluation of the Policy

We generate counterfactual profits for our strategies in an as-if valuation of market clearing based on the available LOB data. To that end, we inject orders generated by the trading strategy introduced in Section 3 into the order book and then clear the market according to the rules of continuous trading. Note that this introduces changes relative to the historically observed traded quantities and prices and yields the profits that could have been made, if the strategy was used. Of course, a limitation of these experiments is that, by the very nature of our analysis and the available data, we cannot take into account the effect that the orders placed by the strategy would have had on the behavior of other market participants.

As discussed in the previous subsection, we use data on intraday updates of day-ahead forecasts for VRES production as signals for our strategy. Based on a preliminary analysis of trading profits and in order to facilitate the discussion of results, we only use the forecast 8 hours before delivery for our policies, i.e., consider ε_t^8 as signal. This is also supported by the results in Section 4.2, which show only a moderate improvement of the hit rate for later forecasts.

Furthermore, the choice ε_t^8 has two advantages: Firstly, it allows the policy to start trading relatively early on the updated information before most other traders update their expectations on renewable production. Secondly, the long period from the arrival of the forecast until gate closure gives the strategy ample time to build up the position and thereby avoid excessive liquidity costs.

We thus fix the time t_1 to start the algorithm at 8 hours before delivery and set t_2 such that the policy has 5 hours to build up the position. After that, the policy waits for 115 minutes and then starts closing the position at t_3 , 65 minutes before delivery. If the position is not closed at t_4 , 35 minutes before delivery, we place a market order to close the remaining position. Note that since the liquidity shortly before gate closure is markedly better than in the early hours of trading, we are able to choose the interval $[t_3, t_4]$ relatively short in comparison to $[t_1, t_2]$. The choice of timing and the 8 hour forecast as signal remains constant for all hourly and quarter-hourly products and all variants of the strategy.

Having fixed t_1, \dots, t_4 , we optimize our strategies by choosing the remaining parameters $\Delta^\pm = (\Delta^+, \Delta^-)$ and $V^\pm = (V^+, V^-)$ to maximize profits using historical training data on days $d \in \mathcal{D}_1$. In particular, we define a set of possible thresholds $\mathcal{L} = \{100 \cdot i : 0 \leq i \leq 20\} \subseteq \mathbb{N}$ and a set of volumes to be traded $\mathcal{V} = \{1, 5\} \cup \{10 \cdot i : 1 \leq i \leq 30\} \subseteq \mathbb{N}$ for hourly products and $\mathcal{V} = \{1, 2, 3, 4\} \cup \{5 \cdot i : 1 \leq i \leq 6\} \subseteq \mathbb{N}$ for quarter-hourly products. We then use a simple grid search separately for hourly and quarter-hourly products to solve

$$(\bar{\Delta}^\pm, \bar{V}^\pm) \in \arg \max \left\{ \sum_{d \in \mathcal{D}_1} \Pi_d(\Delta^\pm, V^\pm) : V^\pm \in \mathcal{V} \times \mathcal{V}, \Delta^\pm \in \mathcal{L} \times \mathcal{L} \right\}, \quad (4)$$

where $\Pi_d(\Delta^\pm, V^\pm)$ is the sum of profits Π_t as defined in (3) for all products t that go into delivery on day d using the parameters V^\pm and Δ^\pm . For the calculation, we set the trading fees to 0.125€/MWh (EPEX, 2020a) and use the quarter-hourly reBAP prices available from <https://www.regelleistung.net/> as balancing prices.

We note that the choice of Δ^\pm determines whether the algorithm acts on a relatively weak signals, i.e., for small values of ε_t^s , or whether a strong signal, i.e., a large forecast error, is required to open a position at t_1 . Clearly, for small Δ^\pm the strategy trades products for which the forecast error might only have a small effect on prices, resulting in a high chance that prices move in the opposite direction due to the influence of other factors such as plant outages or changes in demand. Furthermore, for small estimates of the forecast error ε_t^s , the probability that the actual forecast error ε_t^0 has the opposing sign is significantly greater than for larger forecast errors as illustrated in the discussion in Section 4.2. For example, if ε_t^s takes a small positive value 8 hours before delivery, forecasting that there will be shortage in production, the actual day-ahead forecast error ε_t^0 might still be negative, i.e., VRES producers might be long.

In contrast, larger values on Δ^\pm make the strategy react only to strong signals increasing the chance that forecast errors ε_t^0 have the same sign as ε_t^s and are driving prices in the anticipated direction in the time window $[t_3, t_4]$. However, if Δ^\pm is chosen too large, then the strategy will rarely open a position decreasing overall profits. The optimization in (4) thus seeks to navigate this trade-off by choosing optimal parameters Δ^\pm .

The second set of parameter chosen in (4) are the traded volumes V^\pm . Large volumes will generate large profits if signals are reliable and the price response is moderate, while small orders that incur less transaction costs are preferable if markets are illiquid. Note that due to the rules for building up a position, it might be that even though V^\pm is large only smaller quantities are actually traded in some hours, where the market is illiquid.

