Technische Universität München

TUM School of Management

Essays on the EPEX Spot Continuous Intraday Market

Thomas Kuppelwieser

Vollständiger Abdruck der von der TUM School of Management der Technischen Universität München zur Erlangung des akademischen Grades eines
Doktors der Wirtschafts- und Sozialwissenschaft (Dr. rer. pol.)
genehmigten Dissertation.

Vorsitzender:	Prof. Svetlana Ikonnikova, Ph.D.
Prüfer der Dissertation:	1. Prof. Dr. David Wozabal
	2. Prof. Dr. Sebastian Schwenen

Die Dissertation wurde am 18.03.2022 bei der Technischen Universität München eingereicht und durch die TUM School of Management am 15.10.2022 angenommen.

Danksagung

Ich danke meinem Doktorvater David Wozabal für die Betreuung der letzten Jahre. Die drei Aufsätze dieser Dissertation sind im Rahmen regelmäßiger virtueller Meetings entstanden. Ein weiterer Dank gilt unserem Koautor Christoph Graf für die Zusammenarbeit beim dritten Aufsatz.

Ohne Joachim Messner und Eugen Psaier, dem Geschäftsführer des Unternehmens psaier.energies aus Brixen in Südtirol, wäre diese Industriepromotion nicht möglich gewesen. Das Unternehmen hat mir die Marktdaten zur Verfügung gestellt, welche als wertvolle Datengrundlage dieser Dissertation dienen.

Ein weiterer Dank gilt meinem Mentor Enrico Edoli - grazie mille Enrico per il supporto durante gli ultimi anni - und meinen Arbeitskollegen Martha Oberarzbacher, Hannes Claut und Hannes Mutschlechner für die vielen fachlichen Gespräche und die Unterstützung, besonders während der schwierigen Anfangszeit meiner Promotion.

Danken möchte ich auch meinen Münchner Freunden für die schöne Zeit während meiner dortigen Aufenthalte. Dank euch bin ich immer wieder gerne nach München gefahren.

Schlussendlich danke ich meinen Eltern Helga und Erwin, meiner Schwester Hanna und meiner Oma Johanna, welche stets an meiner Seite standen.

Zusammenfassung

Intraday Strommärkte bieten die Möglichkeit die Produktions- und Verbrauchsprognosen kurz vor Lieferung anzupassen. Deshalb sind diese Märkte besonders für erneuerbare Energiequellen wichtig, aber auch für unerwartete Ausfälle von Kraftwerken. Interessanterweise werden in Europa zwei Marktdesigns für Intraday Märkte verwendet: Auktionen und kontinuierlicher Handel. Beide Designs haben Vorteile und Nachteile. Im Europäischen Intraday Markt hat sich jedoch der kontinuierliche Handel etabliert, mit der Möglichkeit zusätzliche Auktionen hinzuzufügen.

Aufsatz 1: Wir analysieren Liquiditätskosten von Intraday Märkten mit einem cost-of-round-trip Maß, welches sowohl für Auktionen als auch für den kontinuierlichen Handel verwendet werden kann. Wir verwenden Daten der italienischen Intraday Auktionen und des deutschen kontinuierlichen Intraday Marktes und erstellen deskriptive Statistiken sowie ein multivariates Regressionsmodell, um einflussreiche Faktoren auf Liquiditätskosten der beiden Märkte zu finden. Zudem verwenden wir eine *double machine learning* Methode, um Störvariablen auszuschließen und den reinen Einfluss des Marktmodelles auf die Liquidität zu untersuchen. Wir zeigen, dass Wochentage, die jährliche Saisonalität, die Stromnachfrage und Temperatur einen Einfluss auf Liquiditätskosten haben. Der Vergleich der beiden Marktmodelle zeigt, dass Liquiditätskosten grundsätzlich in den italienischen Intraday-Auktionen geringer sind. Interessanterweise nehmen die Kosten der Auktionen nahe an der Lieferung zu. Die Liquiditätskosten im kontinuierlichen deutschen Intraday-Markt hingegen werden gegen Marktschluss immer kleiner und sind kurz vor Stromlieferung geringer als in der entsprechenden italienischen Intraday Auktion.

Aufsatz 2: Wir stellen die erste algorithmische Handelsstrategie für den kontinuierlichen Intraday Markt basierend auf Wettervorhersagen vor. Die Strategie benötigt dafür kein Portfolio - weder Verbrauch noch Produktion - und generiert einen Profit mithilfe von aggregierten Produktionsprognosen der wetterabhängigen erneuerbaren Energiequellen. Dafür stehen mehrere Intraday-Aktualisierungen der Produktionsprognosen von einem professionellen Anbieter zur Verfügung und wir werten die Strategie out-of-sample über ein gesamtes Jahr auf der Grundlage der einzelnen Gebote des deutschen Orderbuchs aus. Unsere Strategie erwirtschaftet einen Gewinn, wodurch gezeigt wird, dass Intraday Märkte nicht *semi-strong efficient* sind. Zudem können die Gewinne mit besseren Produktionsprognosen gesteigert werden und deshalb ist die Qualität der Vorhersagen ein wichtiger Faktor für profitable Handelsstrategien. Das könnte zu einem Wettrüsten für häufigere und bessere Produktionsprognosen führen, welches die effizient des Marktes, die Qualität der Preissignale und die Liquidität des Marktes steigern würde.

Aufsatz 3: Der kontinuierliche Handel hat sich mittlerweile als Marktmodell der europäischen Intraday Märkte etabliert. In diesem Aufsatz schlagen wir hingegen regelmäßig organisierte Auktionen als Alternative vor. In einer Auktion werden die Angebote gesammelt und dann erst gecleart. Das führt zu niedrigeren Liquiditätskosten, zuverlässigeren und weniger störende Preissignale, und die längere verfügbare Berechnungszeit erlaubt es zudem die technischen Einschränkungen des Stromnetzes besser zu berücksichtigen und die Stromflüsse optimal zu steuern. Wir verlgeichen in einer empirischen Studie die kontrafaktischen Ergebnisse von regelmäßig organisierte Auktionen mit jenen des kontinuierlichen Marktes. Das Handelsvolumen ist beim kontinuierlichen Markt höher, Liquiditätskosten sind jedoch bei den Auktionen geringer und die Preissignale beinhalten weniger Rauschen. Zudem diskutieren wir die Verwendung des kontinuierlichen Marktes mit den vorhandenen Netzwerkbedingungen und technischen Einschränkungen der konventionellen Anlagen. Zusammenfassend zeigen unsere Ergebnisse, dass das Sammeln von Angeboten in regelmäßig organisierten Auktionen besonders für schwach gehandelte Intraday Märkte von Vorteil wäre.

Abstract

Intraday electricity markets offer the possibility to adjust production and demand forecasts shortly before physical delivery. Hence, these markets are important to integrate renewable energy sources, but also to balance unforeseen outages of power plants. Interestingly, two market designs for intraday markets are used in Europe: auction markets and continuous trading. Both concepts have advantages and disadvantages. The continuous trading design was chosen for the European intraday market, with the possibility of optional additional auction markets.

Essay 1: We analyze liquidity costs on continuous and auction-based intraday power markets using a cost-of-round-trip measure that works for both market designs. We use data from the Italian auction-based intraday market and the German continuous market and present descriptive statistics as well as multivariate regression models to analyze determinants of liquidity costs in both markets. To test for differences in liquidity due to market design, we employ a double machine learning technique controlling for several confounding variables. We show that weekly patterns, yearly seasonality, electricity demand, as well as the influence of temperatures significantly affect liquidity costs. Comparing liquidity costs in both market, we find that, overall, liquidity costs are lower on the Italian market. However, Italian costs increase towards later auctions, while the costs on the German continuous intraday market decrease and reach their low close to physical delivery, where costs are lower than on the last Italian market trading the corresponding products.

Essay 2: 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.

Essay 3: Continuous trading is currently becoming the standard for intraday electricity markets. In this paper, we propose frequent auctions as a viable alternative. We argue that batching orders in auctions potentially leads to lower liquidity cost, more reliable, less noisy price signals, and allows for better alignment of market outcomes with the technical realities of the grid. In an empirical study, we compare the German continuous intraday market with counter-factual outcomes from frequent auctions. We find that while traded volumes tend to be higher for continuous trading, liquidity costs are lower and price signals contain less noise for the auction market. Furthermore, we critically discuss the suitability of continuous trading in the presence of network constraints and technical restrictions of conventional units. Taken together these findings suggest that in sparsely traded intraday markets, pooling orders in frequent auctions may be beneficial.

Contents

List of Figures

1 Introduction

The share of renewable energy strongly increased in the European power mix during the last years. Wind parks, solar farms, biomass plants and hydropower plants are important technologies to reduce CO_2 emissions and help to reach international climate goals. However, the production of wind parks and solar farms are strongly dependent on weather conditions and are therefore difficult to integrate into the complex power grid. Production forecasts of these power plants are continuously updated based on new available weather forecasts and the predictions are getting more precise with shorter lead-times. Intraday markets were introduced in Europe to provide market participants the possibility to trade forecast errors until shortly before physical delivery. These markets are analysed in this dissertation with the focus on the German continuous intraday market.

This introduction provides the preliminaries for the three essays of this dissertation. The basic challenges of the electricity sector, which are important to guarantee a secure power supply, are described in Section 1.1. Section 1.2 describes the transition of the European energy sector from a monopolistic structure to a liberalised market. Section 1.3 describes the current European market design with sequential power markets, which allows to trade long-term contracts, but also short-term contracts with a finer granularity. The homogeneity of the European market design and the small differences in the single markets will be discussed in Section 1.4 by introducing the German and Italian power markets. In Section 1.5 the limit order book of the continuous intraday market is described, while Section 1.6 analyses basic measures of the limit order book. Section 1.7 introduces the notion of liquidity, which is an important property of markets in general. Finally, all research questions, which are tackled in the three essays, are collected in Section 1.8.

1.1 Challenges of the Electricity Sector

Power can be traded on markets similarly to other commodities like oil, gas or metals. Nonetheless, there is a huge difference to those other commodities, as storage of power is cost intensive leading to the strict constraint that production has to meet consumption at each point in time. A violation of this constraint would disturb the frequency of the entire grid; the consequence could be a blackout.

One particularity of the electricity sector is the power grid. Power lines are used to connect customers with power plants and each power line has a maximum capacity. Customers and power plants form the nodes of the power grid, and power lines define the vertices. Large power grids were developed during the last years to transport power over huge distances. Hence, production has to match consumption at each point in time, but also the constraints of the power grid have to be considered. Therefore, also electricity markets have to be tailored to these technical realities.

The spot market allows to trade electric power in blocks of 15 or 60 minutes and these products are also useful to manage the stability of the power grid. Hence, all controllable power plants are scheduled with a granularity of 15 or 60 minutes to balance the power grid roughly within each block. Other flexible power plants are used on-top to balance the power grid at each time within the block.

One could easily solve the linear power flow problem knowing the exact consumption by adjusting the output of each power plant. However, the real problem is not that simple, because a big challenge in the energy sector is uncertainty. There are many unexpected events which have to be considered to guarantee a secure power supply:

- **Uncertain demand**: Consumers do not have to deliver power schedules and have free choice in their use of electric power. Power providers have to schedule the demand of their customers. It is difficult to predict the demand of one single household, but pooling consumers makes it possible to obtain good estimates. However, the demand in the whole power system remains random.

- **Unexpected plant failures**: On the other side, power plants are scheduled to meet the predicted demand and the in-feed into the power grid has to be guaranteed to obtain a balanced system. Despite that, sometimes power plants have unexpected technical problems and are no longer able to fulfill their power schedules. In that case, consumption has to be reduced or another power plant has to be activated.

- **Failures of the power grid**: It is not enough to meet demand at each time, but the power schedules have to consider the constraints of the power grid with their maximum available power line limits. In the case of a power line failure, new power schedules are needed to adapt the power flow under the new available grid constraints. The blackout in 2006 for example was caused by a planned outage of two power lines which had unexpected consequences.

- **Integration of uncertain renewable energy sources**: The share of green energy increased during the last years to reduce CO_2 emissions. The disadvantage of these power plants is that the production of photovoltaic plants and wind parks are strongly depending on the weather. Hence, weather forecasts are used to predict the production of renewables. The precision is getting better with shorter lead-times, and therefore, we obtain quite good production forecasts shortly before physical delivery. However, the late information of the precise production forecasts and the resulting uncertainty makes it difficult to integrate renewable energy sources into the power grid. Moreover, wind parks have periods when they produce a lot of energy, but there are also calm wind periods where these parks produce less energy. Of course, another type of power plant has to provide the needed capacity in these periods. Another important topic is the expansion of the existing power grid, because wind parks are able to produce huge amounts of energy which has to be transported to the right place. Due to the lack of missing power line capacities, some wind parks have to be switched off at windy days to avoid a blackout.

The four challenges require flexible power which is able to balance power at each point in time. Conventional plants are mainly used to balance the system, but also power from gas turbines, pumped storage hydropower, batteries or flexible consumers is considered. Marginal costs of these sources are usually high and it is convenient to minimize the usage of balancing power.

1.2 Liberalisation of the European Energy Sector

Interestingly, different market designs were developed around the world to handle the above-mentioned challenges of the electricity sector. One solution is to introduce a single company which is responsible for the power supply and is organised as a national or regional monopoly. Another option is a competitive energy sector, where customers can choose between different power providers.

While a monopolist is able to adjust the production of its own power plants to meet the demand of customers, a competitive model needs power markets allowing power providers to exchange power. Power markets are able to organize the scheduling process with price signals, where the higher the prices are getting, the more power is missing in the system and vice-versa. Europe decided to transform the monopolistic energy sector into a competitive market.

A few decades ago, before the liberalization of the energy sector started, local companies were founded to set up regional power supplies in Europe. They started to build new power grids and power generators, and the amount of provided electricity was continuously growing. However, these companies were quite small and had difficulties to guarantee a stable power supply due to missing interconnections between power grids. Hence, a transition from local companies to a few quasi-monopolistic companies can be observed. These larger companies were able to manage, interconnect and extend the separated power grids. These companies managed the power grid, generation and were responsible to guarantee a secure power supply, which was important for the growing economy in each country. Each European country developed its own power system with different rules and particularities, and customers were forced to sign a contract with the corresponding regional company.

The UK was a pioneer in liberalizing the power system and introduced three separated generating companies in 1989. These companies published a uniform price through the Electricity Pool and the demand side was able to purchase energy at the offered price. However, three companies were not enough to spark competition, because they were able to raise the price to increase their profit. In 2001, NETA (New Electricity Trading Arrangements) replaced the Energy Pool and included the demand side, and market participants were responsible to balance their position with penalties for imbalances in each settlement period. Transmission and distribution network operators set standard charges to suppliers for the use of the networks to allow the entrance of new providers Littlechild [1992].

Norway was another European pioneer and collected important experience for the liberalisation of the European energy sector, but also for the introduction of interconnected power markets. They started the liberalisation of the power system in 1992 and NordPool was established in 1993 as the Norwegian power market. In 1996 Sweden joined the market and formed the first international power market including cross-border capacities.[1]

The observed results of the UK with lower energy prices due to competition provided the European Union with strong arguments to start the liberalisation of the energy sector. Moreover, the Nordic power market convinced the EU to implement an interconnected European electricity market with convergent prices across the EU. The new power market should provide the European economy with low and transparent energy prices, and the interconnected market should allow to raise the share of renewable energy sources. As described in the previous section, the integration of renewable power sources is challenging, but it is easier to handle renewables in a larger power grid. On one hand, the advantage of different locations introduce diversification and maybe a wind park in the South is producing while the one located in the North is not. On the other hand, wind parks work more efficient in the North while photovoltaic plants are more efficient in the South with more sunny hours during the year. Summarizing, a large power grid allows to handle a larger share of renewable energy sources and this allows to reduce CO_2 emissions.

The transition to the liberalised market was realised with three main European directives: European Directive 96/92/EC[2] in 1996, 2003/54/EC[3] in 2003 and 2009/72/EC[4] in 2009. The progress started to introduce standardized rules and the removal of monopolies by the unbundling of generation, transmission, distribution and retail Pollitt [2019]. The aim was to allow customers to choose

[1] See https://www.nordpoolgroup.com/About-us/History/

[2] See https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A31996L0092

[3] See https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32003L0054

[4] See https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=celex%3A32009L0072

between competitive power providers, but transmission and distribution still remains in the hands of monopolies called Transmission System Operators (TSOs) and Distribution System Operators (DSOs). The TSO is responsible for the secure power supply and manages the high-voltage power grid, which is used to transport power over long distances. The DSOs are responsible for the local power grids in medium- and low-voltage and distributes power to the final customers.

As already mentioned, countries such as Norway, Sweden and the UK liberalised their energy sector already in the early 1990s. Other countries had still to introduce national laws to unbundle generation, transmission, distribution and retail, because these components were largely in the hands of monopolists. After the introduction of the European directives, all countries started to implement the agreed conventions in their local energy sector and the process of building an European energy sector started. However, all countries developed slightly different versions based on their historic energy sector, which provided the initial energy mix and infrastructures. France for example was dominated by nuclear power, Germany by coal resources, Norway by hydro power, and the UK, Spain and Italy by gas plants. Moreover, cross-border trading was controlled between bilateral monopolists on both sides of the border.

Another important step was the introduction of national power markets which guarantee transparent prices and allowed the foundation of new power providers, because they need a platform to trade electric power. The new markets also allow the foundation of companies which are exclusively specialised on production or consumption. Finally, the transparent prices can be used directly or as an index for buying or selling power.

