



Copyright Infopro Digital Limited 2023. All rights reserved. You may share using our article tools. This article may be printed for the sole use of the Authorised User (named subscriber), as outlined in our terms and conditions. <https://www.infopro-insight.com/termsconditions/insight-subscriptions>

Research Paper

On the potential of arbitrage trading on the German intraday power market

Elisabeth Finhold, Till Heller and Neele Leithäuser

Fraunhofer Institute for Industrial Mathematics (ITWM), Fraunhofer-Platz 1, 67633 Kaiserslautern, Germany; emails: elisabeth.finhold@itwm.fraunhofer.de, till.heller@web.de, neele.leithaeuser@itwm.fraunhofer.de

(Received June 9, 2023; revised September 7, 2023; accepted September 12, 2023)

ABSTRACT

Pair trading on the German intraday power market is a commonly used risk-averse, heuristic trading strategy. However, due to myopic decision-making and a lack of foresight, the profit obtained is far from optimal. By comparing this strategy with the ex post optimal solution (ie, a strategy with perfect foresight), we show on a set of 15 selected days from 2020 to 2022 that the predictive information lets us generate on average more than five times as much profit by excessively buying and selling the same contracts for a trading interval of five minutes. Another problem with pair trading in practice is the possibility of unbalanced auction wins. We show that an unbiased loss of up to 10% has a negligible impact on the obtained profit. In contrast, we also show the value of frequent optimization updates by simulating strategies with only sporadic participation in the market. While this is hardly beneficial for the pair-trading strategy, the ex post optimal profit increases on average by 30% when the time between two trades is halved.

Keywords: intraday market; arbitrage trading; risk analysis; flexibility marketing; ex post analysis; perfect foresight.

1 INTRODUCTION

The intraday market is the most important market for balancing expected short-term imbalances and providing flexibility. Trading on the intraday market has increased significantly in recent years.¹ This is partly due to the increasing share of renewables, but also to storage options such as batteries, which enable the short-term marketing of available flexibility.

On the continuous German intraday power market on EPEX SPOT,² traders can trade hourly, half-hourly and quarter-hourly contracts. For each contract, the orders are collected in a so-called limit order book. As a result, different prices can be seen in one order book. When matching orders in an order book, the price is determined by pay-as-bid. The trading volume on the continuous intraday market is significantly larger than that of the day-ahead market, in some hours more than double. Even locally optimal trading on the intraday market is a very complex problem, mainly because of its online characteristics. The input data used for decision-making changes within milliseconds, as buy and sell orders arrive and are matched continuously. Moreover, the dynamics of price curves for individual contracts are very difficult to analyze and predict (see, for example, Grindel and Graf von Luckner 2022; Uniejewski *et al* 2019). Market participants therefore often use strategies defined by price thresholds (Bertrand and Papavasiliou 2018).

A slightly more sophisticated strategy is called arbitrage pair trading, which is a widely used trading strategy in (financial) markets where the aim is to find spreads between similar commodities; for example, stocks (see, for example, Elliott *et al* 2005). For an introduction to and a discussion of arbitrage pair trading, we refer the reader to Pole (2011) and Vidyamurthy (2004). Arbitrage pair-trading strategies can be grouped into different approaches in terms of the commodities between which arbitrage opportunities are sought. Their similarity can be defined by a distance approach (see, for example, Bowen *et al* 2010; Gatev *et al* 2006), a cointegration approach (see, for example, Vidyamurthy 2004) or a stochastic spread approach (see, for example, Elliott *et al* 2005). In the context of intraday pair trading, different timescales have been used on various minutes-scale data (see, for example, Bowen *et al* 2010; Dunis and Lequeux 2000) and even on high-frequency data (see, for example, Holy and Tomanová 2018).

When trading on energy markets, there is another important distinction to be made between arbitrage trading strategies. The temporal approach is to buy energy at a low price and sell it later at a higher price. Naturally, these trading strategies need the possibility of storing energy over time and involve some kind of price forecast

¹ See, for example, EPEX SPOT (2021, 2022, 2023), who report increases in trading volumes of 21%, 11% and 9% for the years 2020, 2021 and 2022, respectively.

² URL: www.epexspot.com/en.

and expected trading behavior. For an analysis of trading behavior on a continuous intraday market, we refer the reader to Scharff and Amelin (2016) and Kiesel and Paraschiv (2017). Kremer *et al* (2021) developed a first econometric price model with fundamental impacts for the intraday market of 15-minute contracts.