In the next section, we will investigate profits obtained from applying our policy calibrated using a set of training days \mathcal{D}_1 to some (possibly) different set of days \mathcal{D}_2 , which are used as test data. If $\mathcal{D}_1 = \mathcal{D}_2$,

then the measured profits are insample profits, i.e., the policy is calibrated using the same data that is used to evaluate profits. If $\mathcal{D}_1 \cap \mathcal{D}_2 = \emptyset$, the profits for the days \mathcal{D}_2 are out-of-sample profits.

5 Results and Discussion

In this section, we first present the results of a case study using 1.5 years of German LOB data from the 01.07.2017 until the 31.12.2018. In Section 5.1, we explore the in-sample profits made by optimally parameterized patient and impatient policies for hourly and quarter-hourly contracts using both the actual forecast error ε_t^0 as well as ε_t^8 . In Section 5.2, we focus on the more profitable patient strategies and partition the data in calibration and test sets optimizing implementable policies, which we evaluate out-of-sample for the year 2018.

We consider exclusively products where the day-ahead forecast, the 8-hour ahead forecast, as well as the actual production of renewables are available. Furthermore, we exclude the third hour on the 29.10.2017 and 28.10.2018 due to data problems connected with day-light saving and the whole of the 27.10.2018 due to missing LOB data. Additionally, we exclude 69 hourly and 190 quarter-hourly products due to an empty LOB shortly before the market closes. This leaves us with 12 492 hourly and 50 055 quarter-hourly products for the period between 01.07.2017 to 31.12.2018, excluding in total 5% of hourly products and 4.85% of quarter-hourly products.

5.1 Insample Results

In this section, we analyze the optimal parameter choice for V^\pm and Δ^\pm as well as optimal profits, setting both the training data, \mathcal{D}_1 , and the test data, \mathcal{D}_2 , to the period ranging from 01.07.2017 to 31.12.2018. Since we use the same data to calibrate the parameters and calculate the profits, the resulting optimal policy violates non-anticipativity and is therefore not practically implementable. In particular, in reality, a trader is forced to choose a trading strategy ex-ante, without knowing market outcomes in the trading period. The results in this section can therefore be regarded as a *in-sample* evaluation of optimal profits.

As discussed in the previous section, we start building up a position 8 hours before delivery for every hourly and quarter-hourly product in the observation period and optimize both the patient and impatient trading strategy. To that end, we evaluate the profit separately for products with positive and negative forecast error for the $21 \times 32 = 672$ (for hourly products) and $21 \times 10 = 210$ (for quarter-hourly products) parameter combinations in $\mathcal{L} \times \mathcal{V}$. The parameters of the policy are kept constant for all products in the observation period.

We start by analyzing the patient strategies based on actual forecast errors ε_t^0 . Figure 4 shows how the choice of parameters influence the profits for the patient strategy with the red triangles marking the maximum profit. Observing results for fixed thresholds Δ^\pm , it can be seen that, as expected, higher volumes lead to higher overall profits but due to limited liquidity, the increase is not linear and from a certain threshold on, there is even an decrease in profits for increasing V^\pm . Similarly, there is a sweet-spot for the required strength of the signal: Profits are initially rising in the threshold Δ^\pm and then start to fall again illustrating the trade off between frequent trading on weaker signals and infrequent trading on stronger signals.