It was already known, that the process towards liberalisation would take years. The unbundling process needs some time, new power providers have to enter the market to allow competition and the introduced markets have to be established. Moreover, new power lines have to be built to increase the cross-border capacities between countries.

Nowadays, most of the European directives are already implemented. Customers can choose between different power providers and the new national power markets are fast growing and connected to a European market. However, the process is not yet completed. There are still plans to couple the remaining markets to the European market and additional power lines are planned to handle the changing energy mix with the increased share of renewables. It is already observable, that large companies such as EdF, E.ON, Iberdrola, ENEL and Vattenfall expanded to other countries in both generation and retail.

1.3 Sequential Power Market Design

There are different types of power plants which contribute to the energy mix. Some plants provide huge amounts of inflexible energy at low marginal costs and are optimal to cover base load. Other plants are flexible enough at moderate prices to cover peak load. Finally, there are highly flexible power plants with high marginal costs which are needed to balance the system. Europe decided to implement sequential power markets which allow to integrate all types of power plants to cover the expected demand and to integrate renewable energy sources with low system costs. For that purpose, future markets, the spot market consisting of the day-ahead market and intraday markets, and balancing markets were introduced.

1.3.1 Future Markets

Futures or derivative markets allow market participants to trade power which will be supplied in the next few years. The two general types of futures are physical futures, which require also the physical delivery of energy, and pure financial futures.

One advantage of future markets is the possibility to limit the price risks of investment costs. The construction of new power plants is cost intensive and future markets allow investors to hedge these initial costs. Future contracts are also suited to hedge fix-costs of primary resources, which are needed for the energy production.

Each production plant owner and customer on the demand side has to sign a contract with a power provider to sell or buy power. The price can be linked on the hourly day-ahead price, on an index or they can agree on a fixed-price. Some production plant owners tend to sell the production with a fixed-price with a fixed margin above the generation costs which also simplifies the calculation of the future cash flow. On the other side, there are also industrial or commercial customers on the demand side which prefer to buy energy based on a fixed-price. These contracts cap the component of their production costs which is based on the resource power. Production plant owners and customers on the demand side can sign these kind of contracts with trading companies. Hence, the financial risk was shifted to the trading company which has the possibility to hedge the position on the future market.

Future markets are important for established players as hedging instruments allowing them to be protected against future price changes. Established players usually have many forms of physical options and forward markets can be used as additional tool for their risk management. Finally, future markets can also be a pure speculative instrument, as financial products do not require physical delivery and allow to bet against future price developments ECA [2015]. Financial futures also allow to speculate on price differences in different countries by combining different products. These financial traders help to increase liquidity of the market.

The EEX (European Energy Exchange) is the largest market for futures in Europe.[5] The available products are years, quarters, months, weeks, weekends and days with the distinction between base and peak products. The limited cross-border capacities between countries require to distinguish between products of different countries. The German Power Futures are the most liquid European power future. The continuous market design allows market participants to monitor the current price and to trade the needed quantity immediately during 8 am until 6 pm, but also allows to register trades between market participants. Trading of a specific product is possible until physical delivery of a product starts. The day-ahead price determines the realized value of each day and the realized value of a product can be determined as soon as the day-ahead price of the end date of the product is available. Afterwards, the financial position between realised price and cleared future price will be compensated. As an example, consider the product $Y2021$ can be traded until 31/12/2020, when the hourly day-ahead prices of 01/01/2021 are published. The day-ahead prices of the last day 31/12/2021 are published on 30/12/2021 and the realised price of $Y2021$ can be calculated as the average of all hourly day-ahead prices. Usually, products on EEX are financial futures, but there is also a physical fulfilment service offered by EEX, which places the corresponding bids on the day-ahead market to force a physical delivery.

There is also the possibility to trade electricity forwards between companies. Buyer and seller agree on a bilateral contract and conclude the trade OTC (over-the-counter). In contrast to future markets, traders are not restricted to standard contracts, but they can also trade products with individual delivery periods and hourly or quarter-hourly profiles.

1.3.2 Day-Ahead Markets

The power schedule is getting more concrete a few days before physical delivery. On the one hand, the demand forecasts are getting better due to the predicted temperature trend, on the other hand, the accuracy of forecasts of renewable production increases with better weather forecasts. The day-ahead market allows to trade the 24-hourly products of the following day and submission of orders is possible a few days before physical delivery. The market is cleared at 12 am the day before physical delivery and determines a transparent price for each hourly product.

The market is organized as an auction market for each hourly product and the clearing price is the intersection between the supply- and demand-curve. In-

[5] See https://www.eex.com/en/markets/power-derivatives-market/power-futures

terestingly, different methods to submit orders were developed in the European day-ahead markets. Some European markets, such as Italy, accept hourly step orders which are defined with side (buy or sell), price and volume. They indicate to trade the provided volume at the given price or a better price, otherwise the order will not be accepted. Other markets, such as Germany, accept hourly linear piece-wise orders, which are defined by side, volume, price where the order starts to be accepted and a second price where the order is fully accepted. Some markets also allow block orders with different conditions on the execution Committee [2019].

Standardized national day-ahead markets were important to start the unbundling of the energy sector. These markets guarantee a transparent price formation and offered a marketplace to trade power allowing the entrance of new market participants. Companies have the possibility to be exclusively on the demand- or supply side, because the power market offers a platform to exchange power. Moreover, the price of the day-ahead market is used to determine the realised price of financial futures, but is also used as reference price for contracts with hourly prices and to define different indexed prices.

The markets in each country grew with the increasing number of trading companies and the traded volume increased with each year. Each country formed their own bidding zone, which is a geographical area with fixed network constraints and one clearing price. Some bidding zones are interconnected and they obtain the same clearing price if the interconnections do not reach their maximum capacity, otherwise they have different clearing prices. The next step of the liberalisation of the energy sector was to couple these established bidding zones to increase competition and liquidity further, and to use the full capacity of the existing power grid and cross-borders.

The Nordic countries were pioneers for coupled power markets. Norway already introduced a power market in 1993 and Sweden joined the market in 1996 and rebranded the market to NordPool forming the first international European power market. Finland and Denmark joined in 1998 and 2000, respectively, and the three Baltic countries joined in 2010, 2012 and 2014.[6]

The coupling of day-ahead markets in continental Europe started in parallel between some neighboring countries. The Trilateral Coupling (TLC) formed a common day-ahead market and connected the markets of France, Belgium and the Netherlands in November 2006. In 2007, MIBEL started to couple Spain and Portugal, and in 2009 the coupling of the power markets of the Czech Republic and the Slovak Republic started. In 2010 Germany, Austria and Luxembourg joined the TLC and the Central Western Europe (CWE) market was founded. Moreover, the Interim Tight Volume Coupling (ITVC) between the CWE region and the

[6]See https://www.nordpoolgroup.com/About-us/History/

Nordic region was implemented. In 2011, Italy and Slovenia started their market coupling.[7]

In 2014 the new Price Coupling of Regions (PCR) project was launched and allowed to couple the Nordic regions with the CWE and MIBEL, and was called North-Western Europe (NWE) price coupling. New markets can be expanded easily, and Italy and Slovenia were included in 2015. The single price coupling algorithm is called PCR EUPHEMIA (acronym of Pan-European Hybrid Electricity Market Integration Algorithm) Committee [2019]. It calculates energy allocation, net positions and electricity prices across Europe, maximizing the overall welfare and increasing the transparency of the computation of prices and power flows resulting in net positions.

The clearing algorithm is able to handle all types of offers, which were introduced above, and calculates the clearing price for each bidding area respecting all given constraints and national particularities. The algorithm starts with a good first solution and continues to improve and increase the overall social welfare using a cutting plane optimization algorithm. The stopping criteria is a time limit or if the full branch and bound tree is explored. Hence, the algorithm takes a maximum 17 minutes to solve the large optimization problem. A large share of the European consumption is already coupled and 1531 TWh were cleared in 2020. However, most times, the installed cross-border capacities between European countries lead to different prices across Europe.

Europe decided to nominate at least one responsible for each bidding area to manage the coupling of the markets and introduced the so called Nominated Electricity Market Operators (NEMO). They receive bids from registered market participants for the day-ahead and intraday market. These bids have to be matched considering the available cross-border capacities. Finally, the resulting clearing prices have to be published and the position of each market participant has to be determined for the settlement. The NEMOs are: BSP, CROPEX, SEMOpx (EirGrid and SONI), EPEX, EXAA, GME, HEnEx, HUPX, IBEX, Nasdaq, Nord Pool, OMIE, OKTE, OPCOM, OTE, and TGE.

The European single day-ahead market coupling (SDAC) is still growing and the following countries are already coupled: Austria, Belgium, Czech Republic, Croatia, Denmark, Estonia, Finland, France, Germany, Hungary, Italy, Latvia, Lithuania, Luxembourg, the Netherlands, Norway, Poland, Portugal, Romania, Slovakia, Slovenia, Spain, Sweden and the UK.

[7]See https://www.entsoe.eu/network_codes/cacm/implementation/sdac/

1.3.3 Intraday Markets

Trading volumes of the day-ahead markets increased within the first years, however the market already closes at 12 am the day before physical delivery. The resulting lead-time (time from the last possibility to trade until physical delivery) reaches from 12 hours for product $H1$ to 36 hours for product $H24$. During this long lead time additional information arrives and allows to improve forecasts for production and demand. Intraday markets allow market participants to adjust the day-ahead profile until shortly before physical delivery.

Traders with a portfolio of uncontrollable renewable power sources, as photovoltaic plants and wind parks, usually sell the production forecast on the day-ahead market. The production forecast is continuously updated with new weather forecasts and real-time production measurements. The accuracy of the forecast is increasing towards physical delivery and the intraday market allows to trade the resulting forecast error. Similarly, updated temperature trends and real-time measurements of consumption units allow to update demand forecasts and to trade them on the intraday market.

Traders have to guarantee the production if the power plant was already scheduled in the day-ahead market or on the forward market with physical delivery. The missing power from unexpected outages can be balanced with other power plants of the same portfolio or by a purchase on the intraday market.

Another important element of the intraday market are flexible power plants, which provide the needed power to trade forecast errors. Prices on the intraday market are usually highly volatile during the trading session which offers the possibility to earn money with flexible power plants. Highly flexible power plants can offer their flexibility until shortly before physical delivery and less flexible power plants can still participate, but have to consider the longer activation time.

The day-ahead market offers exclusively hourly products. However, production forecasts are getting better a few hours before physical delivery and this also allows to predict power output with a finer granularity. While Italy exclusively offers hourly products, German power providers have the possibility to trade half-hourly and quarter-hourly products to schedule ramps for the increasing production of photovoltaic plants towards noon and decreasing production towards sunset.

Interestingly, two prevailing systems of intraday markets were developed in Europe. Most European countries implemented the intraday market as a continuous market. It allows market participants to submit offers during the trading session and offers are cleared immediately if the conditions are fulfilled. However, in Italy, Spain and Portugal sequential intraday auction were introduced to allow market participants to submit and modify offers, and the market is cleared seven times a day.

Europe has a strong desire to couple intraday markets, and therefore, they decided to use the continuous market design for the SIDC (Single Intra-Day Coupling).[8] Most European countries already offer a continuous intraday market and cross border capacities can be used to compensate power schedules shortly before physical delivery on a European level. Hence, Italy, Spain and Portugal already introduced a continuous market and reduced their historic intraday auction markets to three auctions.

The continuous intraday market is a pay-as-bid clearing, where the highest buy-prices and lowest sell-prices are served first. Each bidding area has its own limit order book, where submitted orders are collected. The SIDC requires a shared limit order book (SOB), a capacity module (CMM) and a shipping module (SM). The capacity module shows if cross-border capacity between countries is available, and provides the information to build the shared limit order book with orders from neighboring countries until the available volume. The available capacity in the CMM is updated if a cross-border clearing shows up and the available orders are updated in the SOB. All concluded trades are registered in the SM, which provide TSOs and NEMOs the set of cleared orders to calculate power flows between countries.

1.3.4 Balancing Markets

Power providers have to balance their position within each period (60 or 15 minutes) and the introduced markets can be used to sell or buy missing power. The last chance to integrate the latest information are intraday markets. The balancing provider takes over the role shortly before physical delivery to guarantee a secure power network by activating the needed flexible power to balance the system.

Balancing power has two directions. There are flexible generators which can increase their production and other generators which can reduce production. Power providers with unbalanced power schedules have to carry these balancing costs which are usually unfavorable compared to the corresponding costs for balancing the position on the day-ahead or intraday market. This forces traders to predict their consumption and production with small forecast errors. Moreover, exact power schedules are also good for the whole power grid to minimize the use of balancing power and to keep system costs low.

The European power market decided to harmonise the imbalance settlement periods of 15 minutes with regulation 2017/2195[9]. Germany already introduced a granularity of 15 minutes and offers market participants on the intraday market to trade products of 15 minutes. However, Italy still has an imbalance settlement period of 60 minutes. These settlement periods provide a rough structure to balance

[8] See https://www.entsoe.eu/network_codes/cacm/implementation/sidc/

[9] https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A02017R2195-20210315

the market, but balancing providers have to balance it in real time.

The balance of the power grid can be measured through the grid frequency. Europe decided to set the frequency of the power grid to 50 Hz and the balancing provider has to regulate balancing power to keep the frequency at a constant level. A lack of power is expressed as a frequency drop, while an excess of power leads to a frequency increase in the whole European power grid. The aim of the TSO is to keep the frequency within a narrow range around 50 Hz.

Balancing providers do not own power plants to balance the system. Owners of flexible power plants or flexible consumption are able to participate at the balancing market, offering the TSO their capacity. Flexible power is usually more expensive and the TSO can reserve balancing power with different levels of flexibility on the market. The balancing provider can accept the power with the lowest marginal costs in combination with the location of the power plant in the power grid. There are difference types of balancing power:[10]

- **Primary balancing power** is provided by highly flexible power plants. This flexible power is used to cover periods until 15 minutes and has to be activated completely within 30 seconds. These power plants are controlled automatically within the European power grid.

- **Secondary balancing power** is provided by medium flexible power plants. The reserved power has to be activated completely within 5 minutes. This power is managed by the corresponding TSO, which has the possibility to active the power to balance the corresponding control area.

- **Tertiary balancing power** is provided by less flexible power plants. Power has to be activated within 15 minutes and the merit-order-list is used to decide the order. This power is used to cover periods from 15 minutes up to several hours.

1.4 European Power Markets

Each country has slightly different power markets, but satisfies the European requirements. In this section, the Italian and German power markets are introduced. Both countries are interesting for this dissertation, because they developed different designs of intraday markets: Italy used sequential auction markets, while Germany introduced a continuous market design.

[10] See https://www.regelleistung.net/ext/static/technical?lang=en

1.4.1 The Italian Power Market

The company ENEL was responsible for the whole power supply in Italy since 1962, when the Italian energy sector was nationalized, until the liberalization of the sector. Beside ENEL, there were just a few cities with their own local power grids and some companies which covered their demand with their own production. The liberalization of the energy sector in Italy started in 1999 and forced ENEL to found the two companies Terna and e-distribuzione to launch the unboundling process. While Terna became the unique TSO of Italy, e-distribuzione became the largest distribution company in Italy, but ENEL is still the largest Italian power company with production plants and consumption Jannuzzi [2021].

Natural gas is the most important energy source in Italy. Italy has no huge coal resources and atomic power is forbidden by law as a result of a referendum after the Chernobyl disaster in 1986. The share of renewable energy sources increased during the last years with photovoltaic plants to benefit from the Mediterranean climate and wind parks in the South. Compared to other European countries, Italy has low CO_2 emissions due to renewable shares and gas plants. As a consequence, Italy has high prices due to the high marginal costs of power produced by natural gas and imports additional energy from neighboring countries.

The Italian power market IPEX (Italian Power EXchange) is managed by GME (Gestore Mercati Energetici) and was launched in April 2004 and is fully operational since January 2005. Initially, there were two auctions: the day-ahead market (MGP - Mercato del Giorno Prima) and the MA (Mercato Aggiustamento) followed shortly after the MGP closed to allow market participants to adjust their power schedules. In November 2009, the MA auction was replaced with two auctions named MI1 and MI2 (Mercato Infragiornaliero) with succeeding closing times allowing market participants to adjust their power schedules. In January 2012, two additional intraday markets MI3 and MI4 were introduced and allowed trading energy of the same day. In 2015 the closing of the MGP was set to 12 am to allow the coupling with the European day-ahead market. The opening times of the intraday auctions were also adapted to the new closing time and a new session MI5 was added. In 2017 the intraday market was expanded to a total of 7 auctions called MI1 - MI7 ARERA [2019]. Due to the historic development, the Italian intraday market was managed with intraday auctions.

Europe decided to introduce an European continuous intraday market, and therefore, Italy introduced a continuous market on 21/09/2021 and the seven intraday auctions were reduced to three auctions MI-A1, MI-A2 and MI-A3. In the following, I will describe the market design before this change, because the old market design is analysed in a part of this thesis.