In contrast, the spatial approach is to find price differences between contracts for the same time period that are traded on different markets, such as different timescale intraday markets (see Metz and Saraiva 2018). Thus, in theory, there is no need for energy storage. Zafirakis *et al* (2016) proposed an analysis of arbitrage opportunities across different European energy markets. They found that in the process of trading more efficiently, these arbitrage opportunities were reduced. In total, several analyses on different markets for the period before 2015 showed that the arbitrage opportunities declined over the years studied (see, for example, Barbour *et al* 2012; Kloess 2012; Steffen 2012). Metz and Saraiva (2018) proposed a mixed-integer programming formulation for arbitrage trading between the German 15- and 60-minute intraday markets on market data until 2016 (see, for example, Metz and Saraiva 2018). However, since then, there have been significant changes to both the market's regulatory framework and its role and size (ie, trading volume).

In this paper we focus on the spatial approach in a specific sense, whereby we consider only arbitrage opportunities between contracts on one market, namely, the German continuous intraday power market for 15-minute contracts.

Theoretically, spatial arbitrage pair trading can be seen as a risk-averse strategy, since every trade yields a profit and does not rely on future market behaviors and price forecasts. The absence of risk of loss, however, holds only in theory, assuming that bids sent simultaneously are either all accepted or all rejected. In reality, it may be the case that only one of a pair of bids can be executed, for example, due to the slow transmission of bids. In this case, the net sum of trades is no longer zero, and losses as well as imbalances can occur. Thus, for the spatial approach as used in this paper, some kind of longer-term storage possibility is also needed. Bowen *et al* (2010) considered a closely related study and discussed the impact of the speed of execution. Moreover, this local greedy strategy is by no means optimal from a global perspective and therefore opportunity costs still arise. Incorporating perfect foresight regarding future prices can provide an insight into the global optimum. Such an approach was used, for example, by Metz and Saraiva (2018), who assumed a perfect knowledge of future prices for different intraday markets.

In this study, we model the intraday market as a rather abstract mathematical model and present an arbitrage pair-trading strategy for the use case of marketing a large battery with bounded flexibilities for each contract. Our main contribution is the quality assessment of this heuristic in comparison with the ex post optimal solution of all strategies that always guarantee a balanced battery at the end of the day on real-world order-book data. In addition to the theoretical trading result, we also

consider the case where not all bids are executed. To the best of our knowledge, an ex post optimal solution has never been computed for the German intraday market. From an application point of view, the optimal solution seems to be of little value as it requires perfect foresight, which is not a given in reality. From a mathematical point of view, the comparison of an online strategy with the ex post optimal solution is known as the competitive analysis and is a standard approach for evaluating an online algorithm (see, for example, Albers 2003; Hertrich *et al* 2022). While this study does not provide any theoretical competitive ratios, it does give numerical insights into the value of accessing the market with a high frequency, the value of price predictions and the risk of loss due to partial bid execution. On a set of 15 days between 2020 and 2022, we show that for high-frequency market access, the pair-trading strategy yields less than 20% of the theoretical optimum, while we obtain around 60% of the theoretical optimum if both algorithms bid only every four hours. As a general rule of thumb, the optimal profit grows by around 30% when the length of the trading period is halved (ie, when the time between trades is halved). This is not true, however, for the pair-trading strategy, where we see only a very small (and not necessarily monotonic) increase in profitability when market access is more frequent.

Overall, the paper investigates the following three research questions.

- (1) Given all orders for contracts on a given day, what profit can be achieved from arbitrage trading with perfect foresight as opposed to a greedy approach with no foresight?
- (2) How does the obtainable profit scale when trading is more frequent?
- (3) How do the achieved profit and the final battery level change when the inexecution of some orders is taken into account?