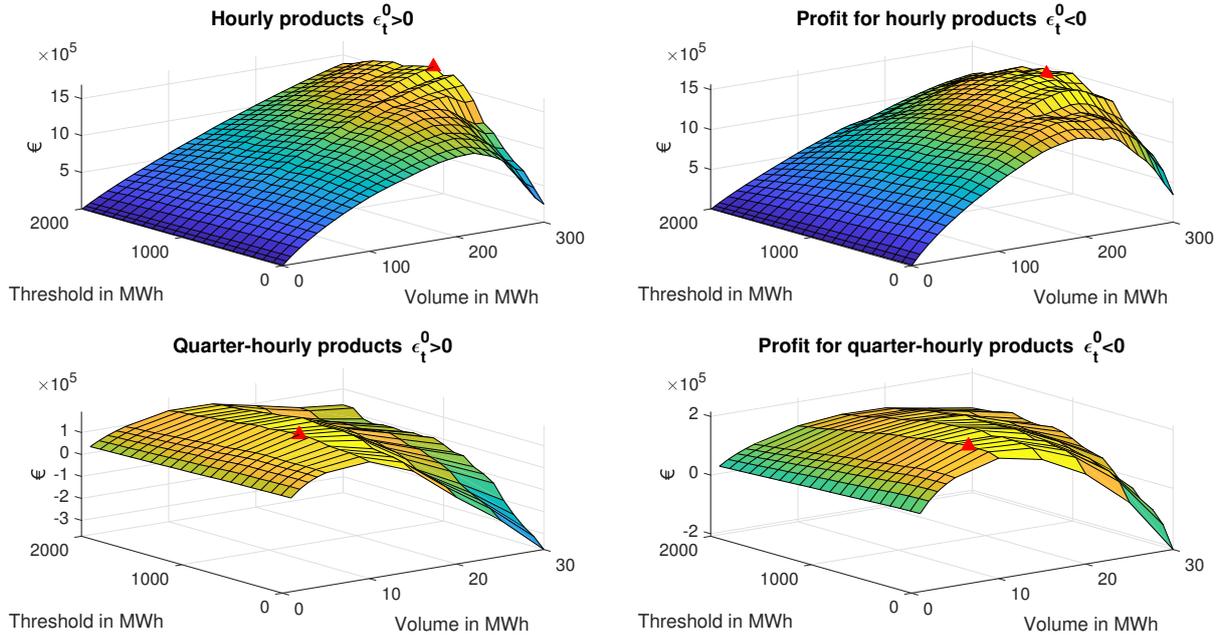


Fig. 4 Optimal profits of the patient trader for real forecast errors for hourly products (above) and quarter-hourly products (below).

			Positive			Negative			Overall
			Profit	V^+	Δ^+	Profit	V^-	Δ^-	Profit
Actual (ε_t^0)	Patient	QH	192 659	10	700	214 774	10	300	407 433
		H	1 686 492	300	1 100	1 560 323	270	1 000	3 246 816
	Impatient	QH	-48 892	1	2 000	-17 350	1	2 000	-66 242
		H	65 167	20	2 000	3 684	1	1 600	68 852
Forecast (ε_t^s)	Patient	QH	48 438	4	200	52 589	4	0	101 027
		H	157 222	200	1 200	331 196	270	1 000	488 418
	Impatient	QH	-30 937	1	2 000	-3 766	1	2 000	-34 703
		H	168	1	1 600	5 607	20	2 000	5 775

Table 2 Profits of insample strategies in € for hourly contracts (H) and quarter-hourly contracts (QH).

The profits and the optimal parameter choices for the considered policies are listed in the first panel of Table 2. The results show that, at least in-sample, a trading strategy that is based on a hypothetical 100% accurate intraday update of the day-ahead weather forecast yields significant positive profits for both hourly and quarter hourly products.

Looking at the profits in detail, two observations can be made. Firstly, hourly contracts are one order of magnitude more profitable than quarter-hourly contracts although there are 4 times more products of the latter. Looking at the optimal parameter choices and in particular at the low quantities traded for quarter hourly products, it becomes clear that this is mostly due to missing liquidity for quarter-hourly products, which start to affect profits already for much lower volumes than this is the case for hourly trading. Secondly, we can observe that the patient trading strategy based on limit orders performs significantly better than the impatient strategy which places market orders. In particular, the results

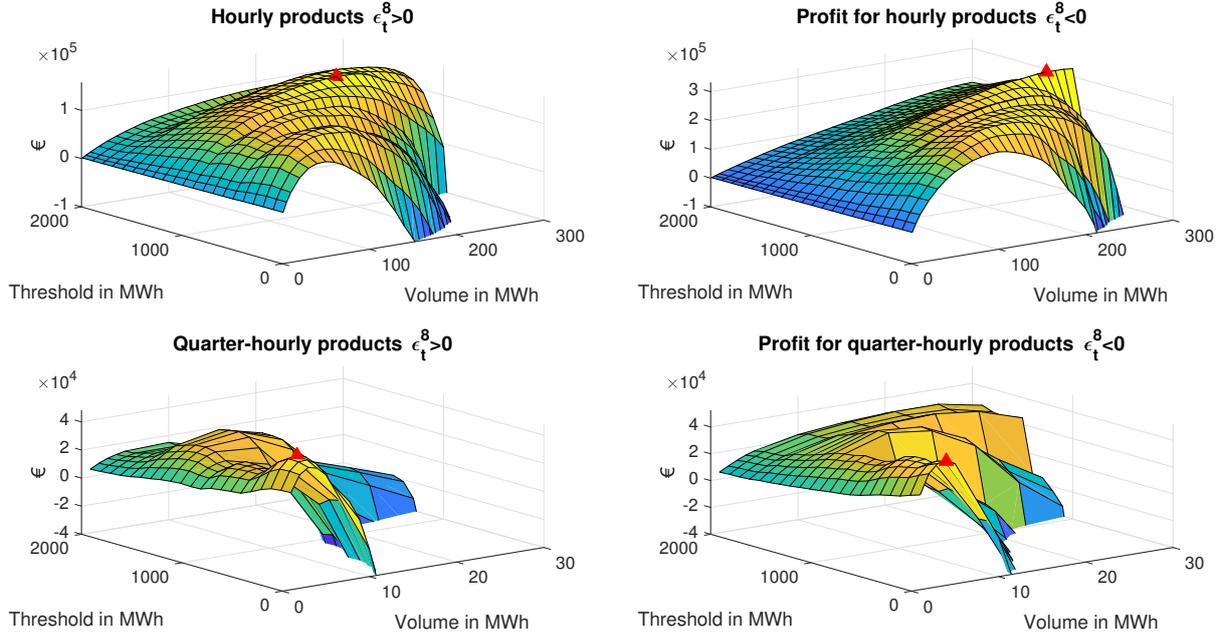


Fig. 5 Optimal profits of the patient trader for forecast errors with an offset of 8 hours for hourly products (above) and quarter-hourly products (below).

suggest that the impatient strategy does not work at all for quarter hourly products and only produces moderate profits for hourly products. Again, this is due to the high liquidity costs in the market which has to be fully born by the impatient strategy.