Italian market operators are able to trade financial power futures on the EEX and futures with physical fulfillment on the market MTE offered by GME. The

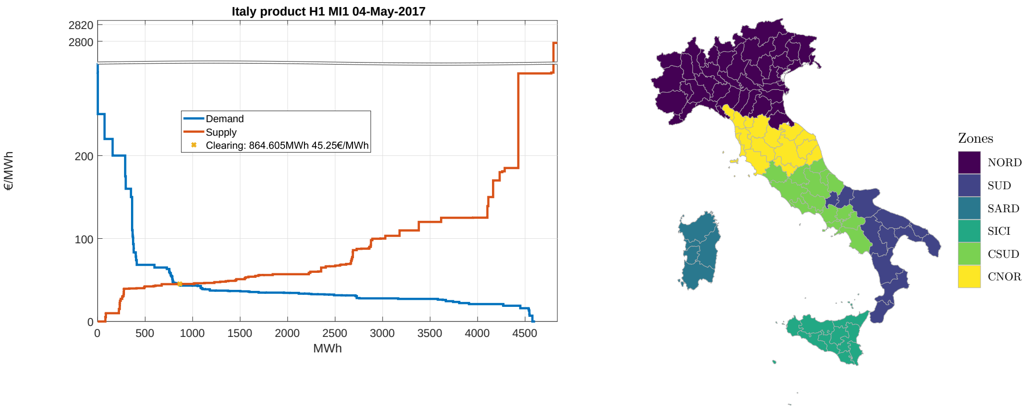

Figure 1: Auction clearing and Italian market zones.

Italian forward electricity market allows to trade base and peak-load with monthly, quarterly and yearly delivery periods on a continuous basis.

Italy has a challenging geographic shape to build a power grid, because it is surrounded by the sea, and therefore, the main power lines have to link the South with the North. Hence, the Italian power grid consists of six bidding zones (NORD, CNORD, CSUD, SUD, SICI and SARD) as shown in Figure 1, four poles of limited production (FOGN, BRNN, ROSN and PRGP) and has connections to neighbouring countries (France, Switzerland, Austria, Slovenia, Corsica (France), Malta and Greece). On 01/01/2021 Calabria was added as additional market zone. Power traders have to balance power of production plants and consumers in each single bidding zone, and are forced to use the intraday market to compensate forecast-errors, i.e. they might need to sell power in one zone and buy power in another zone.

The day-ahead market MPG (Mercato del Giorno Prima), which closes at noon the day before physical delivery, gives power providers the possibility to exchange hourly products. The Italian day-ahead market is special, because Italy consists of different bidding zones and power is sold at the zonal price where the production plant is located. In contrast, consumption (pumped storage excluded) can buy power in each zone at one hourly price called PUN (Prezzo Unico Nazionale) which is the volume weighted price of each zone. This guarantees that no region has economic advantages.

The seven Italian intraday auctions, called MI (mercato infragiornaliero), have just two different opening times, namely 12:55 and 17:30. Moreover, we define the lead-time of a product as the time between the last possibility to trade and the physical delivery. Table 1 summarizes the characteristics of the Italian markets.

Market	**Products**	**From**	**To**	**Results**	**LT in hours**
MI1	$H1-H24$	12:55 (d-1)	15:00 (d-1)	15:30 (d-1)	-
MI2	$H1-H24$	12:55 (d-1)	16:30 (d-1)	17:00 (d-1)	$7\frac{1}{2}$ up to $10\frac{1}{2}$
MI3	$H5-H24$	17:30 (d-1)	23:45 (d-1)	00:15 (d)	$4\frac{1}{4}$ up to $7\frac{1}{4}$
MI4	$H9-H24$	17:30 (d-1)	3:45 (d)	4:15 (d)	$4\frac{1}{4}$ up to $7\frac{1}{4}$
MI5	$H13-H24$	17:30 (d-1)	7:45 (d)	8:15 (d)	$4\frac{1}{4}$ up to $7\frac{1}{4}$
MI6	$H17-H24$	17:30 (d-1)	11:15 (d)	11:45 (d)	$4\frac{3}{4}$ up to $7\frac{3}{4}$
MI7	$H21-H24$	17:30 (d-1)	15:45 (d)	16:15 (d)	$4\frac{1}{4}$ up to $7\frac{1}{4}$

Table 1: Operating times of the Italian intraday market, where LT stands for lead-time.

As can be seen the lead times range from 4.25 to 10.5, which is relatively long and makes it difficult to include the latest production-updates of renewable energy sources as solar-plants and wind-parks.

Registered power providers are able to submit offers for their production plants and consumption units to the Italian power markets. The offers consists of the side (buy or sell), a quantity and a price. The policy of the Italian power market is very transparent and all submitted offers are available on the web.

The Italian MI markets are organized as uniform price auctions. Market participants submit offers during the market operation time specifying price and quantity. The clearing-algorithm is launched after the market closes. All sell and buy offers are ordered by their price, and define in combination with the corresponding cumulative volumes the *supply-curve* for the sell-offers and the *demand-curve* for the buy-offers, which are visualised in *Figure* 1. The intersection between the two curves gives the cleared volume and the clearing price of the market session. The resulting price is the uniform price for all accepted offers and all offers beyond the clearing price are rejected.

Power line capacities are added as constraints to the clearing-algorithm. If the transmission limit between two areas reached its maximum, the missing or extra power in the corresponding zone can not be compensated by another zone and we obtain different prices in the two areas. This phenomenon often happens in the Italian intraday markets.

Terna is the responsible TSO for the whole Italian power grid. It has the possibility to procure the resources on the Dispatching Services Market (MSD), which are needed to manage and control the system. Terna acts as the central counter party and can accept offers through pay-as-bid. There is a single sitting for the six markets, MSD1 - MSD6, where market participants can submit orders. Market participants are informed sequentially if their bids were accepted after each market is closed. Moreover, there exists a continuous balancing market, where market participants can submit offers until 60 minutes before delivery.[11]

1.4.2 The German Power Market

The German power grids were managed by a few regional power providers. These regional providers were responsible for generation, transmission and distribution of their area and customers were forced to sign a contract with the corresponding company. The liberalisation of the German electricity sector started in 1998.

In the past, Germany was using mainly coal, gas and atomic power plants to cover electricity demand. A few years ago, Germany decided to start the transition form conventional to renewable energy sources, and wind parks already produced the largest share as single resource in 2020. Recently, Germany planned to decrease the number of coal and atomic power within the next years. The huge demand and large share of renewable energy sources led to developed power markets with high trading volumes.

In 2000 the LPX (Leipzig Power Exchange) was founded in cooperation with NordPool to introduce a market place for hourly products. Additionally, EEX (European Energy Exchange) introduced a spot market in 2000 and a market for forward contracts in 2001. The market area was too small for two power markets, and therefore, the markets were merged to the EEX in 2002. In September 2008, the EPEX Spot was founded to merge the power spot markets of the French Powernext and the EEX. It manages trading in the German-Austria, French and Swiss market zones since 2009. EEX provides still the forward electricity market to trade electric power of the following six years ahead. In 2015 the APX Group (spot markets for electricity in the Netherlands, UK and Belgium) was integrated into the EPEX Spot and formed a power exchange for Central Western Europe (CWE) and the UK.[12]

In January 2009 EPEX Spot introduced the intraday market for hourly products with a continuous market design. In January 2011 the German continuous market was complemented with additional quarter hourly products. In October 2012 Austria was integrated and formed the German-Austrian intraday market, which was existing until they were separated in 2018. In December 2014 an additional auction for quarter hourly products was introduced in Germany to set a price-level of quarter hourly products. On 30/03/2017 half hourly products were introduced.

Nowadays, EPEX Spot offers a platform for hourly, half-hourly and quarter-hourly products in Germany. The day-ahead market is used exclusively for hourly products, closes at noon, and is followed by an intraday auction for quarter hourly products and a continuous intraday market for all three product types.

[11] See https://www.mercatoelettrico.org/en/mercati/mercatoelettrico/mpe.aspx

[12] See https://www.eex-group.com/en/about/milestones

Product	From	To (within Germany)	To (control area)
Hour	15:00 (d-1)	30 minutes before delivery	5 minutes before delivery
$\frac{1}{2}$ hour	15:30 (d-1)	30 minutes before delivery	5 minutes before delivery
$\frac{1}{4}$ hour	16:00 (d-1)	30 minutes before delivery	5 minutes before delivery

Table 2: Operating times of the German continuous intraday market.

Quarter-hourly products were introduced to compensate ramping-times. Less flexible power plants need some time to increase or decrease production, but also the typical production of a photovoltaic plant is increasing towards noon and decreasing afterwards. EPEX Spot introduced an auction for quarter-hourly products, which closes at 3 pm. The *intraday auction* is organised as a uniform price auction, where market participants submit offers indicating price and quantity, and the intersection of increasing supply- and decreasing demand-curve determines the clearing price. This auction allows traders to trade ramping-times and to obtain a price signal for the following continuous intraday market for quarter hourly products, which opens at 4 pm Neuhoff et al. [2016].

Table 2 summarizes the characteristics of the German continuous intraday market. The continuous intraday market for hourly products opens at 3 pm to provide a further possibility to modify the current production-profile with a hourly resolution. The continuous intraday market for quarter-hourly products opens at 4 pm and half-hourly products can be traded from 3:30 pm. The lead-time of the continuous market was reduced from 45 to 30 minutes on 16/06/2015. On 14/06/2017 the lead-time was further reduced to 30 minutes (or 5 minutes within the control area), which is very useful to integrate fluctuating renewable energy sources.

A submitted order in the continuous market is cleared immediately if the price is better than the best price of an order in the limit order book, otherwise the order is stored in the limit order book. All orders are firstly ordered by the price and secondly by the start validity time, and therefore, the clearing price is the price of the first entered order and the cleared volume is the minimum volume of the two orders. *Figure* 2 shows the best available buy and sell order. A clearing shows up at each intersection of the two curves.

Due to the historic regional monopolies, the German area is still divided in four balancing areas managed by four TSOs. The four German TSOs are Amprion, 50Hertz Transmission, TransnetBW and Tennet as shown in the right plot of Figure 2. However, all German TSOs form a union, where positive and negative balancing power can be exchanged. There are regular open calls for tenders, where balancing power can be offered to the four TSOs organised as pay-as-bid.

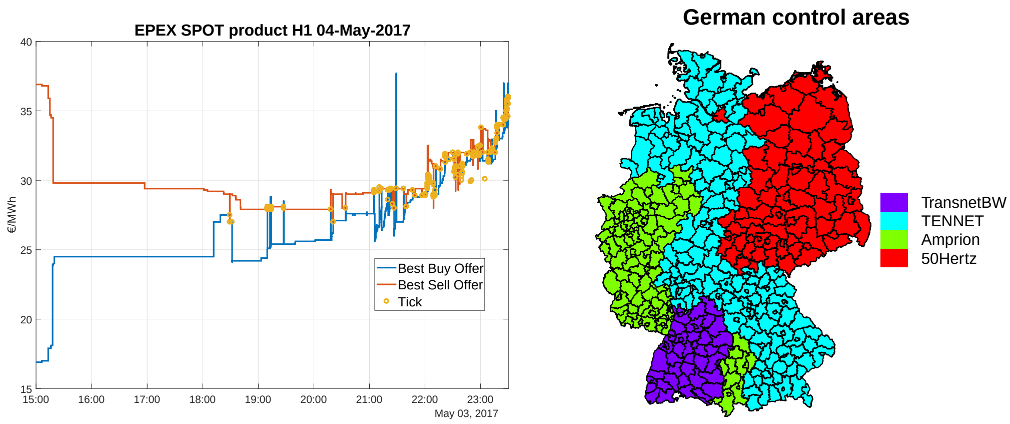

Figure 2: Continuous clearing and German control areas.

1.5 Limit Order Book

Limit order books are used by exchanges with a continuous market design to list the orders for assets like power contracts, stocks and currencies. Traders can submit limit orders indicating side (buy or sell), price, quantity and validity time to the limit order book. A new order is immediately cleared if it can be matched against an existing order of the limit order book. Otherwise, the new order will not be cleared, but will remain in the limit order book until the end of the provided validity time or if another new order leads to the clearing.

1.5.1 Functionality and Order Types

The limit order book is empty at the beginning of each trading session. The basic element of the limit order book is a *limit order*, which can be submitted to the limit order book during the trading session. Traders can accept the price of available limit orders or they can submit a limit order indicating their desired price. I will describe the basic functionality of the EPEX Spot continuous intraday market with the following example.

Example 1. *The trading session starts at* 3 *pm with an empty limit order book. At that moment no trade is possible. However, trader* 1 *has the intention to sell a volume* V_1 *MWh of the contract* $H12$ *at a fixed price* P_1 *and submits the order to the limit order book. Now, it would be possible to buy the volume* V_1 *immediately at the price* P_1*. Trader* 2 *has the intention to buy the volume* V_2 *MWh of the contract* $H12$ *and can choose to accept the price* P_1 *until the available volume* V_1 *MWh or to submit a limit order with a lower price* P_2*. Trader* 2 *decides to submit the limit order with the lower price* P_2*. Hence, the new situation of the limit order book* *consists of two limit orders providing a price* P_1 *for buying and* P_2 *for selling the contract* $H12$*. If trader* 3 *would also like to buy a volume* V_3 *MWh of* $H12$ *with a better price* P_3 *than* P_2*, but lower than* P_1*, she has to submit an additional limit order. The resulting situation allows to buy the product* $H12$ *at* P_1 *and sell at* P_3*.*

As described in the example, the limit order book is growing with additional sell- and buy-limit orders. Trades show up if a new limit order with a better price than the best available limit order will be submitted. The cleared limit order from the limit order book defines the clearing price. Hence, the situation of the limit order book is continuously changing with the submission of new limit orders, clearings or cancelling of valid limit orders. Both sides have to be sorted at each modification of the limit order book to determine the best sell- and buy-price.

Submitted limit orders remain in the limit order book until the end-validity time is reached or they are cleared. However, market participants have the possibility to deactivate submitted orders and re-activate them later. Deactivated orders are no longer considered by the clearing algorithm. Market participants can also modify the price of a limit order, or they can cancel an order before the end-validity time is reached.

One special feature of the continuous market design is the continuous monitoring of the best prices, allowing traders to make decisions based on this information. Some markets show exclusively the best prices with the corresponding volumes and other show the best n limit orders of each side. Hence, the continuous market design is transparent and allows to trade immediately - in contrast to the auction market, where traders submit their orders without seeing the other submitted orders and the market is cleared at the closing time.

Market participants can also decide to place a *market order* indicating a fixed volume and the side, which will be matched against the best available order of the limit order book. Market orders can be interpreted as limit orders with the maximum price. These orders are immediately executed, but traders have to be aware that they accept the resulting market price.

One can distinguish between passive and active traders. Passive traders prefer to submit limit orders to the limit order book and wait until their order will be cleared. Obviously, there is the risk that the order will not be cleared. In contrast, active traders submit orders with an immediate clearing and accept the current market price. Hence, they submit market orders or limit orders with a price at or beyond the best price.

Beside limit orders and market orders, the EPEX Spot Continuous intraday market allows additional order types Martin and Otterson [2018]:

- IOC (Immediate Or Cancel): order is immediately executed or cancelled

- Market Sweep Orders: user-defined block orders with the execution restriction IOC (Immediate Or Cancel) are executed immediately and, as far as possible, against respective single-contract orders
- Iceberg or Hidden-Quantity Orders: are large volume orders and are entered sequentially piece-by-piece in the order book
- FOK (Fill Or Kill): immediately entirely executed or entirely cancelled
- AON (All Or Non): order is executed completely or not at all

Another important product type for power markets are block orders, which allow to trade blocks of products, where all or none of the single products are accepted or rejected. Block orders are special for power markets and are important for less flexible power plants, which can not be switched on or off just for one single product.

The continuous intraday market consists of many limit order books, because there is one limit order book for each product. Hence, at each day, there are 24 limit order books for hourly products, 48 for half-hourly products, 96 for quarter-hourly products, but also for block-orders.

1.5.2 Description of the Data

The limit order book of the German continuous intraday market is not publicly available and has to be purchased at EPEX. The company offers limit order books for each country and the German limit order books of the years 2015 - 2018 were available for this research.

The purchased data-set consists of all submitted orders to the German bidding area. In a part of our available period, Austria, Luxembourg and Germany formed a unique bidding area, but the limit order book contains exclusively orders submitted to the German bidding area. If a German order was matched for example with an Austrian order, just one of the two cleared orders is contained in the purchased data.

The limit order books consist of a few gigabyte of data and the amount of submitted orders increased with each year, because activity increased on the continuous intraday market, but also half hourly products were introduced and the lead-time decreased from 45 minutes to 5 minutes before physical delivery. Moreover, the XBID project was successfully launched and algorithmic trading is getting a powerfull tool for traders, which increased activity further.

The limit order book is available as csv-files and each row of the order book describes a new, modified or executed order containing information about:

- Instrument Type: hour, half hour or quarter hour

- Delivery Instrument: specific product
- Delivery Date: fixed day
- Start Validity Date: offer is valid from this time, which is precise up to milliseconds
- End Validity Date: offer is valid until this time, which is precise up to milliseconds. This time will be replaced by the time of an eventual cancelling, execution or partially execution event.
- Cancelling Date: a trader is able to cancel an offer before the initially given validity date is reached or after the offer was partially executed.
- Is Executed: not, fully or partially executed
- Status: active or not active
- Side: sell or buy
- Price: the price ranges from -9999€/MWh to 9999€/MWh with a granularity of 0.01€/MWh
- Executed Price: matched price
- Volume: total volume, which is precise up to one decimal point
- Executed Volume: matched volume
- Order ID: ID to identify each order
- Initial ID: ID which was assigned to the order when it was entered for the very first time
- Parent ID: order number of the previously modified order

Moreover, observed clearings are also contained in the data as logs showing the exact clearing time, the cleared volume and clearing price. This information is provided by the additional column *Is Executed* and allows to extract all trades with the exact trading time. The *transaction data* shows the same clearings, but also the matched counterpart even if it was an order from abroad. This data can be used to analyse trading volumes, trading activities and the distinction between buy- and sell-orders.