2 ARBITRAGE PAIR-TRADING MODEL

In the following, we present our arbitrage pair-trading model. First, we introduce some notation. The set of time points considered is denoted by \mathcal{T} . The set of contracts is denoted by \mathcal{C} . Each order i has a volume v_i (in MWh), a price p_i (in €/MWh) and a start and end time. The trading decision variables are denoted by v_i^t for all orders i and all time points t . We say that an order is active at time t if t is in the interval between the start and end times. The set of all orders is denoted by \mathcal{O} . The active buy (respectively, sell) orders at time t are denoted by \mathcal{O}_t^B (respectively, \mathcal{O}_t^S) and are, in addition, restricted to contract C by $\mathcal{O}_t^B(C)$ (respectively, $\mathcal{O}_t^S(C)$). Further, the maximal trading volume in each contract $C \in \mathcal{C}$ is denoted by F . We define the total trading volume of a contract C at t as the sum of trading volumes over all orders

TABLE 1 Parameter and variable descriptions for the arbitrage pair-trading model.

Notation	Description
t, \mathcal{T}	Current time index, set of all time points
C, \mathcal{C}	Contract, set of all contracts
\mathcal{O}	Set of all orders
\mathcal{O}_t^B (\mathcal{O}_t^S)	Set of all buy (sell) orders at t
$\mathcal{O}_t^B(C)$ ($\mathcal{O}_t^S(C)$)	Set of all buy (sell) orders at t restricted to C
v_i	Volume of order i in MWh
p_i	Price of order i in €/MWh
F	Available flexibility in MW
$\mathcal{P}_t(C)$	Position for C at t in MWh
v_i^t	Trading volume of order i at t in MWh

(ie, $\sum_{i \in \mathcal{O}_t^B(C)} v_i^t - \sum_{i \in \mathcal{O}_t^S(C)} v_i^t$). Due to the technical constraints of the underlying battery, we also require that a bid on a contract must first be bought before it can be sold again. We further assume that the size of the underlying battery is 48 times the maximal trading volume. This ensures that battery limit constraints are feasible at each point in time without explicitly considering the limits in our model. This even holds for the extreme schedule in which we are buying the maximum possible volume for the first 48 quarter hours and selling in the remaining quarter hours. For a smaller battery, the battery limits would have to be incorporated into the model. However, it is likely that the outcome will not change significantly with a smaller battery, as we rarely reach the battery's limits. Further, note that having a larger battery would not change the outcome if we require a balanced battery at the end of the day.

2.1 The ex post problem formulation

As mentioned above, we are interested in an evaluation of the profit potential of the German intraday power market from an ex post perspective (ie, how much profit could be obtained under perfect foresight). Such an approach is used as a comparison in situations when the price development can be modeled as a time series (see, for example, Metz and Saraiva 2018). In our approach, however, perfect foresight means that all future changes in the order book are already known. Thus, this cannot be modeled as a simple time series, which makes the ex post optimization problem very large, since all available orders have to be considered at every possible trading time.

Now, we formulate the model as follows:³

$$\min_{v_i^t} \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{O}} p_i v_i^t, \quad (2.1)$$

$$\text{such that } \sum_{t \in \mathcal{T}} v_i^t \leq v_i \quad \forall i \in \mathcal{O}, \quad (2.2)$$

$$\sum_{i \in \mathcal{O}_i^B} v_i^t - \sum_{j \in \mathcal{O}_i^S} v_j^t = 0 \quad \forall t \in \mathcal{T}, \quad (2.3)$$

$$\mathcal{P}_t(C) = \mathcal{P}_{t-1}(C) + \sum_{i \in \mathcal{O}_i^B(C)} v_i^t - \sum_{j \in \mathcal{O}_i^S(C)} v_j^t, \\ \forall t \in \mathcal{T} \setminus \{0\}, C \in \mathcal{C}, \quad (2.4)$$

$$-F \leq \mathcal{P}_t(C) \leq F \quad \forall t \in \mathcal{T}, C \in \mathcal{C}, \quad (2.5)$$

$$v_i^t \geq 0 \quad \forall i \in \mathcal{O}, t \in \mathcal{T}. \quad (2.6)$$

Constraints (2.3) ensure that the overall traded volume of each order is bounded by the available volume. Constraints (2.4) are induced by the pair-trading strategy; that is, they ensure that the total trading volume over all contracts at each point in time is zero. Constraints (2.5) and (2.6) ensure that the position for each contract is feasible at every trading time point; that is, that it lies in the interval from $-F$ to F . We start the optimization runs in a balanced state each day; that is, the position of each contract is zero.