Next, we analyze the policy for the more realistic case that the signal is based on an updated forecast instead of the actual production, i.e., we use ε_t^8 instead of ε_t^0 as a signal. We again plot the relationship of the parameters of the patient strategy and the profit in Figure 5. The plot exhibits many of the same characteristics as Figure 4 with the difference that higher volumes V^\pm lead more quickly to less profits, i.e., optimal volumes tend to be smaller. This is due to the lower quality of the signal which in many cases leads to a lower than expected forecast error causing losses for policies that bid too aggressively based on ε_t^8 .

Turning to the value of the strategy in panel 2 of Table 2, we observe that, compared to the strategy based on ε_t^0 , profits are significantly lower for the patient trader and stagnate at low levels for the impatient trader. Again, as for ε_t^0 , the hourly strategies yield higher profits but the relative gap is smaller than for the perfect forecast. Although the signal is of a lower quality, surprisingly, the optimal parameters are rather similar to those found for ε_t^0 , although optimal volumes tend to be slightly lower, explaining parts of the lower profits.

The difference between the profits of the strategies based on ε_t^0 and ε_t^8 can be interpreted as a lower bound on the monetary potential of improved weather forecasting, which is substantial for the patient trader.

To put the profits in perspective to the required capital, we evaluate daily capital requirements as the sum of the cost of opening the positions for all products traded on a day, netting out positive and negative costs. The results are displayed in Table 3 and indicate that, on average, the strategy requires

			Mean	Max	Min	Std
Hour	Patient	ε_t^8	-22 163	5 798	-210 712	40 655
		ε_t^0	-57 795	117 889	-38 7446	75 280
	Impatient	ε_t^8	-68	0	-1015	159
		ε_t^0	-2 246	5 277	-23 798	4 015
Quarter Hour	Patient	ε_t^8	404	21 450	-24 865	6 256
		ε_t^0	1 597	49 375	-33 657	10 711
	Impatient	ε_t^8	-141	2 613	-19 710	1 047
		ε_t^0	-276	4 069	-10 835	1 424

Table 3 Amount of net capital invested per day for the different trading strategies.

a negative amount of capital with low positive maximal values. The profits displayed in Table 2 can therefore be realized with a small amount of risk capital and offer a high return on investment.

5.2 Out-of-Sample Results

In this section, we evaluate strategies out-of-sample in the time period from 01.01.2018 until 31.12.2018. More specifically, we study non-anticipative strategies, i.e., make sure that decisions at any point in time only depend on information available at that time [Shapiro et al. \(2009\)](#). Since the impatient strategy performs poorly in-sample, we exclusively focus on the patient strategy for the experiments in this section.

We use a rolling window setting for the out-of-sample evaluation of our strategy and re-optimize the parameters Δ^\pm and V^\pm every day using the last six months of data for the calibration. More specifically, we start our evaluation on the 01.01.2018 using 180 days of training data spanning the period from the 04.07.2017 until 30.12.2017 to calibrate Δ^\pm and V^\pm by grid search as in (4). We then evaluate the profits of the resulting strategy on the 01.01.2018 and proceed to the 02.01.2020 by including the 31.12.2017 in the training sample while removing the 04.07.2017 and retrain our policy to obtain out-of-sample profits for the 02.01.2020. In this manner, we build up out-of-sample profits for every product traded in the year 2018.

Figure 6 shows the results of our experiment for hourly products. The first panel displays the development of cumulative profits of the strategy based on the signal ε_t^8 and ε_t^0 . Looking at the graph for ε_t^8 , it becomes clear that while profits over one year of trading are significantly positive and close to €200,000, there are single days with large losses and extended time periods where the strategy did not generate profits. Comparing with the profits of the strategy that uses ε_t^0 , we see that, as in the insample results, a perfect intraday update of the weather forecast increases the profits by one order of magnitude. Furthermore, the strategy that is based on ε_t^0 exhibits a much smoother increase in cumulative profits with fewer losses. This suggests that the losses for ε_t^8 are mainly due to inaccurate forecasts and suggests that better forecasts can not only increase the profits of the strategy but also reduce the variance of daily profits and therefore the inherent risk of trading.