1.5.3 Basic Technical Concepts

At each time, all buy-orders have to be sorted by their increasing price and all sell-orders have to sorted by their decreasing price. The start validity time is considered as second order criterion if offers have the same price. The third order criterion is the submission-time of the limit order. The two sorted sides are called *order-stacks*, where the best sell-order is called *best ask* and the best buy-order is called *best bid*. Clearings follow exactly the order of the specific order-stack and guarantee the best available price.

More formally, all buy- and sell-orders of the limit order book are sorted at each time t by their price $\cdots < P^t_{-2} < P^t_{-1} < P^t_{-0} < P^t_0 < P^t_1 < P^t_2 < \cdots$, where P^t_{-0} is best bid-price and P^t_0 is the best ask-price. Hence, the set

$$S(t) = \{(P^t_0, V^t_0), (P^t_1, V^t_1), \cdots, (P^t_{N_S}, V^t_{N_S})\}$$

is the sell-stack with N_S orders and

$$B(t) = \{(P^t_{-0}, V^t_{-0}), (P^t_{-1}, V^t_{-1}), \cdots, (P^t_{-N_B}, V^t_{-N_B})\}$$

is the corresponding buy-stack of the limit order book at time t with N_B orders.

An important quantity in the analysis of the limit order book is the *bid-ask-spread*. It can be defined at each time t by the best bid price P^t_{-0} and the best ask price P^t_0:

$$BAS(t) = P^t_0 - P^t_{-0}.$$

If traders would like to buy a fixed quantity smaller than the minimum of the two best orders and sell it immediately after, the trader has to pay the bid-ask spread times the traded volume. Moreover, an active trader has to pay the additional bid-ask spread compared to a passive trader. The bid-ask-spread from the limit order book situation in the left plot of Figure 3 can be calculated as follows:

$$BAS = 30.5€/MWh - 27.5€/MWh = 3€/MWh.$$

The bid-ask spread just considers the prices of the best offers and ignores the corresponding volumes. The price of each available limit order just covers the provided volume, and clearing larger volumes requires to accept limit orders with prices beyond the best price. I will describe this concept in the following example.

Example 2. *The left plot of Figure 3 shows the situation of the order book of the contract* $H12$ *and delivery date* 20/02/2015 *at the fixed time* 19/02/2015 16:21:51.293. *The four best offers of the buy-stack are shown in red and the sell-stack in blue. The best bid is a order with* 2.8 *MWh for* 27.5 *€/MWh and the best ask is the order with* 2 *MWh for* 30.5*€/MWh. The best offers of each order stack are*

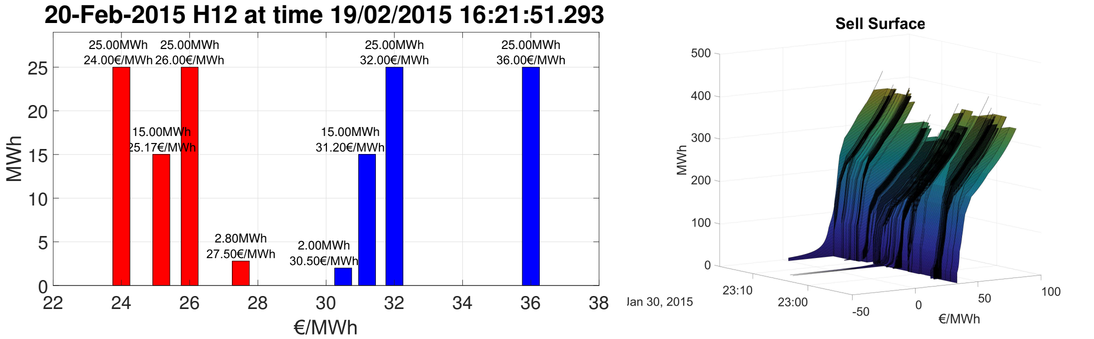

(a) Order book situation at a fixed time.

(b) Temporal development of the sell-stack.

Figure 3: Situation of the limit order book.

also visible to traders in real time. Traders can trade the corresponding volume at the provided price, but they have to be aware that trading larger volumes lead to worse prices. If for example a trader would like to buy $4MWh$*, the first* $2MWh$ *can be bought at* $30.5€/MWh$ *and the second* $2MWh$ *can be bought at* $31.2€/MWh$ *and this leads to an average price of* $30.85€/MWh$*.*

The example shows that the bid-ask spread can be misleading if the volume of the best offers is small. Trading larger volumes causes larger price impacts, and therefore, the *liquidity spread* for a volume V at time t can be defined as follows:

$$BAS(t,V) = \frac{\sum_{i=0}^{N_S} P_i^t \bar{V}_i^t}{V} - \frac{\sum_{i=0}^{N_B} P_{-i}^t \bar{V}_{-i}^t}{V},$$

where $\bar{V}_i^t = max(0, min(V_i - \sum_{k=0}^{i-1} V_i))$ and V_i is the corresponding volume to P_i of all sorted orders of the sell-side, and analogously for the buy-side. The liquidity spread of 4 MWh of the limit order book situation in the left plot of Figure 3 can be calculated as follows:

$$\begin{aligned} BAS(t,4) = & \frac{2MWh \times 30.5€/MWh + 2MWh \times 31.2€/MWh}{4MWh} - \\ & - \frac{2.8MWh \times 27.5€/MWh + 1.2MWh \times 26€/MWh}{4MWh} = \\ & = 30.85€/MWh - 27.05€/MWh = 3.8€/MWh. \end{aligned}$$

The volume of a market order has a strong impact on the final cost, because the liquidity spread for 2 MWh ($3.8€/MWh$) is much higher than the basic bid-ask-spread ($3€/MWh$) at the fixed market situation. Hence, one has always to check the orders beyond the best limit order.

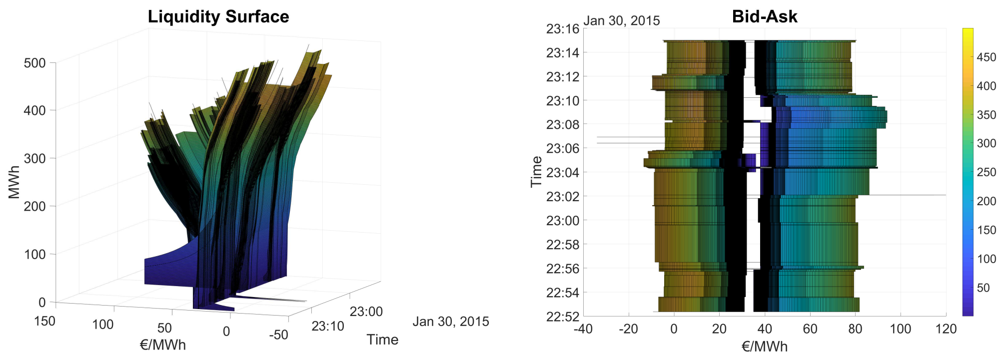

(a) Sell-Side & Buy-Side.

(b) Bid Ask.

Figure 4: Interaction of the two objects.

Another important property of a market is the depth of the sell- and buy-stack at each time t. The depth shows the maximum available volume of each side. The two quantities in MWh can be defined as follows

$$D_S(t) = \sum_{i=0}^{N_S} V_i^t \quad \text{and} \quad D_B(t) = \sum_{i=0}^{N_B} V_{-i}^t.$$

The left plot of Figure 3 shows the situation of the limit order book at a fixed time. However, the limit order book changes with each new order submission or if an end-validity time of an active order is reached. The sorting of the two order stacks has to be launched at each modification of the limit order book and the depth changes. The left plot of Figure 3 shows the full sorted sell order stack along trading time. The best sell-offers are located on the bottom. The spikes on the bottom are limit orders with very low prices to enforce a clearing with a limit order of the buy-stack and these orders are immediately cleared and cancelled from the sell-stack. The height of the surface shows the depth of the sell-stack.

The buy-side can be visualised analogously and the left plot in Figure 4 shows the two surfaces. There are a few peaks which lead to clearings, but the two surfaces are separated beside these punctual intersections. Since all orders are sorted, the two surfaces are increasing and trading larger volumes leads to worse prices.

The right plot of Figure 4 shows the same surface as the plot on the left, but as projection from above. The plot shows the prices for selling and buying power at each moment along trading time. The horizontal length of the white space along trading time indicates the bid-ask-spread at each time, while the length of each surface indicates the depth of each side of the limit order book.

1.5.4 Implemented Clearing Algorithms

The clearing algorithm of the European intraday market is slightly more complex than a clearing algorithm for continuous traded financial assets due to block orders and cross-border trading. The limit order books of different countries are separated, but the shared limit order book of one specific country considers cross-border capacities and shows the best orders of the neighboring countries until the available cross-border capacity. This allows traders to match an order from abroad.

Unfortunately, limit orders from neighboring countries are not contained in the available limit order book. Nevertheless, these orders are needed to recalculate the observed clearings. I implemented an algorithm which matches buy- and sell-orders of executed trades. The algorithm adds a new order if there is a missing counter part, because these missing orders can be identified as missing orders from abroad. Unfortunately, the exact information of the validity time and the full volume is not available, but the identified order was for sure available in the limit order book at the clearing time.

Order types are not indicated in the purchased limit order book. However, there is a possibility to identify iceberg-orders by analysing the log information of observed clearings. Each limit order can be identified by the order ID. A limit order can be classified as iceberg-order if the order is fully cleared and afterwards appears again in the limit order book with the initial volume. This method allows to classify cleared volumes of iceberg-orders, but other order types as IOC or FOK can not be identified.

This dissertations analyses the purchased limit order book in different manners. I implemented one code to analyse the available data at each modification of the limit order book. The second code is a replication of the clearing algorithm and allows to evaluate trading strategies, where the impact of additional submitted orders is considered. Both implemented codes use the data of the limit order book and sort the buy- and sell-orders at each modification of the limit order book to provide the buy- and sell-stack at each time of the trading session. The order ID is an additional sorting criterion if multiple orders have the same price and start validity time.

The code to analyse the limit order book is implemented in *matlab*. The script allowed me to analyse properties of the market as the cleared volume, the exact clearing times of observed trades, the distribution of entered offers and clearings along trading time, the bid-ask-spread, and the offers beyond the bid-ask-spread. The buy- and sell-stacks are obtained by considering all active limit orders at each modification of the limit order book.

The second code is able to test trading strategies and was implemented in *java*. Java has the *TreeMap* object, which is efficient to handle the ordering of the order stacks after each modification of the limit order book. The implemented code is able to recalculate the observed clearings. The buy- and sell-stack are updated at the start and end-validity times of the submitted limit orders and logs of observed clearings are ignored. Hence, this algorithm allows to add additional clearings, which have an impact on the current buy- and sell-stack, but also affects the future evolution of the limit order book.

I also implemented the clearing algorithm of the Italian intraday auction and recalculated all observed clearings using the public available orders. Moreover, I expanded the implemented *java* code of the continuous clearing to simulate an auction clearing using the data from the limit order book.

There are several problems with the purchased limit order book. The data is suited to recalculate all observed clearings, but the initial end-validity time of the submitted offers are overwritten if the order was cleared or partially cleared. Hence, this information is misleading, because it is not possible to distinguish between an overwritten end-validity time and a real one. Moreover, the restriction to one specific country becomes a problem with increasing cross-border trades, because one would have to purchase all limit order books to understand the interactions between countries.

1.6 Analysis of the Limit Order Book

The public available data of the EPEX Spot Continuous Intraday market are reduced to 9 statistics for each product (Low, High, Last, Weight. Avg., Index, ID_3-Price, ID_1-Price, Buy Vol and Sell Vol). Data containing more details has to be purchases from EPEX. Most published studies were based on the public available data of the continuous intraday market and made analyses of the published statistics Hagemann and Weber [2015], Hagemann [2013]. Other studies as Janke and Steinke [2019], Hagemann and Weber [2013] used the *transaction data* to include the distribution of the clearings along trading time. However, just a few studies Balardy [2022], Martin and Otterson [2018], Kath and Ziel [2020], Glas et al. [2020], Kiesel and Paraschiv [2017] used the whole limit order book for their analysis as it is done in this dissertation.

1.6.1 Along the Trading Session

The data of the limit order book can be used to analyse the order book situation at each time. All registered clearings, the best bid and best ask of a fixed product along trading time are plotted in Figure 5. The bid-ask spread is very large at the beginning of the session, but decreases towards the end of the trading session. This plot just considers the best orders and ignores orders beyond the best bid and best ask. On this level, activity at the beginning of the trading session is rather low as can be observed by the constant price levels on both sides for a few hours.

However, activity on the market increases towards the end of the trading session with more clearings.

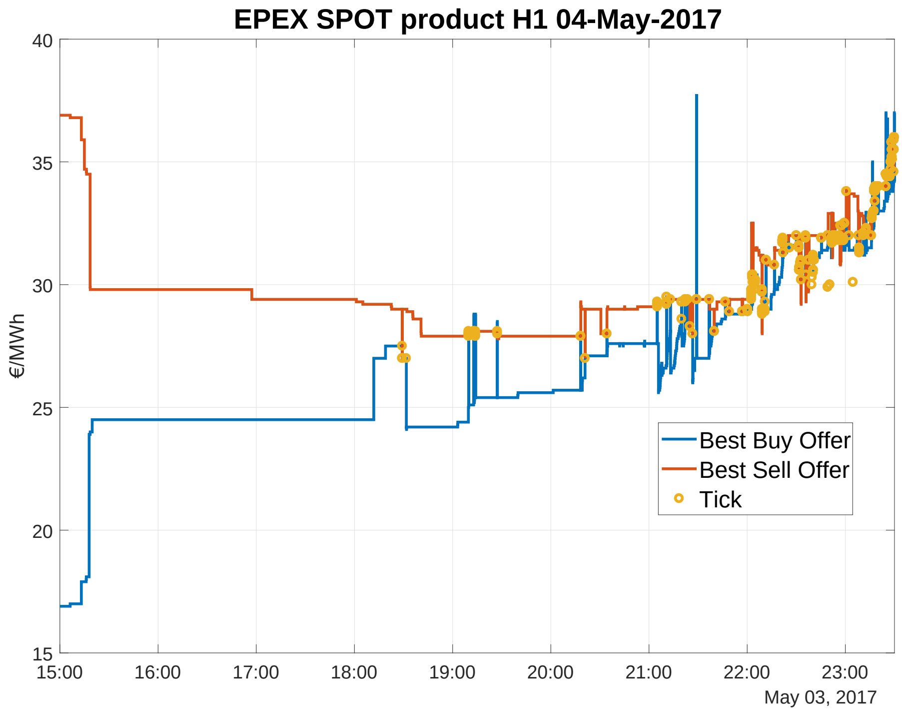

Figure 5: Clearings.

The continuous market opens just 3 hours after the day-ahead market closes and traders might place limit orders of flexible power plants, which were not accepted on the day-ahead market. Moreover, traders of renewable energy sources obtain updates of the production forecast, but there is no significant update a few hours after the day-ahead market closes. Physical delivery is still far away and there is a lot of uncertainty in the power forecasts, which decrease towards physical delivery. Hence, the strategy of traders of renewable energy sources is to wait until the accuracy of the power forecast increases to avoid buying and selling a quantity multiple times.

Most stock markets also use the continuous market design to trade shares and developed methods can be applied on the power market. However, the markets are slightly different. Financial markets have limited trading times, while the power market is open 24/7. Another difference are the traded products, because products of the financial market remain the same at each trading day, while products on the power market expire shortly before physical delivery and new contracts of the following day are tradable starting from the next trading day.

1.6.2 Market Properties

The markets for hourly products open at 3 pm and close 5 minutes before physical delivery. I define the activity of the market by the amount of inserted limit orders in a 15 minutes interval. In the upper heat-map of Figure 6, the activity of hourly products of the German LOB between 16/07/2015 and 31/12/2015 is evaluated. At the beginning of the trading session a few orders are submitted. This initial phase is followed by a period with low activity until a few hours before physical delivery. About $5-6$ hourly before the closing of the market, activity starts to increase and reaches the maximum around the closing of the market.

In the lower heat-map of Figure 6, the observed trading volumes within each 15 minutes intervals are shown. In contrast to the activity, just a small volume is cleared at the opening of the market. However, similar as the activity, trading volumes increase towards the end of the trading session. Hence, there is no much interest in the continuous intraday market after the market opens, but activity and the intention to trade power starts around $5-6$ hours before the market closes. This can probably be explained with the increasing accuracy of production forecasts of renewable energy sources and the intention to be balanced to avoid balancing costs.

Another interesting property of the limit order book is the depth, which indicates the available volume of each side. I calculated the mean available volume for each interval of 15 minutes over the period 17/06/2015 - 31/12/2015 for each product and plotted the mean in Figure 7. The available volume is strongly increasing after the market opens and reaches its maximum a few hours before physical delivery. At the same time trading volumes start to increase leading to a weakly decreasing limit order book depth towards the end of the trading session.