2.2 The greedy algorithm

Since the ex post approach is not applicable in practice, we formulate an iterative procedure for the pair trading on the intraday market. For this, we optimize on a regular basis (eg, once every hour). We call the distance between such optimizations the (trading) period length and test different period lengths. The problem introduced in Section 2.1 now changes as follows. Instead of considering a set of time points \mathcal{T} , we now consider only one point in time, t . Thus, the constraints (2.6) ensuring that the battery level lies in the feasible flexibility interval now reduce to only one per contract. The constraints (2.5) were used for linking different points in time together. Since we now consider only one point in time, these constraints are removed and the values of $\mathcal{P}_t(C)$ are now updated between optimization steps. We will refer to this

³ Note that although this model allows a formulation as a minimum costflow problem in a suitable graph, we nevertheless stick to a linear program formulation since this alternative formulation did not yield any significant improvement in runtime by applying generic flow algorithms in our computational studies. However, it is an open question as to whether algorithms adapted to the underlying graph could provide an improvement.

strategy as the greedy algorithm. Note that for the first trading time point, we start with a position of zero for each contract, which initially allows us to find arbitrage opportunities. Since we are maximizing the obtained profit at each trading time point and a solution that does not trade at all is also feasible, we obtain a nonnegative profit at each trading time point.

2.3 Pair trading under uncertainty

When bidding on the intraday market, uncertainty has a big impact on trading strategies. Recall that it is not only price developments, but also whether orders are placed fast enough to trade successfully that is uncertain. Regarding the latter uncertainty, we use the abovementioned greedy approach and simulate a failure (ie, not being fast enough) for randomly drawn orders. Note that, in contrast to the idea behind our approach to spatial pair trading, the battery will most likely not be in perfect balance after a day of trading (ie, we need to take an imbalanced battery into account).

3 CASE STUDY

On the intraday market at the European Power Exchange (EPEX SPOT), short-term trading of electricity is possible for 15-, 30- and 60-minute contracts.⁴ In this paper, we focus on the 15-minute contracts. The number of contracts is given by the 96 quarter hours of a day. We refer to a contract by the start time of delivery (ie, the contract 12:00:00–12:15:00 is abbreviated to 12:00). Note that the continuous exchange for the 15-minute contracts opens at 15:00 Central European time (CET) (or 15:00 Central European summer time (CEST), respectively) on the previous day and ends 30 minutes before delivery for trades within Germany and 5 minutes before delivery for trades within the same bidding zone.⁵ By “gate closure”, we mean the point at which trading ends on the exchange prior to physical delivery and by “time to maturity” we mean the time remaining until the gate closure of a trading contract. Since trading begins on the previous day, a trading range of at most 32 hours and 40 minutes is considered for the day under consideration (namely, for the contract 23:45 with a gate closure at 23:40). The minimal amount that can be offered is 0.1 MWh. On the continuous intraday market, the prices are determined by pay-as-bid.

3.1 Data and assumptions

The data obtained from EPEX SPOT contains all the orders for quarter-hourly contracts for a given day.⁶ Based on these order books, we create a highly realistic com-

⁴ See footnote 2.

⁵ Until 16:00 CE(S)T on June 1, 2021.

⁶ See footnote 2.


mon data basis for both algorithms. We note that some phenomena, such as iceberg orders, may not have been represented in their entirety when reconstruction was not possible due to missing or inconsistent data (see also Grindel and Graf von Luckner 2022; Martin and Otterson 2018). We consider orders that are visible in the order book to be tradable if the difference between their valid start time and valid end time is greater than 100 ms. Further, we remove all entries that are incomplete. Recall that the daily trading start time changed from 16:00:00 CET to 15:00:00 CET on June 1, 2021. For a more detailed simulation and an analysis of the order books of the continuous intraday market we refer the reader to Martin and Otterson (2018).

3.2 Numerical results

We randomly sample 15 days between 2020 and 2022 (a period for which we have the data). Table 2 (see the online appendix) states some relevant data for the days selected. We consider trading period lengths of 240, 120, 60, 30, 15 and 5 minutes; that is, we take snapshots of the order book at the respective time points and solve the proposed optimization problem based on all orders available at those times. Due to an exponential running time and memory usage, we do not consider smaller frequencies. All computations consider a market offering of 1 MW per contract.