Turning our attention to panel 2 and 3 of Figure 6, which display the size and the value of the open position after time t_2 for the strategy based on ε_t^8 , we see that the strategy takes long and short positions of up to 200 MWh with a roughly equal share of long and short positions. The position values suggest

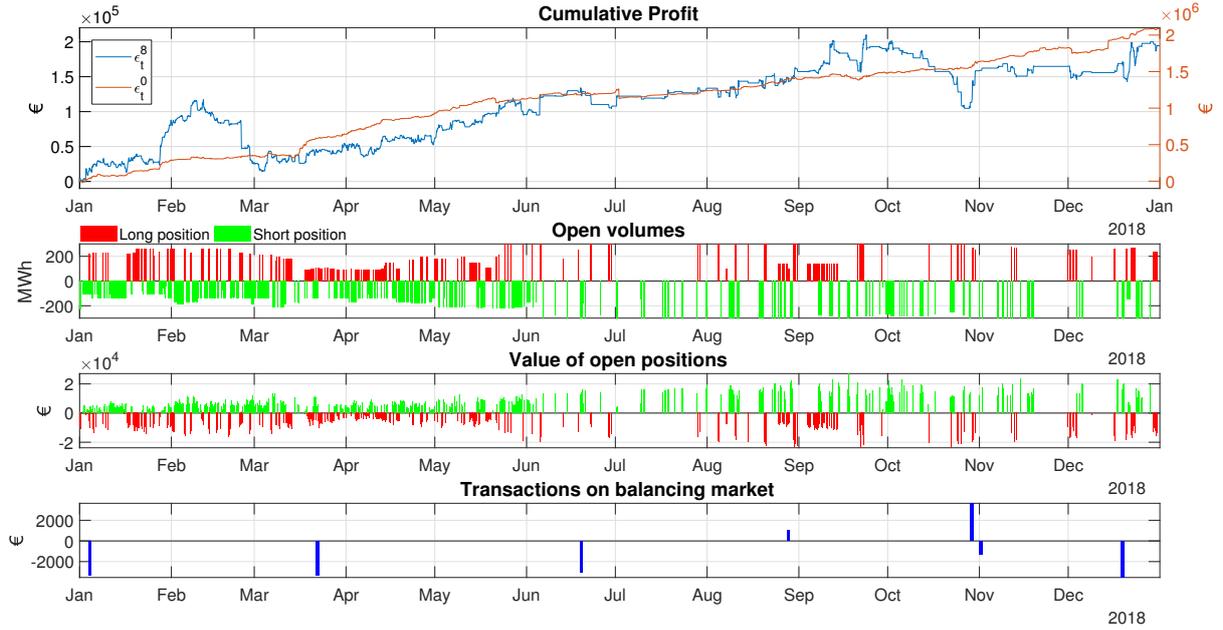


Fig. 6 Cumulative profit of the optimal insample and out-of-sample strategy for hourly products in the first panel. Panels 2 and 3 display the opened volume and the financial value of the positions held by the out-of-sample strategy. Panel 4 shows daily payments on the balancing market for the out-of-sample strategy.

that the capital at risk for single products does not exceed €20,000. It can also be observed that there is a change in the strategy within the observation period: in the first few months the algorithm triggers frequently and short positions tend to be smaller than long positions. In the summer months, there is generally less trading activity, possibly due to lower wind production which lead to smaller forecast errors.

Finally, the last panel of Figure 6 displays netted daily payments from balancing for products for which the position cannot be closed until gate closure. As can be seen, there are only 7 days with a requirement for balancing. In most of these instances the payment is negative, i.e., the trader has to pay to the grid operator for balancing. However, as balancing is rare and none of the single payments to the balancing market exceed €5,000, we conclude that balancing fees are not a major driver of profits for the chosen strategies.

Figure 7 presents an analogous analysis for trading of quarter-hourly products. The plot of the cumulative profits of the strategy reveals that, consistent with the insample results, the strategy is less profitable for quarter-hourly products than for hourly products. As with the insample results and the results on hourly products, the strategy based on perfect forecast is one order of magnitude more profitable than the strategy based on ϵ_t^8 and at the same time is less volatile.

A closer look at the cumulative profits over time reveals that, although the trading of quarter-hourly products yields only roughly one fourth of the profits that can be earned with hourly products, individual earnings for each product fluctuate much less than in the case for hourly products. This is due to the generally smaller positions taken by the optimal strategies which lead to less exposure to market risk as evidenced by panels 2 and 3 of Figure 7. Observing these plots also reveals that there are less seasonal trends in the traded quantities for the quarter-hourly strategy. Finally, the last panel of the figure