1.7 Liquidity

Liquidity is an important market property for power markets, but also for financial markets in general. Put simply, liquidity describes the possibility to quickly buy or sell an asset without affecting the market price. However, it is challenging to analyse liquidity quantitatively, because it is multidimensional. Kyle [1985] determined the following four dimensions:

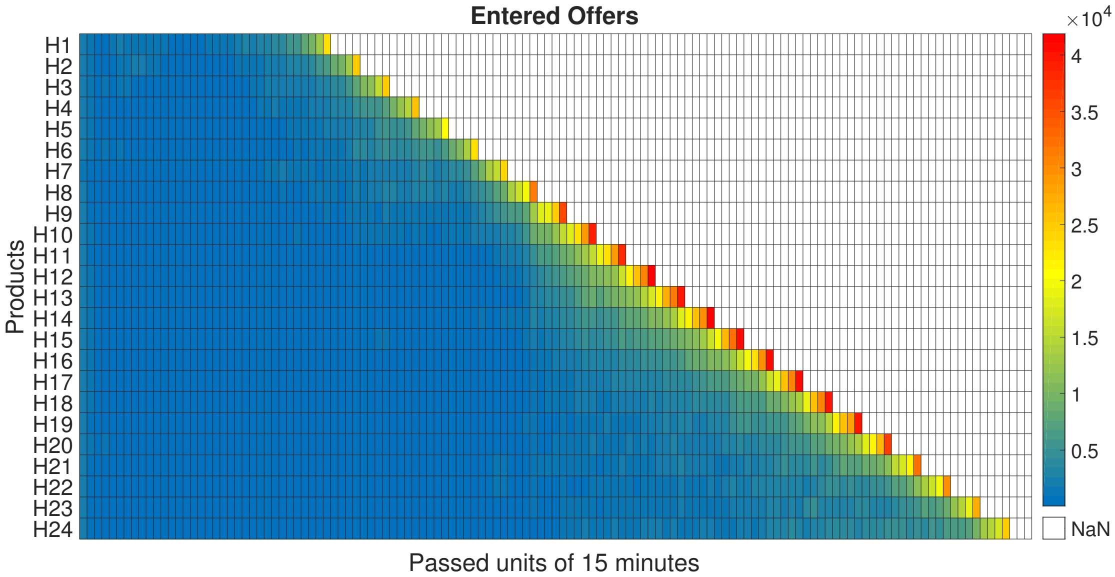

(a) New offers.

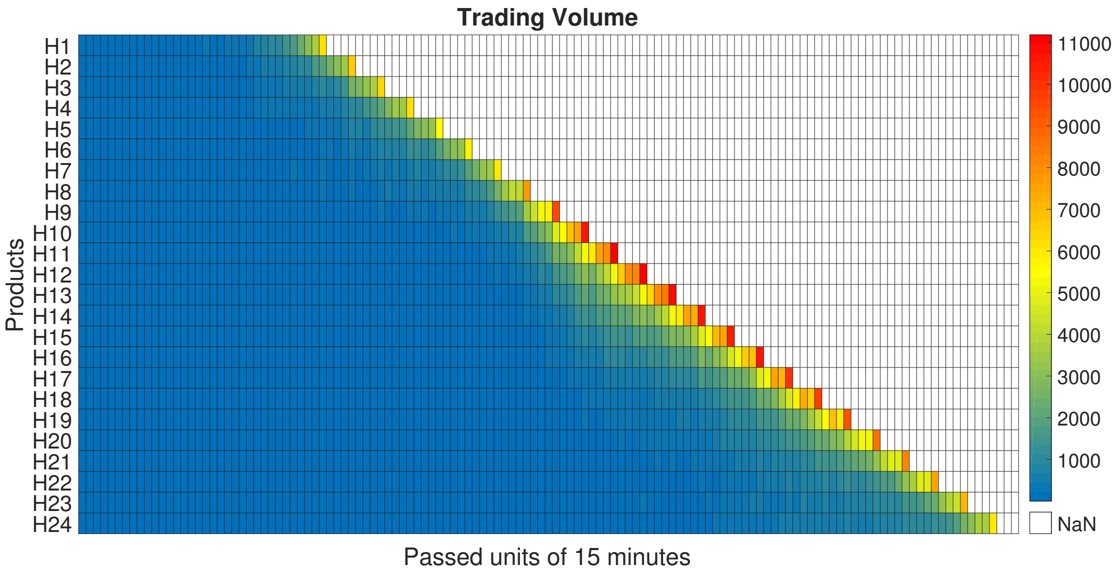

(b) Registered trades.

Figure 6: Activity in the limit order book of hourly products.

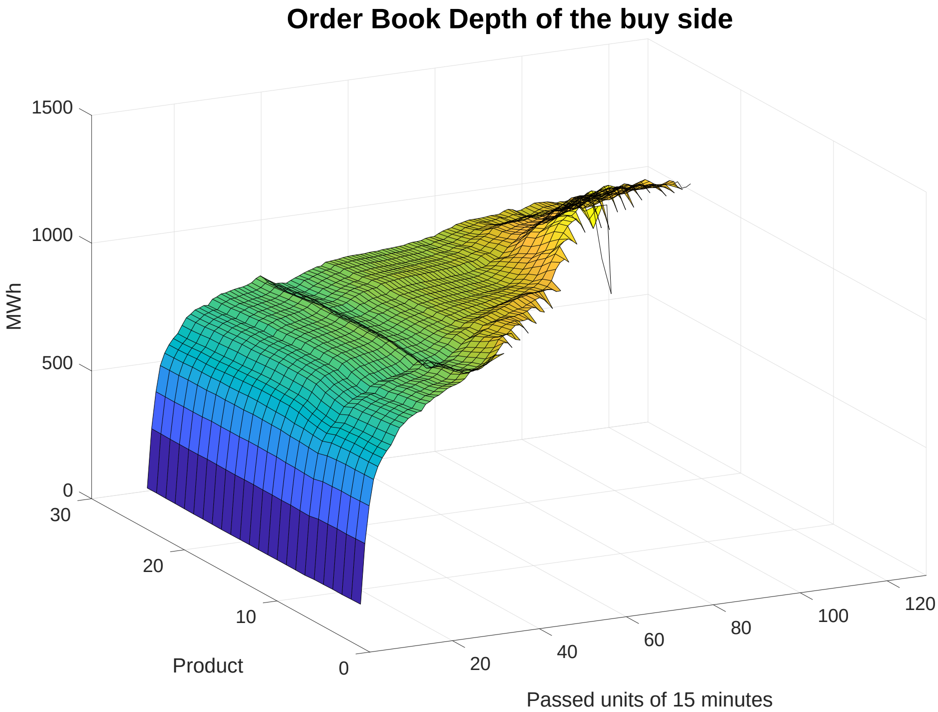

Figure 7: Depth.

- Tightness: trading costs for small quantities are small (bid-ask spreads).
- Depth: trading costs for large quantities should also be small.
- Resiliency: any deviation of the market price, which might come from market orders flow, should return quickly to the initial price-level.
- Immediacy: the ability to execute a transaction immediately at the prevailing price.

No single measure can be found in the literature to describe all dimensions of liquidity. Hence, a combination of many quantitative measures is considered to determine the liquidity of a market.

A basic analysis of liquidity can be done with summary statistics of the market. These statistics usually provide the overall traded volume on the buy- and sell-side, and statistics of prices, i.e., minimum, maximum, volume-weighted price, and last price. A market with large trading volumes and a small variance of prices can be classified as a liquid market. Hagemann and Weber [2015] analysed trading volumes of European intraday markets, which is still used in reports of EPEX to show growing liquidity of their intraday markets during the last years.

The transaction data consists of all cleared orders with the exact trading time, side, volume and price. This data-set can be used to produce more detailed analysis. Moreover, it allows to analyse prices and trading volumes along with the trading session and to analyse liquidity with a temporal behaviour. It is possible to determine the passive order and the corresponding price, which allows to reconstruct the price level of the two sides at observed clearings. Hagemann and Weber [2013] used the transaction data to analyse liquidity with more details. All clearing prices of the transaction data were classified as prices of the best-bid or best-ask. This method does not allow to estimate the bid-ask-spread for periods with no transactions, but during periods with many trades it is possible to obtain good estimates of the bid-ask spread.

The limit order book is needed to analyse the market's liquidity with all details. All submitted orders allow determining the depth of each side at each moment. Moreover, it allows to calculate the exact bid-ask spread at each time, but also orders beyond the bid-ask spread can be considered to evaluate hypothetical trading costs of larger volumes. Detailed data became available recently and Balardy [2022] was the first one to calculate the exact bid-ask spread.

1.8 Research Questions and Methodology

This dissertation analyses European intraday markets considering detailed order data from the Italian and German markets. The main sub-project of this dissertation was to implement the continuous clearing algorithm and the mechanism of the auction clearing. As input for the two algorithms, all submitted orders to the Italian intraday auction and the German continuous intraday market for the years 2015–2018 were available. The implemented environment allows to reproduce observed clearing prices but also to manipulate the market by adding additional orders or cancelling submitted orders. Most of the existing literature used transaction data to analyse intraday markets, but the implemented environment allows for a more detailed market analysis. The implemented algorithms allowed to answer the following three research questions:

> RQ1: How to compare liquidity of the German EPEX Spot Continuous intraday market and the Italian intraday auctions?

Liquidity is an important property for intraday power markets and market operators of the continuous intraday market praise the liquidity of their market. They use the trading volume as an indicator for liquidity, which strongly increased during the last years. However, liquidity is multidimensional, and more indicators should be analysed. This was not possible in the past due to missing data because the limit order book data became available just a few years ago. The implemented environment allows to analyse more dimensions of liquidity considering the buy- and sell-stack at each moment during trading time. Moreover, liquidity can be analysed along trading time. Another interesting topic is to compare liquidity of the two prevailing market designs for intraday markets considering all offers submitted to the Italian and German intraday market for hourly products.

> RQ2: Is it possible to develop a speculative trading strategy for the German EPEX Spot Continuous intraday market?

High-frequency trading is an interesting topic in financial markets. These strategies are purely speculative and trades can be triggered by signals from the limit order book or news from outside. These trading strategies take advantage of speed for being the first ones that use the newly available information and close the opened position later when the signal was absorbed by the market leading to the expected price change. The implemented environment made it possible to backtest these types of trading strategies with all its details considering liquidity costs, trading fees, and the impact of additional orders on the future development of the limit order book.

> RQ3: How could liquidity of the EPEX Spot Continuous intraday market be increased by introducing a new market design?

There are two prevailing market designs in Europe and each design has advantages and disadvantages. While the continuous intraday market suffers from low liquidity at the opening of the market, the auction market has a long lead-time making it difficult to integrate renewable energy sources. The detailed submitted orders of the German continuous intraday market were used to analyse the results of a newly proposed market design and the resulting clearings were compared with the current continuous trading outputs.

This dissertation is based on three main sections, where each section discusses one of the three introduced research questions. The last section summarizes the results and provides ideas for future research.

2 Liquidity Costs on Intraday Power Markets: Continuous Trading Versus Auctions

... written in cooperation with Prof. Dr. David Wozabal and published in Energy Policy.

We analyze liquidity costs on continuous and auction-based intraday power markets using a cost-of-round-trip measure that works for both market designs. We use data from the Italian auction-based intraday market and the German continuous market and present descriptive statistics as well as multivariate regression models to analyze determinants of liquidity costs in both markets. To test for differences in liquidity due to market design, we employ a double machine learning technique controlling for several confounding variables. We show that weekly patterns, yearly seasonality, electricity demand, as well as the influence of temperatures significantly affect liquidity costs. Comparing liquidity costs in both market, we find that, overall, liquidity costs are lower on the Italian market. However, Italian costs increase towards later auctions, while the costs on the German continuous intraday market decrease and reach their low close to physical delivery, where costs are lower than on the last Italian market trading the corresponding products.

2.1 Introduction

In the last two decades electricity markets world-wide have moved from being dominated by highly vertically integrated monopolies to competitive markets populated by many diverse players. To satisfy these companies' requirements, electricity trading takes place in multi-settlement markets that allow trading products with different temporal granularities and with different times to maturity. In particular, the growing share of variable renewable production led to the rising importance of spot markets, making it possible to adapt traded positions until close to delivery as new information arrives.

While in the US the day-ahead market is immediately followed by the real-time balancing market [Ela et al., 2014], European market designs feature a spot market that is split into a day-ahead market and an intraday market where power can be traded until shortly before physical delivery. Currently, there are two prevailing designs of intraday markets in Europe. While most European countries use continuous trading, Italy, Spain, and Portugal mainly use staggered intraday market auctions.

Clearly, the benefits of intraday trading are closely tied to the liquidity of the market, i.e., the ability of firms to trade while experiencing only minimal adverse price effects. Furthermore, liquid markets are less prone to market manipulation and gaming by pivotal players.

However, liquidity in most European intraday markets remains rather low. Weber [2010] finds that markets in Germany and several other European countries are not sufficiently liquid. Garnier and Madlener [2015] conclude that due to this illiquidity, current intraday markets are of limited use in balancing short-term forecast errors in demand and variable renewable production. It is therefore interesting to policy makers and industry professionals alike to identify factors that drive liquidity in the two market designs and understand how the designs themselves influence liquidity.

Consequently, the issue of liquidity in intraday markets has recently attracted some attention in the academic literature. Weber [2010] analyzes the integration of wind energy considering different European market designs and finds that the intraday auctions in Spain are the most attractive in terms of trading volume. Based on transaction data from the German intraday market, Hagemann and Weber [2013] investigate liquidity in intraday power markets using established measures from financial markets. Neuhoff et al. [2016] find that the additional auctions for 15 minutes contracts in the German intraday markets increased liquidity and market depth while reducing price volatility.

Balardy [2022] is one of the first, who uses the German limit order book (LOB) data to analyze liquidity in terms of bid-ask-spreads and market depths. The author finds a positive relation between bid-ask spreads and risk as well as a negative relation between bid-ask spread and adjustment needs, activity, and competition in the market. von Luckner et al. [2017] use the LOB to find an optimal market maker pricing and analyze the market order intensity and the bid-ask spread. Hagemann and Weber [2015] analyze intraday trading volumes on auction-based and continuous intraday markets, and observe higher volumes on the auction-based intraday markets. The authors conclude that this difference is not due to the difference in market design but rather due to idiosyncratic factors affecting the two markets.

The literature on electricity forecasting is in many ways related to our paper. Most models for price forecasts are time-series models using exogenous variables, some of which we also use in our models. For example, as in Narajewski and Ziel [2020b] and Uniejewski and Weron [2018], we use time dummies for Saturday, Sunday and Monday, and the day-ahead forecast for load, solar production and wind power as covariates in our regression models. Marcjasz et al. [2020] use dummies for each weekday, forecasts for load, solar production and wind production and its forecast errors, and balancing volumes. Janke and Steinke [2019] use the forecasts of demand and renewable production, and hourly dummies for each hour.

Despite the importance of the topic, the literature analyzing liquidity costs in intraday power markets remains scarce. To the best of our knowledge, this is the first paper to compare liquidity costs of the two markets in a statistically sound way using the complete order book data of the continuous intraday market and all submitted orders of the intraday auction.

In this paper, we contribute to the discussion by the first analysis of intraday electricity market liquidity that is based on a cost-of-round-trip (CRT) measure which captures all quantitative aspects of liquidity both in auction markets as well as for continuous trading. We provide a univariate analysis of the CRT which is complemented by regression models that explore possible drivers of liquidity costs on the German and Italian market. We find that, depending on the market, liquidity cost are driven by weekly patterns, yearly seasonalities, electricity demand, as well as temperatures.

To directly compare the cost of liquidity and thus measure the impact of market design, we use a state-of-the-art double machine learning method proposed in Chernozhukov et al. [2018] controlling for possible confounding factors identified in the analysis for the CRT for the two markets. Comparing the two markets, by and large the Italian auction-based market exhibits lower CRTs. We observe this result in a univariate analysis and confirm it in a multivariate analysis controlling for the confounding factors identified above. However, this effect gets progressively weaker for larger traded volumes and as trading time approaches physical delivery. In particular, it can be observed that the German continuous intraday market consistently exhibits lower costs for high volumes close to delivery.

Our findings suggest that a combination of several auction-based intraday markets with continuous trading might be able to leverage the benefits of both systems. In particular, auctions can be used to increase liquidity and therefore decrease trading costs by pooling orders for products which are far from delivery. These auctions could be complemented by continuous trading close to delivery, where market participants have the opportunity to trade the forecast errors for demand and variable renewable production at a point in time when accurate forecasts are available [see Ocker and Jaenisch, 2020, for a similar proposal]. In fact, Spain already implemented such a hybrid system when it joined the cross-border intraday market project XBID in June 2018. This proposal is close to the literature on optimal implementations of the European target model for a single coupled intraday market as laid out in the European Commission Regulation (EU) 2015/1222. Bellenbaum et al. [2014] discuss different intraday market designs meeting these requirements and come to the conclusion, that a hybrid between continuous trading and auctions potentially combines the advantages of both designs. Similarly, Ehrenmann et al. [2019] propose to add additional auction markets to the existing continuous market, as auction markets are more suitable for small market participants. The authors see a clear advantage of this setting, but the question remains at which time of the day to introduce auction markets and how many. A possible solution that leverages the advantages of both continuous trading and auctions is to have a large number of *frequent* auctions as proposed in Budish et al. [2015] for financial markets and in Deutsche Börse Group [2018] for the intraday power market. Such a design would alleviate some of the problems of continuous trading while still providing market participants with ample opportunities to trade.

The paper is organized as follows. In Section 2.2, we briefly describe the Italian and German intraday markets. Section 2.3 describes the market data and our set of explanatory variables. In Section 2.4, we introduce the cost-of-roundtrip measure and specify the econometric models used to determine the factors driving liquidity costs in both markets as well as the application of double machine learning, which we use to determine the effect of market design on liquidity costs. Section 2.5 discusses the empirical results. Finally, Section 2.6 concludes, discusses limitations and policy implications.

2.2 Background: Market Designs in Germany and Italy

In this short section, we discuss the relevant facts about the Italian auction-based intraday market and then proceed to discuss the German continuous intraday market. We collect key characteristics of the two markets for the year 2018 from ENTSO-E [2019], GME [2019], Burger [2019] in Table 3, and calculated the Italian weighted prices based on the national price. Note that the traded volumes of the day-ahead market and the intraday market of hourly products of the two markets are comparable. Consumption and production of renewables are higher in Germany, and Italy is a net importer of electricity while Germany generates high volumes for export, since it has significant overcapacities in cheap base-load production. As a result, average spot market prices in Germany are lower than in Italy.