Figure 1 gives an illustrative example of the different evolutions of the algorithms when applied to the same data. We choose the randomly selected sample date of April 11, 2021 to allow the reader to compare different plots for the same data. While it can be seen that the greedy strategy trades the 00:00 contract against the 18:00 contract at the beginning, this cannot be observed for the remaining points in time due to the small sample of contracts listed. Note that despite the very different bidding behaviors, all three contracts end at the same position. This is not true for all contracts, as can be seen in Figure 2 (and in Figure 9 in the online appendix). The color scale in Figure 1 has a slightly different meaning than it does in Figure 9: it is not the absolute position of a contract at a given trading time but the difference of positions between two trading times (ie, the total trading volume for the specific contract at a given trading time). For example, if at a certain trading time we buy 1 MW and 0.5 MW for the 10:00 contract and we sell 0.75 MW for the 12:00 contract and 0.75 MW for the 18:00 contract, then we would see a dark red (1.5 MW) entry for the 10:00 contract and a light blue (-0.75 MW) entry for the 18:00 contract, while the 12:00 contract would not be displayed at all. While our flexibility is limited to ± 1 MW per contract, total trading volumes can exceed these values when the previous position in that contract is different than zero.

FIGURE 1 Evolution of three different contracts in the respective algorithmic strategies for the sample date (April 11, 2021) for a period length of five minutes.



Note that trading time points are given in coordinated universal time. The gray-shaded area marks the time after the maturity of the contracts where trading is not possible anymore. The color denotes the final position of the specific contracts.

FIGURE 2 Final positions per contract for both the ex post strategy (solid blue line) and the greedy strategy (dashed orange line) for a sample date (April 11, 2021) and period length (five minutes).






FIGURE 3 Distribution of the absolute profits obtained for different frequencies by the greedy and ex post algorithms, as well as the relative profit (ie, the ratio of absolute profit for the greedy algorithm to the absolute profit of the ex post algorithm).



3.2.1 Profit

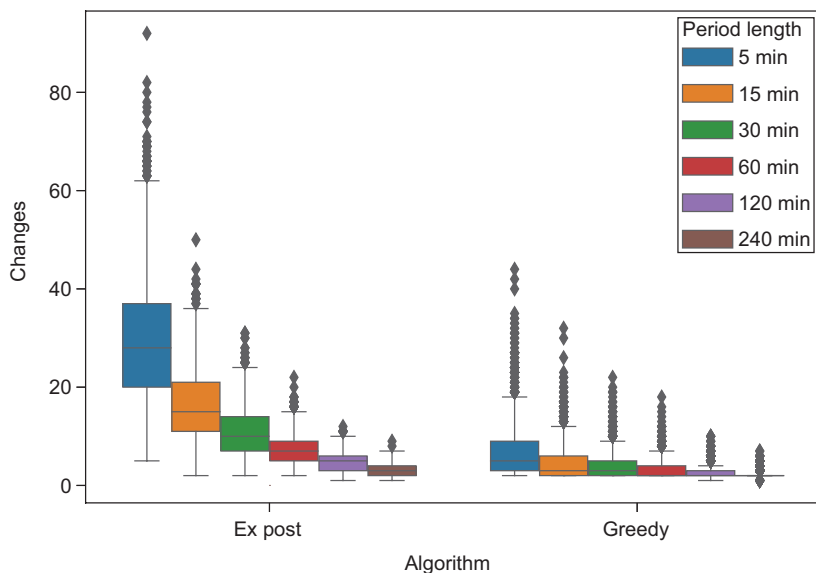
The distribution of the profits obtained for the different frequencies of the two algorithms over the considered days is shown in Figure 3. There are several metrics presented. The boxes show the 25% and 75% quantiles, while the horizontal black line in the middle of the box represents the 50% quantile. The vertical lines above (below)

FIGURE 4 Distribution of profit per position change of the ex post and greedy algorithms for different period lengths (shaded areas give the 95% confidence interval).



depict the distance to the largest (smallest) data point that falls within 1.5 times the interquartile range, which again is the distance between the 75% quantile and the 25% quantile. Both the absolute profit and the range of the obtained profits increase for higher frequencies for the ex post algorithm. The mean profit of the ex post algorithm is smallest for the largest period length and increases by 25–40% when the period length is halved, whereas the profit obtained by the greedy algorithm does not change significantly for different period lengths.


The profit per position change increases for algorithm variants, with longer period lengths. Hence, if optimized less frequently, more profitable trades are concluded (see Figure 4). The main reason for this is that we already see most of the particularly profitable trades for long period lengths; when optimizing more often, we do see some additional options, but on average these are less profitable. The difference between the two variants is not very large; that is, the ex post algorithm sees slightly more profitable trades, but most of the higher profit comes from the higher number of trades. Thus, the greedy algorithm does not make significantly less profitable decisions but merely blocks many future trades by trading too soon.