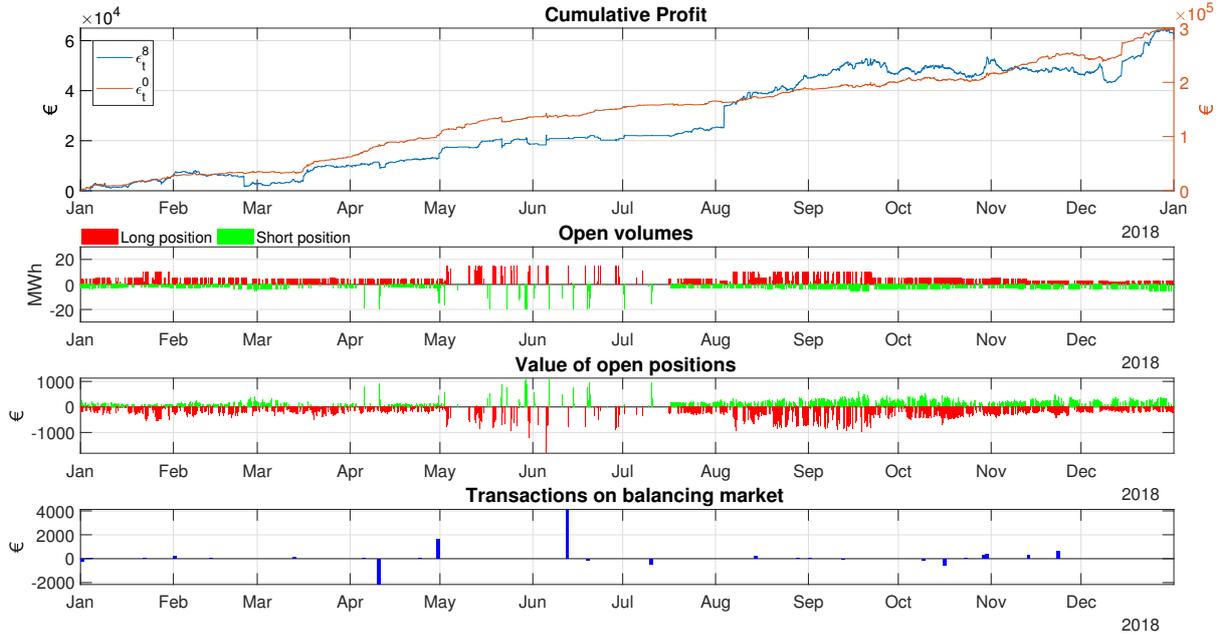


Fig. 7 Cumulative profits, traded volumes, value of traded positions, and daily balancing payments for quarter-hourly products (see Figure 6 for a more detailed description of panels).

	Hour		Quarter Hour	
	ϵ_t^8	ϵ_t^0	ϵ_t^8	ϵ_t^0
Profit	194 385	2 087 823	62 724	297 656
Balancing Costs	-9 865	31 202	4 214	8 055
Mean	22.29	239.43	1.8	8.52
Standard Deviation	968	2 110	44	99
p-value of t-test	0.0316	0.0000	0.0000	0.0000
Minimum	-21 220	-93 030	-1 731	-2 717
1% quantile	-2 814	-3 740	-98	-246
10% quantile	-394	-929	-22	-30
Median	0	0	0	0
90% quantile	522	1 824	28	69
99% quantile	3 137	5 600	118	300
Maximum	15 908	32 174	1 836	3 518
Number of products	8 288	8 288	33 189	33 187
Number of traded products	2 853	4 732	21 425	21 044
Number of individual trades	136 863	311 802	223 593	367 719

Table 4 Descriptive statistics for the profits of different strategies and the number of traded products and trades.

documents that, similar to the case for hourly products, balancing occurs infrequently and therefore only plays a minor role.

Table 4 provides detailed figures for overall profits, balancing costs, and summary statistics for profits per product for both hourly and quarter-hourly trading. Looking at the summary statistics of profits per product confirms that trading quarter-hourly products yields profits with a lower dispersion and

therefore lower capital requirement. Furthermore, conducting t-tests, we see that all average per-product profits are significantly greater than zero at least at the 0.05% level and, due to their lower standard deviation, the significance is greatly increased for quarter-hourly products.

We observe that the number of traded products is nearly twice as high for the strategies based on ε_t^0 as opposed to ε_t^8 . Furthermore, due to the lower thresholds for trading, the relative amount of traded products is larger for the quarter-hourly products. Despite this and the fact that there are more quarter-hourly products, the number of single trades that get cleared as result of our strategy is nearly as high for hourly products as for quarter-hourly products. This is due to the larger quantities traded for the hourly products which often cannot be cleared at once but require trades with a large number of counter-parties dispersed over a larger span of time.

6 Conclusion & Outlook

In this paper, we propose a simple parametric trading strategy for continuous intraday trading on power markets based on intraday updates of forecast VRES production. Our strategy generates significant out-of-sample profits for one year of trading by an arbitrage trader that owns no production assets, has no own demand, and operates on the German intraday market.

Our results show that one of the most important factors to consider when trading on the intraday markets is the lack of liquidity and the resulting transaction costs. In particular, any algorithmic trading strategy has to cope with the limited liquidity of the market, which on the one hand side drives price variability and thereby may favorably influence profits but on the other side makes it harder to capitalize on informational advantages, as any speculative trading strategy has to overcome the bid-ask spread.

We mitigate these problems by designing a patient trading strategy that uses limit orders instead of market orders and allows for an extended time to trade waiting for favorable orders to arrive on the respective other side of the market. We show that this patience is key to making profits and that the impatient strategy incurs substantial liquidity costs that absorb most of the profit that can be generated with weather related information.

Additionally, our results demonstrate that the German intraday market for power is not semi-strong efficient, since publicly available data on weather forecasts can be used to define a trading strategy that generates significant profits while requiring a relatively small amount of risk capital. Furthermore, there would be a substantial potential for even more profitable trading, if weather forecasts were to further improve.

This implies that trading strategies similar to the one presented in this paper, could be a driver for continued innovations in short-term forecasting of VRES production as traders compete in the accuracy of their forecasts. This might trigger an arms race in weather forecasting with market participants trying to capitalize on ever improving forecasts. Algorithmic traders would consequently help the market to process information more efficiently thereby generating price signals of a higher quality and at the same time improve market liquidity.