2.2.1 The Italian IPEX

The Italian spot market offers a platform to trade electricity for delivery in hourly granularity. The day-ahead market in Italy closes at noon on the day before delivery and is followed by seven intraday auction markets, called MI (mercato infragiornaliero). Bid prices are constrained between €0 and €3000 while bid quantities are restricted to multiples of 1 kWh. For more details see GME [2016].

The Italian power grid consists of the six market zones NORD, CNORD, CSUD, SUD, SICI, and SARD. The MI markets are organized as *uniform price auctions* that aggregate the bids of all zones. The left plot in Figure 8 shows the cleared volume and the clearing price of an exemplary market session. If the resulting national market outcome is physically infeasible due to lack of transmission line

Quantity	Italy	Germany
Consumption (TWh)	322.2	538.1
PV infeed (TWh)	22.9	41.2
Wind infeed (TWh)	17.3	107.2
Imports (TWh)	47.1	31.5
Exports (TWh)	3.3	82.7
Day-ahead trading volume (TWh)	212.9	234.5
Intraday trading volume (TWh)	25.4	37.8
Volume weighted day-ahead price (€/MWh)	62.22	43.26
Volume weighted intraday price (€/MWh)	61.05	46.6

Table 3: Summary of annual key characteristics of the two markets for 2018. The German day-ahead volume includes Austria and Luxembourg and the trading volume for the German continuous market is restricted to hourly products

.

capacities between the zones, the result is made feasible by altering the market outcome resulting in different zonal prices for the different Italian market zones. For our analysis, we disregard this complication, by only considering the *national price*, which considers all submitted offers without taking into account the effects of transmission limits between zones.

Table 4 summarizes the characteristics of the Italian intraday market. The lead-time, defined as the time between the last possibility to trade the specific product and its physical delivery, range from 4.25 to 10.5 hours. Since wind power forecasts significantly improve approaching delivery [e.g., Hannele Holttinen, 2013], this relatively long lead-time make it hard to incorporate the last and therefore most precise production forecasts.

2.2.2 The German EPEX Spot Market

The German day-ahead market closes at noon of the previous day and is followed by an auction for quarter-hours of the next day at 3 p.m. and a continuous intraday market. For a detailed description we refer to the operational rules in EPEX [2019] and to Table 4 for a summary of trading times.

In contrast to the Italian MI markets, the German intraday market is based on continuous trading with a limit order book (LOB) much like in financial markets. Next to hourly products 1/2-hour and 1/4-hour products are traded. We do not include these products in our analysis, since shorter deliveries serve different purposes than hourly products. In particular, firms use sub-hourly products to model the ramps of their production or consumption, which is possible only to a

Market	Products	Opening	Closing	Results	Last Update	Lead-Time (h)
			Italian Markets			
MI1	$H1 - H24$	12:55 (d-1)	15:00 (d-1)	15:30 (d-1)	-	-
MI2	$H1 - H24$	12:55 (d-1)	16:30 (d-1)	17:00 (d-1)	$H1 - H4$	$7\frac{1}{2}$ up to $10\frac{1}{2}$
MI3	$H5 - H24$	17:30 (d-1)	23:45 (d-1)	00:15 (d)	$H5 - H8$	$4\frac{1}{4}$ up to $7\frac{1}{4}$
MI4	$H9 - H24$	17:30 (d-1)	3:45 (d)	4:15 (d)	$H9 - H12$	$4\frac{1}{4}$ up to $7\frac{1}{4}$
MI5	$H13 - H24$	17:30 (d-1)	7:45 (d)	8:15 (d)	$H13 - H16$	$4\frac{1}{4}$ up to $7\frac{1}{4}$
MI6	$H17 - H24$	17:30 (d-1)	11:15 (d)	11:45 (d)	$H17 - H20$	$4\frac{3}{4}$ up to $7\frac{3}{4}$
MI7	$H21 - H24$	17:30 (d-1)	15:45 (d)	16:15 (d)	$H21 - H24$	$4\frac{1}{4}$ up to $7\frac{1}{4}$
			German Markets			
Auction	$QH1 - QH96$	d-45	15:00 (d-1)	15:10 (d-1)	-	-
Cont. H	$H1 - H24$	15:00 (d-1)	D-5'	-	$H1 - H24$	$\frac{5}{60}$
Cont. QH	$QH1 - QH96$	16:00 (d-1)	D-5'	-	$QH1 - QH96$	$\frac{5}{60}$
Cont. HH	$HH1 - HH48$	15:30 (d-1)	D-5'	-	$HH1 - HH48$	$\frac{5}{60}$

Table 4: Operating times of the German and the Italian intraday markets. The table reports the traded products, the opening and closing times of the markets (d-1 indicating a time on the day before delivery), the time when the results are announced, the list of products that are traded the last time on the respective market, as well as the lead time for the products that are traded the last time. H indicates a hourly product, HH stands for half-hour and QH for a quarter-hourly product while D signifies the time of delivery.

small extent with hourly products. To be comparable to the Italian market, isolate the effect of market design on liquidity, and avoid diluting our analysis by mixing in different aspects, we therefore only consider hourly products in our analysis. The market for a specific product closes 30 minutes (or 5 minutes within the control area) before delivery, which facilitates trading forecast errors of fluctuating renewable energy sources.

Market participants can submit buy and sell offers for prices ranging between −9999.9€/MWh and 9999.9€/MWh, with a minimum bid size of 0.1MWh, and several specified order types [Martin and Otterson, 2018]. A submitted bid/offer is cleared immediately if the price is better than the best price of an offer/bid in the LOB. If there is no such matching order, the new order is stored in the LOB and matched with orders arriving at a later point in time. The right plot in Figure 8 shows the best available bid and ask price over time with each tick representing a match between a newly placed order and an order in the order book generating a trade.

2.3 Data

In Section 2.3.1, we discuss the market data which we use for the Italian and German intraday market. In Section 2.3.2, we introduce variables which we use

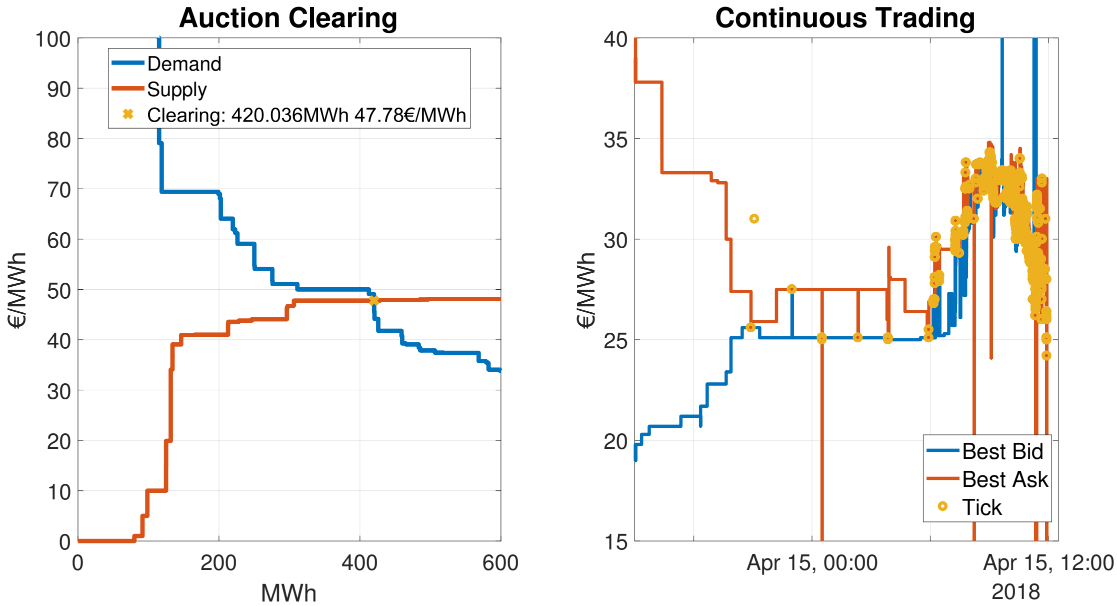

Figure 8: Clearing of the Italian MI3 intraday auction (left) and German continuous trading for the 13^{th} hour on the 15.04.2018. The yellow marker on the left signifies the uniform clearing price of the auction. The markers on the right represent price ticks, i.e., instances when orders were cleared in the German market.

in Section 2.4 and Section 2.5 as controls in our comparison of the two market designs.

2.3.1 Market Data

All offers submitted to the Italian intraday market are available on the website of the Italian Power Exchange (IPEX). The offers contain information about the side (sell or buy), product/hour, intraday market (MI1-MI7), zone, price and volume and can be used to calculate the *national price*.

The LOB of the German continuous intraday market can be purchased from EPEX SPOT SE. The data-set includes information about the side (sell or buy), product/hour, validity period, control area, as well as the price and volume of every submitted bid/offer. We note that the EPEX allows for the submission of so called iceberg orders, for which the bid quantity is only gradually revealed as parts of the order get executed. We only consider those parts of iceberg orders that were actually executed in our analysis. For more information about the LOB-data we refer to Martin and Otterson [2018].

The German intraday trading system was subject to frequent changes in the

	Variable	Frequency	Unit	Source
$R_t^{S,I}$	Italian solar production	hourly	MWh	https://transparency.entsoe.eu
$F_t^{S,I}$	Italian solar forecast	hourly	MWh	https://transparency.entsoe.eu
$R_t^{W,I}$	Italian wind production	hourly	MWh	https://transparency.entsoe.eu
$F_t^{W,I}$	Italian wind forecast	hourly	MWh	https://transparency.entsoe.eu
$R_t^{D,I}$	Italian demand	hourly	MWh	https://transparency.entsoe.eu
$F_t^{D,I}$	Italian demand forecast	hourly	MWh	https://transparency.entsoe.eu
$R_t^{S,G}$	German solar production	$\frac{1}{4}$-hourly	MWh	https://transparency.entsoe.eu
$F_t^{S,G}$	German solar forecast	$\frac{1}{4}$-hourly	MWh	https://transparency.entsoe.eu
$R_t^{W,G}$	German wind production	$\frac{1}{4}$-hourly	MWh	https://transparency.entsoe.eu
$F_t^{W,G}$	German wind forecast	$\frac{1}{4}$-hourly	MWh	https://transparency.entsoe.eu
$R_t^{D,G}$	German demand	$\frac{1}{4}$-hourly	MWh	https://transparency.entsoe.eu
$F_t^{D,G}$	German demand forecast	$\frac{1}{4}$-hourly	MWh	https://transparency.entsoe.eu
D_t	Daylight of Munich	daily	days	https://galupki.de
T_t^I	Temperature of Milan	hourly	°C	www.arpalombardia.it
T_t^G	Temperature of Berlin	hourly	°C	www.dwd.de
W_t	Weekends	daily	Boolean	-

Table 5: Overview of data used in the analysis.

recent years with effects on market liquidity, especially shortly before delivery. In order to have a dataset with consistent market rules, we restrict our analysis of both markets to the time from 20.11.2017, a few days after the trading system M7 (version 6.0) was launched to the 15.06.2018, when the XBID project was introduced.

2.3.2 Explanatory Variables

Table 5 provides an overview of the variables which potentially have an impact on the CRT and which we control in our comparison of the two market designs in Section 2.5.

Motivated by Goodarzi et al. [2019], Kulakov and Ziel [2020] who show that forecast errors in renewable production influence intraday prices and by Balardy [2022] who observes an impact of renewable energy sources on bid-ask spreads, we include data on forecasts and actual production of variable renewables in both countries. Since we exclusively analyse hourly products, we consider the average over the four quarter-hourly quantities to obtain hourly values. We use the day-ahead forecasts for renewable production as published by ENTSO-E. While the forecasts used by individual market participants for trading might be different, we think that the chosen forecast captures the overall sentiment of the market well.

Temperature influences power markets, because power is used for temperature regulation of buildings. Hence, we introduce a heating- and a cooling-function as described in Fan and Hyndman [2012] for Italy and Germany. The cooling function of Italy C_t^I and Germany C_t^G are defined as $\max(T_t - 19.5°\text{C}, 0)$, where T_t is the hourly temperature at time t in Italy (Milan) or Germany (Berlin). Analogously, we introduce the heating function for Italy H_t^I and Germany H_t^G as $\min(T_t - 17.5°\text{C}, 0)$. The choice of the two cities as temperature proxies is motivated by the fact that Milan is the leading industrial city in Italy and Berlin is the largest German city. A more detailed modeling of the influence of temperatures could be based on weighted temperatures from several areas in Germany and Italy as was for example done in Graf and Wozabal [2013], Kovacevic and Wozabal [2014], Pape et al. [2016]. However, for the purpose of this paper we stick to the abovementioned simple approach.

Prices on power markets follow a seasonal and weekly pattern. Hence, as in Kovacevic and Wozabal [2014] and Graf and Wozabal [2013], we use a variable containing the length of daylight D_t in units of days to capture annual seasonality of the observations. As these quantities are similar for both countries, we use the day-length of Munich located in the south of Germany for both markets. Moreover, as in Narajewski and Ziel [2020b] and Uniejewski and Weron [2018] we introduce dummy variables $W_t = (W_t^{Mon}, W_t^{Sat}, W_t^{Sun})$ for Monday, Saturday, and Sunday for weekends W_t to model weekly price patterns.

To capture the overall market size and therefore the scarcity of supply in a given period t, we use the forecast as well as the actual demand for Italy and Germany. An alternative way to capture the scarcity in an electricity system would be the so called *load-supply-ratio* (LSR) as defined by Pape et al. [2016]. The LSR takes into account detailed modeling of supply and demand and is a more accurate measure of scarcity than mere electricity demand. However, the demand is easier to include in our analysis, since it requires much less detailed data.

2.4 Methodology

In this section, we first detail how we measure liquidity costs in the two markets by a cost-of-round-trip measure in Section 2.4.1. In Section 2.4.2, we introduce a multivariate regression model to analyse the impact of possible confounding factors in the comparison of the two market designs. Finally, in Section 2.4.3, we discuss a double machine learning method in order to measure whether the continuous markets in Germany or the auction markets in Italy lead to higher CRTs.

2.4.1 Liquidity Measures

Market liquidity describes the possibility to quickly buy or sell an asset without affecting the market price. This rather vague definition of liquidity does not lend itself to a quantitative analysis of the phenomenon. In fact, there is no single established quantitative measure of liquidity in the literature that captures all aspects of market liquidity.

Hagemann and Weber [2013] introduced six dimensions of liquidity for continuous energy markets using established measures from the literature on financial markets. The first dimension is *tightness* and is measured using bid-ask spreads defined as the difference between the best bid and best ask price. The second dimension is *resiliency* describing the market's ability to bounce back to an equilibrium price after a temporary distortion. The third dimension is *price impact* or *market depth* and describes the impact of large orders which might require several offers beyond the best price to be cleared. The fourth dimension is known as *short-run price volatility*. The fifth dimension captures delay and search costs describing the propensity of traders to delay trades to obtain better prices. The sixth dimension describes trading activity in the form of traded volume, number of trades, and number of active traders.

Irvine et al. [2000] introduced a CRT-measure as the *per dollar cost of roundtrip trade of D dollars*. In particular, the number of shares that corresponds to the dollar amount D are calculated based on the best-bid and best-ask, and afterwards the LOB is used to calculate the resulting cost of buying and selling the determined number of shares. Since the interpretation in terms of quantities is more natural in power markets, we modify this definition by proposing a CRT measure which depends on volume V instead of the amount of money and captures all aforementioned cost related dimensions of liquidity. Moreover, we modify the measure to be applicable to both continuous trading as well as auction markets.

Conceptually, the CRT is the per unit cost incurred by buying a certain quantity V of power and then immediately selling it again. Note that in a liquid market CRT is close to zero. Choosing a small V yields measurements close to the bid-ask spread while larger volumes increasingly measure the depth of the order book plus all additional costs.

More formally, we define a volume oriented measure by sorting the buy- and sell side of the LOB at each point in time t by price to obtain $\cdots < P^t_{-2} < P^t_{-1} < P^t_{-0} < P^t_0 < P^t_1 < P^t_2 < \cdots$, where P^t_{-0} is the highest bid-price and P^t_0 is the lowest ask-price. We denote the corresponding bid quantities by Q^t_i. For a given quantity V in MWh, we define how much of an order i would be cleared when placing a market order of size V by

$$\bar{Q}^t_i(V) = \min\left(\max\left(V - \sum_{k=0}^{i-1} Q^t_k, 0\right), Q^t_i\right), \ \bar{Q}^t_{-i}(V) =$$
$$= \min\left(\max\left(V - \sum_{k=-i+1}^{-0} Q^t_k, 0\right), Q^t_{-i}\right).$$

We then define the cost-of-round-trip measure for a fixed value V as

$$CRT_t(V) = \underbrace{\frac{1}{V}\sum_k P_k^t \bar{Q}_k^t(V)}_{\text{average cost}} - \underbrace{\frac{1}{V}\sum_k P_{-k}^t \bar{Q}_{-k}^t(V)}_{\text{average revenue}}. \tag{1}$$

In a continuous market it is possible to execute the buy and sell decisions that are used to define the CRT, making equation (1) directly applicable. However, we note that, in principle, a trader in a continuous market has the option to *spread* her trades over a longer period of time, waiting for more orders on the other side of the market to arrive. In this way, some of the liquidity costs measured by the CRT can be avoided at the cost of the risk of adversely changing prices during the extended time of bidding. The CRT on the continuous intraday market can therefore be seen as an overestimation that accurately reflects liquidity costs only for an *impatient trader* placing market orders.