FIGURE 5 Distribution of the mean changes over all contracts for the different period lengths and algorithms.

3.2.2 Position changes

If the total trading volume of a contract differs from zero for a given trading time, then the position changes accordingly. We count these occurrences per contract and call them position changes. In general, the number of position changes increases with smaller period length for both algorithms, but it does so significantly more for the ex post algorithm than for the greedy algorithm (see Figure 5). A heat map of the positions of the contracts throughout the day can be seen in Figure 9 (in the online appendix), in which it is clear that the ex post algorithm skips the first profit opportunities in order to trade more profitable positions later. It is also noticeable that position changes often take place in contracts whose time to maturity is small (see Figure 6) and whose delivery times are close together (see Figure 9). This phenomenon is much more pronounced for the ex post algorithm than for the greedy algorithm. This is particularly interesting in the context of Kremer *et al* (2020), who show that there is a significant and positively correlated effect on the price changes of neighboring contracts. In addition, the increasing number of position changes on contracts with short maturity is consistent with Kremer *et al* (2021), who show that both the number of trades and the trading volume of 15-minute contracts with short

FIGURE 6 Evolution of position changes regarding their time to maturity (shaded areas give the 95% confidence interval).

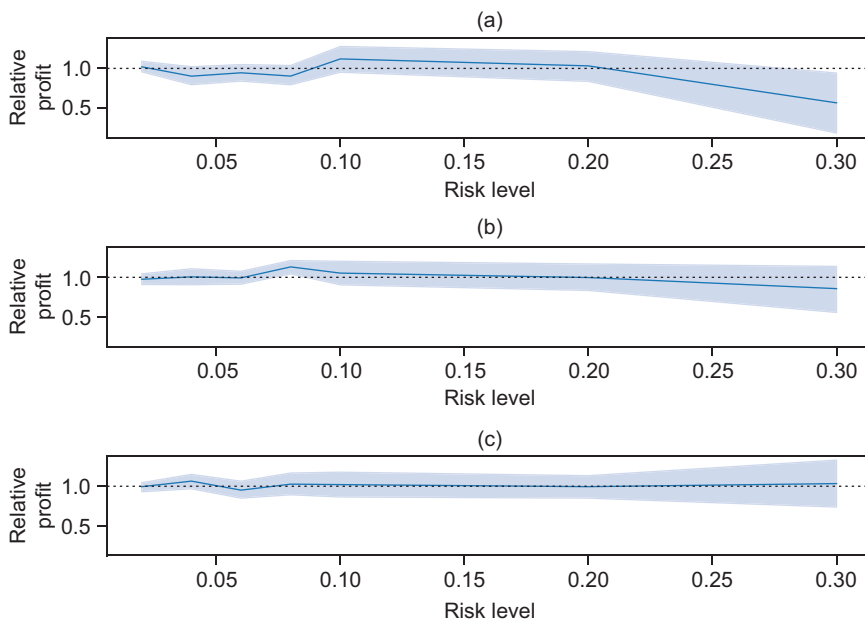


maturity increase toward gate closure. In Figure 2, which shows the final trading position of all contracts for both the ex post and the greedy algorithm, it is noticeable that the positions differ in very few contracts, even if much of the trading by the ex post algorithm happens close to gate closure.

3.2.3 Risk analysis

Due to the slow execution speed of the submitted orders, we may end up not being able to execute all of the intended trades. Thus, in order to make our trading setting more realistic, we consider the inexecution of some of the trades for each trading time point using the greedy algorithm. For this, the inexecution of an order at each trading time is simulated by drawing equally distributed random numbers from the interval $[0, 1)$: the order is executed if the random number is greater than or equal to a considered risk level, and is not executed if it is less than this level. This means that for a risk level of 0.0, all orders are executed, whereas for a risk level of 1.0, no orders are executed. We consider risk levels ranging from 0.0 to 0.1 in steps of

FIGURE 7 Distribution of relative profits obtained by the greedy algorithm for different period lengths for a sample date (April 11, 2021).



(a) 240 minutes. (b) 120 minutes. (c) 60 minutes. All profits are normalized by the profit obtained from the respective run with no risk considered.