Additional market liquidity would in turn make weather-based trading easier and more profitable as is demonstrated by, for example, the higher profits generated by our algorithm for the more liquid hourly products as opposed to the less liquid quarter-hourly products. Hence, such a trend could, at

least for a while, feed itself and therefore has the potential to lead to a much more responsive intraday market. Therefore, as opposed to the arguably adverse welfare effects of the *arms race for speed* that characterizes algorithmic trading on financial markets (Budish et al., 2015), this development would likely unlock positive welfare effects.

In our study, we take great care to evaluate the proposed trading strategy as realistically as possible. To that end, we use detailed limit order book data on submitted orders to calculate profits based on an exact implementation of the EPEX clearing algorithm. Furthermore, we make sure that all our policies are non-anticipative, enforcing a strict separation of training and test data.

However, there are still some limitations in our study. Most importantly, we work with historical order data to compute counterfactual profits of our strategy in an as-if fashion. This analysis by design cannot take into account the reaction of other market participants to our trading strategy. A completely different experimental design would be required to overcome this shortcoming.

Another shortcoming of our analysis concerns the quality of the order book data. In particular, we only use German orders even if a small amount of orders is cleared against order from other countries. Although we reconstruct the foreign orders that were historically cleared against German orders, we cannot completely capture the influence that orders from order books of other countries would have had on our results if we had executed our trading strategy. However, due to transmission line restriction, the fraction of German orders cleared with orders from other countries is rather small (below 5%) and we therefore think that our results are robust with respect to this influence.

Furthermore, the order book data supplied by the EPEX is imperfect in many ways impeding a fully accurate what-if analysis. In particular, the end validity date of cleared orders is overwritten with the clearing time which makes it impossible to reconstruct the actual end-validity dates of cleared orders. Additionally, it is hard to correctly identify iceberg orders and market orders from the data. However, since, apart from very few exceptions, our implementation of the clearing algorithm correctly reconstructs historically observed prices, we are confident that the cumulative impact of these issues on our results is negligible.

Our research opens some avenues for further research in weather-based automated trading algorithms on intraday power markets. In particular, it is easy to conceive improvements in the proposed trading strategies. One obvious example is the inclusion of maximum and minimum prices to build up a position as additional parameters of the strategy, preventing trades at unfavorably high or low prices.

This and other possible refinements would lead to a larger number of parameters of the strategy and would therefore necessitate a more sophisticated optimization of the strategy. Possible improvements in this direction could be based on machine learning techniques such as artificial neural networks or reinforcement learning (e.g. Bertrand and Papavasiliou, 2019). Alternatively, one could employ state-of-the-art black box solvers such as CMAES (see Hansen et al., 2010) to find optimal parameters.

Another large area of improvement is in the use of data. Firstly, it is conceivable that the quality of the order book data will improve in the coming years making more accurate analysis of the profits possible and mitigate most of the data related problems described above. Furthermore, as more data becomes available the training of strategies will become more easy and the results more reliable. Secondly, a more careful selection of training data might benefit the performance of the algorithm. For the present paper, we simply use the last 180 days of data to train our strategy for all products. This implies that

data from different times of the day, weekdays, and seasons is used indiscriminately to train the strategy for all products in the test data. Making sure that the training data matches the test data more closely and thus enabling different strategies for different weekdays, seasons, and products has the potential to increase trading profits.

Acknowledgements We are thankful to Meteologica for kindly allowing us to use their renewable production forecasts, which are central to the trading strategies proposed in this article. In particular, we want to express our gratitude to Edoardo Giacomini who was enormously helpful in obtaining and understanding the forecasts.

Data Deposition Information The data that support the findings of this study are available from Meteologica (weather forecasts) as well as from the *EPEX Spot* (order book data) but restrictions apply to the availability of these data, which were used under license for the current study, and so are not publicly available. Data on weather forecasts are however available from the authors upon reasonable request and with permission of *Meteologica*.

References

- C. Balardy. An empirical analysis of the bid-ask spread in the german power continuous market. Working paper, 2018.
- G. Bertrand and A. Papavasiliou. Adaptive trading in continuous intraday electricity markets for a storage unit. *IEEE Transactions on Power Systems*, 2019.
- T.K. Boomsma, N. Juul, and S.-E. Fleten. Bidding in sequential electricity markets: The nordic case. *European Journal of Operational Research*, 238(3):797–809, 2014.
- E. Budish, P. Cramton, and J. Shim. The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. *The Quarterly Journal of Economics*, 130(4):1547–1621, 7 2015.
- Bundesnetzagentur. Beschluss bk6-12-024. <https://www.regelleistung.net/ext/static/rebap>, 2012. [Online; accessed 21-February-2021].
- EPEX. Epex spot operational rules. https://www.epexspot.com/sites/default/files/download_center_files/EPEX%20SPOT%20Market%20Rules_6.zip, 2020a. [Online; accessed 21-November-2020].
- EPEX. Epex spot annual report 2019. https://www.epexspot.com/sites/default/files/download_center_files/Epex-spot-2019_200703_Planche.pdf, 2020b. [Online; accessed 22-November-2020].
- European Commission. Commission regulation 2015/1222. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32015R1222>, 2015. [Online; accessed 22-November-2020].
- E. Garnier and R. Madlener. Balancing forecast errors in continuous-trade intraday markets. *SSRN Electronic Journal*, 6, 01 2014.
- S. Glas, R. Kiesel, S. Kolkmann, M. Kremer, N.G. von Luckner, L. Ostmeier, K. Urban, and C. Weber. Intraday renewable electricity trading: Advanced modeling and optimal control. In *Progress in industrial mathematics at ECMI 2018*, pages 469–475. Springer, 2019.