To use the CRT in an auction market, we add a market order for buying V units to the existing orders and record the marginal price instead of per unit cost when clearing the auction modified in this way. We then subtract the hypothetical sell price of V units which we calculate adding a market order of size V on the sell side instead and divide the result by V.

The resulting CRT-measure of the auction market consists of one value for each market and volume. In contrast, the CRT-measure of a certain product in a continuous market is a function of time and potentially changes with each modification of the LOB. As is illustrated in the right panel of Figure 8 large market orders might lead to temporary extreme values of the CRT-measure distorting our measurement. We therefore use the mean over 15 minutes instead of $CRT_t(V)$ at any fixed time t. To this end, we consider a discrete form of the continuous time varying CRT-measure by considering averages over 15 minute intervals before time τ

$$CRT_\tau(V) = \frac{1}{15}\int_{\tau-15}^{\tau} CRT_t(V)\, dt = \frac{1}{15}\sum_{k=2}^{N} \frac{CRT_{t_k}(V) + CRT_{t_{k-1}}(V)}{2}(t_k - t_{k-1}),$$

where $t_1, \ldots, t_N$ are the N points in time where the LOB changes in the 15-minute time interval $[\tau - 15, \tau]$. In the following, we use the index τ in CRT_τ to refer to a 15-minute average and CRT_t to refer to an instantaneous CRT at time t. The computed average thus reflects the *expected CRT* a trader would have to pay, if she picks a random trading time in the given time interval.

The Italian intraday market has seven fixed times when the market is cleared. We use clearing times of MI2 to MI7 to analyze the two markets, i.e., measure the CRT for the German markets at the times when the Italian markets are cleared.

The reason for the exclusion of MI1 is that the German intraday auction closes nearly at the same time as MI1, which results in less liquidity on the continuous market at this point in time and thus a distortion. We compare the CRT of the remaining Italian intraday auctions with the mean German CRT over the 15 minutes before the closing of the Italian intraday auction. To this end, we define $\mathcal{D}_h^I$ as the closing times of the Italian intraday markets, where the hourly product h is traded. For example, $\mathcal{D}_1^I = \{16:30\}$ while $\mathcal{D}_{24}^I = \{16:30, 23:45, 3:45, 7:45, 11:15, 15:45\}$, where the first two time stamps are from the day before delivery.

The German continuous intraday market allows participants to trade until 30 minutes before physical delivery on a national market. Hence, we will also compare the first two 15-minute CRT-means within the last hour of the German continuous intraday market with the CRT-measure of the last available market of the Italian intraday auction for the corresponding product. Correspondingly, the points in time which we consider for the German market are $\mathcal{D}_h^G = \mathcal{D}_h^I \cup \{h-60, h-45\}$.

We generate observations corresponding to $V = 0.1$MWh, which is the smallest value that can be traded on the German intraday markets as well as for $V =$ 5MWh, 10MWh, 15MWh and 50MWh. On some days the order book does not contain orders of combined size V on either the bid or the ask side at a time $t_i \in [\tau - 15, \tau]$. For our analysis, we calculate over 313 million clearings for the German market. In 0.0466% of these cases at least one side of the limit order book is empty and we exclude these timestamps in our calculation of the 15-minute intervals. In further 0.0804% of the cases not the whole quantity V is available on at least one side of the market. To define CRT for these cases, we use the last available price to clear the remaining quantity in order to calculate a CRT.

2.4.2 Analysis of the CRT

In this section, we analyze the impact of the variables described in Section 2.3.2 on the CRTs of the two markets. To this end, we define an index $\mathcal{J} = (V, h, \tau)$ for every volume V, product $h = 1, \dots, 24$, and time to delivery $\tau \in \mathcal{D}_h^G$ or $\tau \in \mathcal{D}_h^I$ and construct the following linear regression models for Italy and Germany

$$CRT_{\mathcal{J}}^G = X_{\mathcal{J}}^G \beta_{\mathcal{J}}^G + \epsilon_{\mathcal{J}}^G \quad \text{and} \quad CRT_{\mathcal{J}}^I = X_{\mathcal{J}}^I \beta_{\mathcal{J}}^I + \epsilon_{\mathcal{J}}^I, \tag{2}$$

where

$$X_{\mathcal{J}}^G = (X_{\mathcal{J}}, C_{\mathcal{J}}^G, H_{\mathcal{J}}^G, R_{\mathcal{J}}^{W,G}, F_{\mathcal{J}}^{W,G}, R_{\mathcal{J}}^{S,G}, F_{\mathcal{J}}^{S,G}, R_{\mathcal{J}}^{D,G}, F_{\mathcal{J}}^{D,G})$$
$$X_{\mathcal{J}}^I = (X_{\mathcal{J}}, C_{\mathcal{J}}^I, H_{\mathcal{J}}^I, R_{\mathcal{J}}^{W,I}, F_{\mathcal{J}}^{W,I}, R_{\mathcal{J}}^{S,I}, F_{\mathcal{J}}^{S,I}, R_{\mathcal{J}}^{D,I}, F_{\mathcal{J}}^{D,I}),$$

$X_{\mathcal{J}} = (1, W_{\mathcal{J}}, D_{\mathcal{J}})$ are the regressors that are market independent, and $CRT_{\mathcal{J}}^G$ and $CRT_{\mathcal{J}}^I$ are the CRTs of the German and Italian market, respectively. All regressors are standardized by subtracting the mean and dividing by the standard deviation. The standardization helps to simplify the interpretation of the effects of covariates with different scales.

We estimate the models in equation (2) separately, for every index $\mathcal{J}$. This yields 420 models for the Italian intraday auction market, and 660 models for the German continuous intraday market, because we additionally analyze the two 15-minutes intervals shortly before physical delivery for the German market. For example, for $h = 1$, we compare the liquidity cost on the two 15-minute intervals that start 60 minutes and 45 minutes before physical delivery on the German market with the latest available intraday market in $\mathcal{D}_1^I$, i.e., MI2.

2.4.3 Double/Debiased Machine Learning

In this section, we describe how we compare the impact of the two market designs on the CRT while controlling for the impact of confounding variables. In particular, we directly compare the CRT in the two markets while controlling for linear and non-linear effects of the regressors introduced in Section 2.3.2. For this purpose, for every volume V, product h, and every trading time $\tau \in \mathcal{D}_h^G$, we combine the data on $CRT_{\mathcal{J}}^G$ and $CRT_{\mathcal{J}}^I$ into a combined $CRT_{\mathcal{J}}^C$ by stacking the two vectors on top of each other. For $\tau \in \mathcal{D}_h^G \setminus \mathcal{D}_h^I$, we use the CRTs of the corresponding last market where the hour was traded on an Italian intraday market.

We then define a sparse matrix

$$X_{\mathcal{J}}^C = \begin{pmatrix} X_{\mathcal{J}} & X_{\mathcal{J}}^G & 0 \\ X_{\mathcal{J}} & 0 & X_{\mathcal{J}}^I \end{pmatrix}$$

by padding market specific observations with zeros. We compute all quadratic interactions to capture non-linear effects obtaining

$$\begin{pmatrix} Y_{\mathcal{J}}^F & \begin{matrix} Y_{\mathcal{J}}^G & 0 \\ 0 & Y_{\mathcal{J}}^I \end{matrix} \end{pmatrix},$$

where $Y_{\mathcal{J}}^G$ and $Y_{\mathcal{J}}^I$ consist of interactions that contain a market specific variable for Germany and Italy, respectively, while $Y_{\mathcal{J}}^F$ contains interactions of variables in $X_{\mathcal{J}}$. Next, we delete all columns with fewer than 10 observations different from zero.

We then replace the zeros of the sparse submatrices with the corresponding mean to obtain

$$\begin{pmatrix} Y_{\mathcal{J}}^F & \begin{matrix} Y_{\mathcal{J}}^G & \bar{Y}_{\mathcal{J}}^I \\ \bar{Y}_{\mathcal{J}}^G & Y_{\mathcal{J}}^I \end{matrix} \end{pmatrix}. \tag{3}$$

We standardize (3) by subtracting the mean and dividing by the standard deviation and denote the resulting matrix by $Y_{\mathcal{J}}^C$.

Note that replacing the zeros by the respective means in (3) ensures that there is no variable in $Y^C_{\mathcal{J}}$, which has a different mean for the subset for Italian and German observations. We introduce a dummy variable G that takes the value 1 for CRT values from the German market and 0 for data from the Italian market. Using these regressors, we specify a combined linear model

$$CRT^C_{\mathcal{J}} = \alpha_{\mathcal{J}} G_{\mathcal{J}} + Y^C_{\mathcal{J}} \beta^C_{\mathcal{J}} + \epsilon_{\mathcal{J}}, \tag{4}$$

which is able to control for interactions between the variables and non-linear effects. Moreover, all regressors have mean zero and the introduced dummy variable $G_{\mathcal{J}}$ is the only available variable to describe the systematic differences in CRTs between the two countries.

Our aim is to obtain consistent estimates of the effect of the market design $\alpha_{\mathcal{J}}$ as well as confidence intervals. Equation (4) has many regressors and we are no longer able to apply OLS due to overfitting. Hence, we would have to select a subset of regressors using a model selection mechanism and then estimate the coefficient α from the reduced model. However, as pointed out by Leeb and Pötscher [2005], model selection distorts inference and especially small parameters cannot be estimated consistently. Additionally, the same data set would be used twice: the first time for model selection and the second time to estimate $\alpha_{\mathcal{J}}$ and its p-value in the resulting regression. Another naive method would be to estimate the model (4) using a LASSO regression and directly analyze $\alpha_{\mathcal{J}}$. However, the resulting estimates are biased due to the L1-regularization term introduced in LASSO.

In order to avoid biased estimates for $\alpha_{\mathcal{J}}$, we use a double machine learning procedure by Chernozhukov et al. [2018] as implemented in STATA. The method uses Neyman-orthogonal moments/scores to eliminate the regularization bias and cross-fitting to eliminate the bias resulting from over-fitting of nuisance functions. In particular, we use LASSO regression for model selection in (4) where the penalty parameter is chosen using 10-fold cross validation. We resample 10 times for the calculation of an unbiased estimate $\tilde{\alpha}_{\mathcal{J}}$ for the parameter $\alpha_{\mathcal{J}}$ in the selected models. We refer to StataCorpLLC [2019] for a detailed exposition of the method.

2.5 Results and Discussion

In this section, we first consider a descriptive analysis of CRT in Section 2.5.1. In Section 2.5.2, we construct two linear regression models to analyze the impact of confounding variables on the liquidity costs of the two markets. Finally, we analyze the difference of the two market designs using double-machine learning in Section 2.5.3.

Subset	N	mean	std	min	25%	50%	75%	max
Average CRTs at $\mathcal{D}_h^I$ and $\mathcal{D}_h^G$								
GER, 0.1MWh	27453	5.85	6.27	0.10	2.40	4.50	7.50	137.70
GER, 5MWh	27453	6.25	6.67	0.10	2.81	4.90	7.95	147.50
GER, 10MWh	27453	6.69	7.00	0.10	3.00	5.02	8.33	162.75
GER, 15MWh	27453	7.18	7.34	0.10	3.40	5.62	8.93	168.50
GER, 50MWh	27453	10.66	10.24	0.10	5.65	8.51	12.85	198.55
ITA, 0.1MWh	17471	1.26	1.72	0.01	0.25	0.72	1.59	27.63
ITA, 5MWh	17471	2.13	2.68	0.01	0.50	1.22	2.82	38.08
ITA, 10MWh	17471	2.74	3.29	0.01	0.67	1.72	3.60	46.79
ITA, 15MWh	17471	3.33	3.88	0.01	0.92	2.11	4.34	50.45
ITA, 50MWh	17471	6.45	6.65	0.01	2.18	4.50	8.26	63.47
Trading Volume Weighted CRTs								
GER, 0.1MWh	4991	1.99	1.23	0.55	1.24	1.64	2.34	14.32
GER, 5MWh	4991	2.20	1.33	0.62	1.40	1.83	2.60	24.31
GER, 10MWh	4991	2.38	1.42	0.72	1.54	1.99	2.79	35.52
GER, 15MWh	4991	2.56	1.49	0.81	1.69	2.17	3.00	39.55
GER, 50MWh	4991	3.90	2.09	1.32	2.67	3.37	4.52	58.61
ITA, 0.1MWh	4991	0.84	0.62	0.01	0.43	0.70	1.07	5.98
ITA, 5MWh	4991	1.36	0.93	0.01	0.73	1.15	1.75	8.29
ITA, 10MWh	4991	1.71	1.13	0.01	0.95	1.46	2.18	9.49
ITA, 15MWh	4991	2.05	1.32	0.01	1.15	1.73	2.63	12.10
ITA, 50MWh	4991	3.84	2.35	0.01	2.21	3.32	4.84	21.25

Table 6: Descriptive statistics of CRT-measures and traded-volumes CRT-measures from 17.11.2017 to 15.06.2018.

2.5.1 Univariate and Bivariate Analysis of CRT

The descriptive statistics of the CRT-measures are summarized in Table 6. The first panel reports the average CRTs as measured at the points in time $\mathcal{D}_h^I$ and $\mathcal{D}_h^G$ which we use in our comparisons between the markets. However, since trading in the German continuous intraday market occurs mostly within the last three hours before delivery, we also define a *trading volume weighted CRT*, which allows us to compare CRTs of a specific product over longer periods of time as

$$CRT_{V,h} = \sum_{\tau} \frac{CRT_{V,h,\tau} Q_{h,\tau}}{\sum_{\tau} Q_{h,\tau}},$$

where $Q_{h,\tau}$ is the traded volume for product h and time to delivery τ. The above sum is over all quarter hours τ where a specific product h is traded. Similarly,

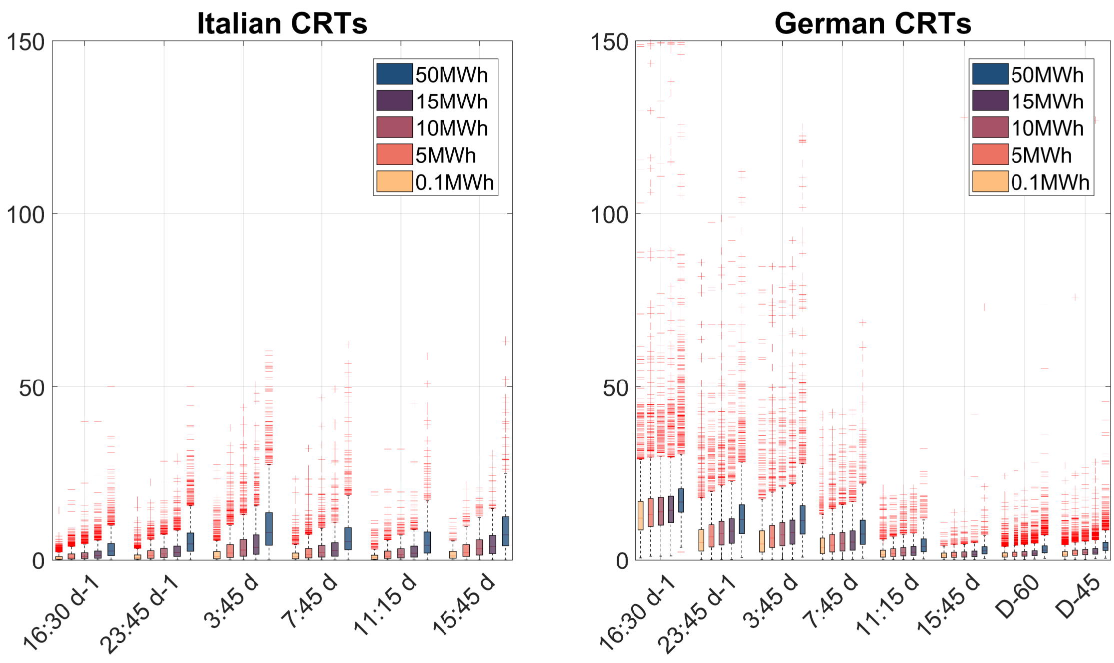

Figure 9: Boxplots of CRTs grouped by trading time and volume.

when computing $CRT_{V,h}$ for the Italian markets, the cleared volumes for each auction market and the corresponding calculated CRTs are used. The results of these computations are reported in the lower panel of Table 6.

The analysis reveals that the CRT for all volumes is higher for the German market on average for both ways of measurement. Comparing the maxima of the distributions, we observe that the corresponding CRTs for the German market far exceed the maximal CRTs observed in the Italian markets. However, the results for the averages are not entirely driven by the right tail of the distribution as the analysis of the other quantiles reveals. Another interesting observation is that while the CRT for the Italian market increases sharply with V, this effect is much less pronounced on the German market, where costs are high even for small volumes due to the bid-ask spread on the German market.

The univariate analysis along the dimension volume does not capture changes with the time to delivery. Hence, we show the dependence of the results on the time to delivery in the boxplots in Figure 9 for the CRTs calculated at $\mathcal{D}_h^I$ and $\mathcal{D}_h^G$. We note that liquidity costs on the Italian market are low during the first two auction markets, and are relatively high for the MI4 and MI7. The German CRTs decrease towards one hour before physical delivery and increase afterwards – this *L-shape* was also observed in Balardy [2022].