0.02 and from 0.1 to 0.3 in steps of 0.1 and we compute 10 runs for each period length–risk pair. For our analysis, we again use April 11, 2021 as a sample date.

In Figure 7, for each risk level the obtained profit relative to the profit obtained by the run without risk consideration is depicted. We observe that for a small risk level, the profit is nearly the same. Thus, the inexecution of some orders has almost no impact on the obtained profit for small risk levels, since there are enough alternative profit opportunities later on. This is in line with the evaluation of the total battery imbalances after one trading day (see Figure 8 in the online appendix). We see that the inexecution of some of the orders has almost no impact for small risk levels. By construction of the inexecution risk, it is of no surprise that the imbalance is distributed around zero. Note that there might be statistical artifacts.

Overall, the key findings are as follows.

- For the ex post algorithm, the total trading volume is much higher and it is mainly the neighboring contracts close to gate closure that are traded.

- The greedy algorithm starts trading too soon, but has a comparable profit-per-position change.
- The final positions do not differ much for either algorithm.
- When the length of the trading period is halved, the profit of the ex post algorithm increases by 25–40%, while the profit of the greedy approach remains almost the same.
- The impact of the inexecution of some orders for small risk levels is almost negligible since there are plenty of profit possibilities and possible future trades that are not blocked by inexecutions.

4 CONCLUDING REMARKS

We investigated the opportunities and risks of arbitrage trading on the German intraday power market for 15-minute contracts. We formulated a general optimization problem for a single trading step as well as an ex post optimization problem.

Our ex post (perfect foresight) results show that arbitrage trading can yield significant profits as a low-risk trading strategy if the right trades are made at the right time. Conversely, results based on the greedy algorithm show that blindly searching for arbitrage opportunities leads to profits that fall far short of those of the ex post strategy. In terms of trading strategies, our ex post results in particular show that by adding predictions of prices or price trends, arbitrage trading can indeed be both a profitable and a low-risk trading strategy. However, if all market participants were to apply this strategy, then in the extreme case these arbitrage effects would disappear from the market and the probability of inexecution (as discussed above) would be higher. When considering shorter period lengths (eg, less than one minute), the profit obtained is expected to be exponentially larger.

Possible improvements to the greedy algorithm include a delayed start to trading, the withholding of flexibility, or even the restriction of products to neighboring products, or the extension toward cross-day arbitrage opportunities. Further, some predictions of future prices or price trends could increase the obtained profit. All of these directions remain open for future research.

DECLARATION OF INTEREST

The authors report no conflicts of interest. The authors alone are responsible for the content and writing of the paper.

ACKNOWLEDGEMENTS

This work was partly supported by the German Federal Ministry for Economic Affairs and Climate Action through grant 01186724/1 (FlexEuro: Wirtschaftliche Optimierung flexibler stromintensiver Industrieprozesse).