- S. Glas, R. Kiesel, S. Kolkmann, M. Kremer, N.G. von Luckner, L. Ostmeier, K. Urban, and C. Weber. Intraday renewable electricity trading: Advanced modeling and numerical optimal control. *Journal of Mathematics in Industry*, 10(1):3, 2020.
- N. Hansen, A. Auger, R. Ros, S. Finck, and P. Pošík. Comparing results of 31 algorithms from the black-box optimization benchmarking bbob-2009. In *Proceedings of the 12th annual conference companion on Genetic and evolutionary computation*, pages 1689–1696, 2010.
- C. Kath and F. Ziel. The value of forecasts: Quantifying the economic gains of accurate quarter-hourly electricity price forecasts. *Energy Economics*, 76:411–423, nov 2018.
- C. Kath and F. Ziel. Optimal order execution in intraday markets: Minimizing costs in trade trajectories, 2020.
- R. Kiesel and F. Paraschiv. Econometric analysis of 15-minute intraday electricity prices. *Energy Economics*, 64:77 – 90, 2017.
- M. Kremer, R. Kiesel, and F. Paraschiv. An econometric model for intraday electricity trading. *Philosophical Transactions of the Royal Society A*, Forthcoming, may 2020a.
- M. Kremer, R. Kiesel, and F. Paraschiv. Intraday electricity pricing of night contracts. *Energies*, 13(17): 4501–0, sep 2020b.
- S. Kulakov and F. Ziel. The impact of renewable energy forecasts on intraday electricity prices. *The Energy Journal*, 10, 2019.
- N. Kumbartzky, M. Schacht, K. Schulz, and B. Werners. Optimal operation of a chp plant participating in the german electricity balancing and day-ahead spot market. *European Journal of Operational Research*, 261(1):390–404, 2017.
- T. Kuppelwieser and D. Wozabal. Liquidity costs on intraday power markets: Continuous trading versus auctions. Working paper, 2020.
- K. Maciejowska, W. Nitka, and T. Weron. Day-ahead vs. intraday—forecasting the price spread to maximize economic benefits. *Energies*, 12(4), 2019.
- B.G. Malkiel and E.F. Fama. Efficient capital markets: A review of theory and empirical work. *The journal of Finance*, 25(2):383–417, 1970.
- H. Martin and S. Otterson. German intraday electricity market analysis and modeling based on the limit order book. In *2018 15th International Conference on the European Energy Market (EEM)*, pages 1–6, June 2018.
- M. Milligan, B.A. Frew, A. Bloom, E. Ela, A. Botterud, A. Townsend, and T. Levin. Wholesale electricity market design with increasing levels of renewable generation: Revenue sufficiency and long-term reliability. *The Electricity Journal*, 29(2):26 – 38, 2016.
- C. Monteiro, L.A. Fernandez-Jimenez, and I.J. Ramirez-Rosado. Predictive trading strategy for physical electricity futures. *Energies*, 13(14), 2020.
- M. Narajewski and F. Ziel. Econometric modelling and forecasting of intraday electricity prices. *Journal of Commodity Markets*, 19:100107, 2020.
- I. Oksuz and U. Ugurlu. Neural network based model comparison for intraday electricity price forecasting. *Energies*, 12(23), 2019.
- M. Pollitt. The european single market in electricity: An economic assessment. *Review of Industrial Organization*, pages 1–25, 02 2019.

-
- T. Rintamäki, A.S. Siddiqui, and A. Salo. Strategic offering of a flexible producer in day-ahead and intraday power markets. *European Journal of Operational Research*, 284(3):1136–1153, 2020.
- A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on Stochastic Programming: Modeling and Theory*. MOS-SIAM series on optimization. Siam, 2009.
- S. Séguin, S.-E. Fleten, P. Côté, A. Pichler, and C. Audet. Stochastic short-term hydropower planning with inflow scenario trees. *European Journal of Operational Research*, 259(3):1156–1168, 2017.
- J. Viehmann. State of the german short-term power market. *Zeitschrift für Energiewirtschaft*, 41(2): 87–103, 2017.
- R. Weron. 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, 2014.
- D. Wozabal and G. Rameseder. A Stochastic Optimization Approach for Optimal Bidding of a Virtual Power Plant on the Spanish Spot Market for Electricity. *European Journal of Operational Research*, 280(2):639–655, 2020.