REFERENCES

- Albers, S. (2003). Online algorithms: a survey. *Mathematical Programming* **97**, 3–26 (<https://doi.org/10.1007/s10107-003-0436-0>).
- Barbour, E., Wilson, I. A. G., Bryden, I. G., McGregor, P. G., Mulheran, P. A., and Hall, P. J. (2012). Towards an objective method to compare energy storage technologies: development and validation of a model to determine the upper boundary of revenue available from electrical price arbitrage. *Energy and Environmental Science* **5**(1), 5425–5436 (<https://doi.org/10.1039/C2EE02419E>).
- Bertrand, G., and Papavasiliou, A. (2018). An analysis of threshold policies for trading in continuous intraday electricity markets. In *2018 15th International Conference on the European Energy Market (EEM)*, pp. 1–5. IEEE Press, Piscataway, NJ (<https://doi.org/10.1109/EEM.2018.8469774>).
- Bowen, D., Hutchinson, M. C., and O’Sullivan, N. (2010). High-frequency equity pairs trading: transaction costs, speed of execution, and patterns in returns. *Journal of Trading* **5**(3), 31–38 (<https://doi.org/10.3905/jot.2010.5.3.031>).
- Dunis, C., and Lequeux, P. (2000). Intraday data and hedging efficiency in interest spread trading. *European Journal of Finance* **6**(4), 332–352 (<https://doi.org/10.1080/13518470050195100>).
- Elliott, R., van der Hoek, J., and Malcolm, W. (2005). Pairs trading. *Quantitative Finance* **5**(3), 271–276 (<https://doi.org/10.1080/14697680500149370>).
- EPEX SPOT (2021). EPEX SPOT annual market review 2020. Press Release, January 14, European Power Exchange SE, Paris. URL: <https://bit.ly/3LKw1de>.
- EPEX SPOT (2022). EPEX SPOT annual market review 2021. Press Release, January 19, European Power Exchange SE, Paris. URL: <https://bit.ly/48GkpSb>.
- EPEX SPOT (2023). EPEX SPOT annual market review 2022. Press Release, January 23, European Power Exchange SE, Paris. URL: <https://bit.ly/46zIZne>.
- Gatev, E., Goetzmann, W., and Rouwenhorst, K. G. (2006). Pairs trading: performance of a relative-value arbitrage rule. *Review of Financial Studies* **19**(3), 797–827 (<https://doi.org/10.1093/rfs/hhj020>).
- Grindel, R., and Graf von Luckner, N. (2022). Forecasting of the ID3 using limit order book data. Working Paper, Social Science Research Network (<https://doi.org/10.2139/ssrn.4017248>).
- Hertrich, C., Weiß, C., Ackermann, H., Heydrich, S., and Krumke, S. O. (2022). Online algorithms to schedule a proportionate flexible flow shop of batching machines. *Journal of Scheduling* **25**(6), 643–657 (<https://doi.org/10.1007/s10951-022-00732-y>).
- Holý, V., and Tomanová, P. (2018). Estimation of Ornstein–Uhlenbeck process using ultra-high-frequency data with application to intraday pairs trading strategy. Preprint (arXiv: 1811.09312).

- Kiesel, R., and 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>).
- Kloess, M. (2012). Electric storage technologies for the future power system: an economic assessment. In *2012 9th International Conference on the European Energy Market*, pp. 1–8. IEEE Press, Piscataway, NJ (<https://doi.org/10.1109/EEM.2012.6254729>).
- Kremer, M., Kiesel, R., and Paraschiv, F. (2020). Intraday electricity pricing of night contracts. *Energies* **13**(17), Paper 4501 (<https://doi.org/10.3390/en13174501>).
- Kremer, M., Kiesel, R., and Paraschiv, F. (2021). An econometric model for intraday electricity trading. *Philosophical Transactions of the Royal Society A* **379**(2202), Paper 20190624 (<https://doi.org/10.1098/rsta.2019.0624>).
- Martin, H., and Otterson, S. (2018). German intraday electricity market analysis and modeling based on the limit order book. In *2018 15th International Conference on the European Energy Market (EEM)*, pp. 1–6. IEEE Press, Piscataway, NJ (<https://doi.org/10.1109/EEM.2018.8469829>).
- Metz, D., and Saraiva, J. T. (2018). Use of battery storage systems for price arbitrage operations in the 15- and 60-min German intraday markets. *Electric Power Systems Research* **160**, 27–36 (<https://doi.org/10.1016/j.epsr.2018.01.020>).
- Pole, A. (2011). *Statistical Arbitrage: Algorithmic Trading Insights and Techniques*. Wiley (<https://doi.org/10.1002/9781119197072>).
- Scharff, R., and Amelin, M. (2016). Trading behaviour on the continuous intraday market ELBAS. *Energy Policy* **88**, 544–557 (<https://doi.org/10.1016/j.enpol.2015.10.045>).
- Steffen, B. (2012). Prospects for pumped-hydro storage in Germany. *Energy Policy* **45**, 420–429 (<https://doi.org/10.1016/j.enpol.2012.02.052>).
- Uniejewski, B., Marcjasz, G., and Weron, R. (2019). Understanding intraday electricity markets: variable selection and very short-term price forecasting using LASSO. *International Journal of Forecasting* **35**(4), 1533–1547 (<https://doi.org/10.1016/j.ijforecast.2019.02.001>).
- Vidyamurthy, G. (2004). *Pairs Trading: Quantitative Methods and Analysis*, Wiley Finance, Volume 217. Wiley.
- Zafirakis, D., Chalvatzis, K., Baiocchi, G., and Daskalakis, G. (2016). The value of arbitrage for energy storage: evidence from European electricity markets. *Applied Energy* **184**, 971–986 (<https://doi.org/10.1016/j.apenergy.2016.05.047>).