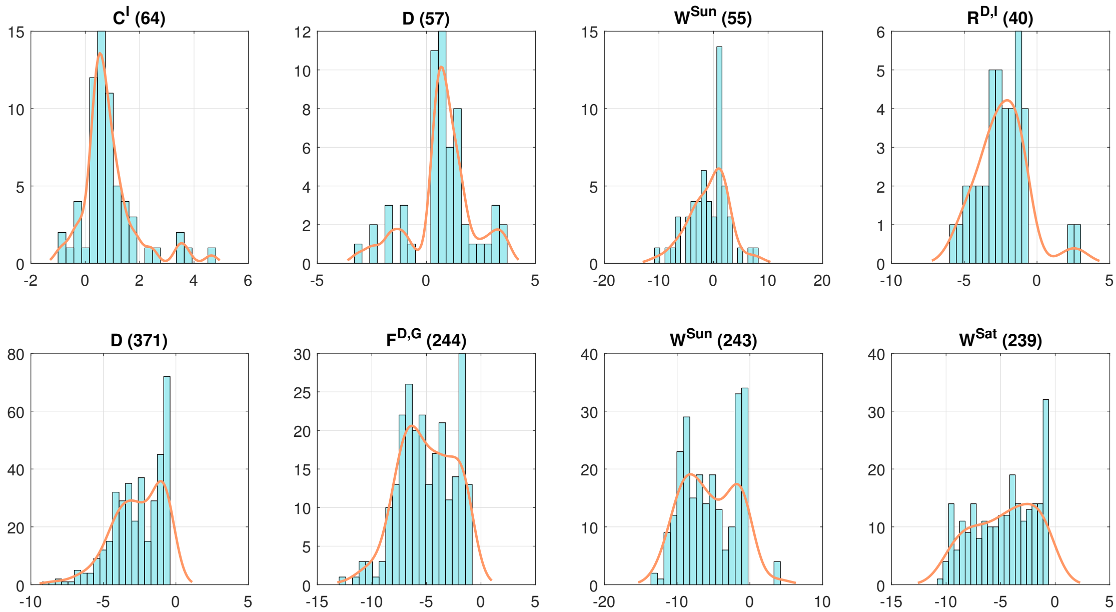

Figure 10: Distribution of the top 4 significant estimates of the selected controls of Italy (above) and Germany (below). The x-axis of the plots represents the values of the estimated coefficients.

2.5.2 Effects in the Individual Markets

In this section, we analyze the effect of the explanatory variables X^I and X^G as introduced in Section 2.4.2 on the CRT in the respective markets. In order to do so, we fit the linear regression models (2) using the *fitlm* function as implemented in *MATLAB R2017a*. We consider the same data-set as used in the previous section grouped by volume, product, and time to delivery.

We consider a regressor to be significant in a regression, if its p-value is smaller than 0.05 and order the regressors according to the number of models that they are significant in. The upper row of plots in Figure 10 shows the distribution of coefficients of the four regressors which are most often significant in the estimated models for the Italian market. The lower four plots repeat this analysis for the German market.

For Germany, the most important regressor is the seasonality D_t modeled as the length of daylight, which has a significant positive impact in 371 out of 660 models. As the estimated coefficients are unambiguously negative, this implies lower liquidity costs in summers.

The next most significant regressor is the forecast demand $F^{D,G}$, which is significant in 244 models and has also a clearly negative coefficient implying that higher (forecast) demands lead to more trading, which in turn decreases liquidity costs.

The last two depicted regressors are the dummies for Sundays and Saturdays which are significant in 243 and 239 models, respectively. On a first glance, the negative signs of the estimated regressors might seem surprising, since there is less trading on the weekend lowering liquidity costs. However, this effect is already captured by the regressor $F^{D,G}$ so that the weekend dummies only measure the weekly patterns which do not directly depend on demand. The dummies for Saturday and Sunday, thus allow for a more moderate increase in liquidity costs on these days as would be modeled by the effect of lower demand alone.

By and large the German market shows clear effects and the corresponding regressors are significant in many of the considered models, which underlines the importance of considering these variables as controls when we measure the effect of market design on liquidity costs in Section 2.4.3.

The situation for the Italian models is not nearly as clear cut. Generally speaking, the proposed regressors are significant much less often and the signs are more ambiguous making easy explanations of the results harder. This is in line with Hagemann and Weber [2015], who find that the trading volume on the Italian auction market cannot be explained very well by fundamental variables.

The most important regressor for the Italian market is cooling C^I which significantly affects liquidity in 64 out of 420 models for the Italian market with a mostly positive sign implying that the increased demand by air-conditioning, which is widely used in Italy, leads to a positive impact on liquidity cost on hot days.

The length of daylight D^I, which is significant in 57 models is the second most important regressor in Italy. As the figure shows, the estimated coefficients are mostly positive indicating a positive impact of the length of daylight on the CRT. This implies a seasonal effect with higher liquidity cost in summers. This is in contrast to the German situation, where the effect on the seasonal variable is reversed.

The Sunday dummy is significant in 55 models. The sign of the regressor is rather ambiguous and hard to interpret, since, similar to the German market, there is an interaction with the realized demand, which is also contained among the top 4 regressors.

Lastly, the realized demand $R_t^{D,I}$ affects the CRT on the Italian market significantly in 40 models, where it mostly has a negative effect on the CRT.

2.5.3 Comparison of Market Designs

Our aim in this section is to analyze the difference of the CRTs of the two markets controlling the effect of confounding variables. For that purpose, we use the function *xporegress* of STATA StataCorpLLC [2019] to estimate the models presented in Section 2.4.3.

The output of our analysis is an estimate, a valid confidence interval, and the corresponding p-value for the parameter α in model (4). Table 7 summarizes the results in form of a heatmap showing estimates and p-values. The columns indicate different hourly products, while the rows indicate time to delivery. To distinguish between the different quantities V, we divide the table into five panels.

We compare CRTs for products with the same time to delivery, and the CRTs for the two 15 minutes intervals on the German continuous intraday market before delivery of a specific product with the last auction market in Italy where the corresponding product is traded. Cells marked grey indicate products that can not be compared, since they are no longer traded on the Italian market. A cell is colored red if the estimate for $\alpha_{\mathcal{J}}$ in the corresponding model is positive, i.e., the CRT in the Italian market is lower than in the German market. Analogously, cells are colored blue if $\alpha_{\mathcal{J}}$ is negative. The intensity of the color reflects the magnitude of the p-value with more intense coloring for lower p-values, i.e., more significant results as indicated in the color map in the last row of the table.

As expected from the univariate results in Section 2.5.1, the majority of cells are red indicating higher cost of liquidity on the German market. Comparing the overall results of the five different panels, this effect weakens for higher volumes V, indicating that the German market is relatively less affected by large volume bids as can also be seen in Figure 9.

Observing the first rows of the five panels, it becomes clear that there is a strong influence of the time to delivery on the estimated parameter $\alpha_{\mathcal{J}}$. In particular, the Italian market has clearly lower liquidity cost at the time of clearing of the first two Italian intraday markets for all volumes V. However, looking at single columns corresponding to products $h = 1, \ldots, 24$, this effect weakens as trading times move closer to delivery. These results are consistent with the analysis in Figure 9 and the fact that traded volumes tend to decrease for later Italian auction markets, while the German market is most active close to delivery.

The last two rows of every panel compare the first two 15-minutes intervals in the last hours before delivery in the German market with the last Italian auction market where the respective hour can be traded. In these 15-minute intervals the German market reaches its highest liquidity and exhibits significantly lower liquidity cost as the Italian markets, except for small volumes.

In summary, the German market gets relatively more liquid towards physical delivery, with higher liquidity in the German market close to delivery and for larger volumes V. This is also supported by looking at single rows where we mostly observe increasing estimates for $\alpha_{\mathcal{J}}$ with increasing products $h = 1, \ldots, 24$.

Looking at the first non-gray blocks in every row corresponding to MI3-MI7,

V=0.1MWh

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
16:30 d-1	10.49	10.85	10.83	11.01	11.76	11.85	12.07	11.96	12.13	13.32	13.39	13.38	13.85	14.09	13.94	14.37	14.28	14.44	14.35	14.49	14.59	14.65	14.48	14.56
23:45 d-1					0.47	1.27	2.23	3.25	4.09	4.72	5.62	6.31	6.88	7.50	6.89	7.40	7.31	7.04	6.63	7.39	7.40	7.63	7.05	7.69
3:45 d									-0.90	0.22	1.69	2.37	5.27	5.77	4.94	5.59	5.44	6.02	5.28	6.45	6.40	6.26	6.04	6.75
7:45 d													0.17	0.46	1.17	1.78	3.51	4.53	3.50	4.02	4.94	5.77	5.23	5.54
11:15 d																	-0.01	0.42	0.46	0.86	1.38	1.91	2.23	2.43
15:45 d																					-0.86	-0.31	-0.01	0.44
D-60	0.69	0.80	0.49	0.78	0.56	0.74	0.61	0.89	-1.46	-1.80	-1.60	-1.52	-0.41	-0.50	-0.57	-0.25	0.18	0.42	0.38	0.39	-0.37	0.07	0.17	0.67
D-45	0.91	1.00	0.85	0.97	0.90	0.91	1.00	1.04	-1.36	-1.57	-1.28	-1.22	-0.09	-0.24	-0.18	0.02	0.54	0.71	0.85	0.77	0.07	0.46	0.80	0.99

V=5MWh

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
16:30 d-1	10.58	11.31	11.24	11.54	12.40	12.60	12.68	12.35	12.62	13.92	13.84	14.16	14.78	14.89	14.80	14.98	14.76	14.91	15.00	14.97	14.94	15.43	15.23	15.57
23:45 d-1					-0.58	0.28	1.63	2.39	3.72	4.69	6.29	7.20	7.83	8.39	7.84	8.15	7.97	8.10	7.94	8.15	8.44	8.45	8.33	9.01
3:45 d									-3.70	-2.09	-0.39	1.14	5.41	6.26	5.64	6.07	6.02	6.42	5.88	6.67	7.30	6.76	7.09	7.83
7:45 d													-0.94	-0.52	0.20	0.92	3.32	4.62	3.58	3.73	4.94	5.86	5.82	6.17
11:15 d																	-0.78	-0.23	-0.25	0.27	0.92	1.85	2.00	2.41
15:45 d																					-2.82	-1.43	-1.01	-0.62
D-60	0.16	0.46	0.09	0.35	-0.42	-0.19	-0.11	-0.22	-4.21	-4.23	-4.10	-3.48	-1.49	-1.51	-1.63	-1.26	-0.59	-0.19	-0.31	-0.29	-2.22	-0.98	-0.71	-0.33
D-45	0.76	0.81	0.59	0.71	0.08	0.16	0.53	0.03	-4.01	-3.88	-3.65	-3.08	-1.05	-1.16	-1.14	-0.89	-0.09	0.24	0.25	0.21	-1.62	-0.41	0.14	0.21

V=10MWh

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
16:30 d-1	10.71	11.41	11.35	11.52	12.68	12.88	12.94	12.65	12.89	14.09	14.09	14.66	15.24	15.29	15.26	15.35	15.05	15.27	15.40	15.08	15.05	15.81	15.54	16.00
23:45 d-1					-1.17	-0.13	1.17	2.02	3.42	4.79	6.36	7.46	8.20	8.94	8.23	8.66	8.49	8.81	8.53	8.76	8.89	9.04	9.10	9.64
3:45 d									-5.39	-3.25	-1.38	0.62	5.26	6.43	5.70	6.17	6.25	6.39	5.91	6.90	7.04	7.04	7.46	8.27
7:45 d													-1.46	-1.15	-0.28	0.33	3.06	4.69	3.59	3.33	4.85	5.81	6.04	6.53
11:15 d																	-1.12	-0.75	-0.77	-0.15	0.68	1.55	1.99	2.45
15:45 d																					-4.23	-2.25	-1.64	-0.75
D-60	-0.07	0.21	-0.18	-0.01	-1.04	-0.65	-0.65	-0.73	-5.91	-5.52	-5.41	-4.43	-2.04	-2.17	-2.20	-2.04	-0.94	-0.71	-0.88	-0.82	-3.60	-1.77	-1.43	-1.02
D-45	0.75	0.58	0.36	0.36	-0.50	-0.27	0.05	-0.48	-5.69	-5.13	-4.90	-3.98	-1.55	-1.80	-1.67	-1.62	-0.41	-0.24	-0.28	-0.29	-2.96	-1.03	-0.51	-0.40

V=15MWh

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
16:30 d-1	10.80	11.68	11.56	11.68	12.90	12.98	13.10	12.82	13.12	14.37	14.47	14.99	15.69	15.72	15.60	15.68	15.41	15.52	15.66	15.04	15.30	16.18	15.98	16.41
23:45 d-1					-1.80	-0.52	0.80	1.49	2.89	4.61	6.33	7.65	8.49	9.38	8.55	9.20	8.77	9.04	8.81	8.97	9.16	9.44	9.58	10.10
3:45 d									-7.12	-4.95	-2.36	-0.23	5.03	6.39	5.43	5.88	6.39	6.60	5.98	6.79	6.61	6.68	7.58	8.44
7:45 d													-1.80	-1.46	-0.73	-0.21	2.82	4.51	3.54	3.03	4.68	5.53	6.16	6.85
11:15 d																	-1.43	-1.20	-1.27	-0.58	0.26	1.37	2.01	2.57
15:45 d																					-5.66	-2.93	-2.17	-0.72
D-60	-0.29	0.11	-0.29	-0.29	-1.67	-1.07	-1.10	-1.32	-7.67	-7.36	-6.60	-5.61	-2.39	-2.52	-2.72	-2.74	-1.31	-1.17	-1.42	-1.28	-5.09	-2.38	-1.97	-1.43
D-45	0.62	0.50	0.29	0.08	-1.10	-0.68	-0.38	-1.07	-7.45	-6.92	-6.05	-5.11	-1.87	-2.12	-2.15	-2.30	-0.74	-0.67	-0.76	-0.70	-4.39	-1.71	-1.03	-0.78

V=50MWh

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
16:30 d-1	12.62	13.54	13.64	13.09	15.04	14.95	15.21	14.73	14.20	15.86	16.76	17.47	18.02	18.19	17.85	17.71	17.54	17.53	17.26	16.37	16.88	18.29	18.36	19.06
23:45 d-1					-4.14	-2.48	-0.35	-0.72	1.67	3.88	6.90	8.68	10.60	11.71	10.58	11.17	11.08	10.58	10.20	9.86	9.72	11.21	12.09	12.89
3:45 d									-13.59	-10.58	-6.47	-2.86	4.52	6.98	5.65	6.24	6.88	6.62	5.72	5.43	4.97	6.92	9.21	10.39
7:45 d													-3.48	-3.15	-2.95	-2.06	1.98	3.64	2.65	1.71	2.60	5.38	7.29	8.70
11:15 d																	-3.54	-3.01	-3.29	-3.39	-2.01	0.69	2.37	3.52
15:45 d																					-10.28	-6.90	-4.38	-0.78
D-60	-0.49	-0.09	-0.72	-0.84	-3.87	-3.08	-2.46	-4.01	-14.33	-14.14	-12.49	-10.63	-4.22	-4.62	-5.35	-5.37	-3.50	-3.11	-3.69	-4.24	-10.53	-6.12	-4.10	-2.23
D-45	0.91	0.51	0.22	-0.46	-3.08	-2.43	-1.47	-3.72	-14.27	-13.38	-11.64	-9.81	-3.49	-3.98	-4.51	-4.77	-2.67	-2.36	-2.68	-3.28	-9.48	-4.97	-2.74	-0.93

Legend

NaN	5e-07	1e-06	5e-06	1e-05	5e-05	1e-04	5e-04	1e-03	5e-03	1e-02	5e-02	1e-01	5e-02	1e-02	5e-03	1e-03	5e-04	1e-04	5e-05	1e-05	5e-06	1e-06	5e-07

Table 7: Results in €/MWh using all possible combinations until quadratic terms with standardized regressors and re-sampling (10).

i.e., the hours that can be traded the last time on an Italian market, we observe that liquidity cost on the Italian market is higher than on the German market. These markets are the last possibility to trade forecast updates of renewable energy sources in the corresponding hours and trading volumes are correspondingly relatively high. Apparently, the high CRTs are thus a consequence of tight market situations caused by either large demands or large free production capacities flooding the market for low prices.

2.6 Conclusions and Policy Implications

This article explores liquidity costs of the German continuous intraday market and the Italian auction-based intraday market. For that purpose, we introduce a cost-of-round-trip measure to analyze liquidity costs. Grouping the data of each market by volume and trading time, we compare cost of liquidity in the two markets using descriptive statistics. Secondly, we analyze the impact of several explanatory variables on the two markets separately. Thirdly, we compare the two market designs by controlling the impact of the confounding variables.

We find that liquidity costs are generally lower in the Italian auction market, whereby the difference tends to decrease with the traded quantity of power and as trading gets closer to physical delivery. The latter finding is consistent with the *L-shape* of the German bid-ask spread observed by Balardy [2022].

Our results show that the cost of liquidity in both countries is influenced by weekly and yearly seasonalities, temperatures via cooling demand, and the overall demand for electricity.

Our study has some limitations. Firstly, the German market provides the possibility to place iceberg orders, i.e., orders where the full volume is not visible but gets revealed gradually as parts of the order are cleared. The existence of a significant amount of these *invisible* orders might lead us to underestimate the liquidity and correspondingly overestimate the CRT on the German market. Secondly, the CRT on the Italian intraday auction markets might be higher due to zonal prices in Italy in auctions where there is congestion of transmission lines between market zones.

Our analysis suggests that a hybrid system might leverage the advantages of both market designs and decrease liquidity costs on intraday markets [Bellenbaum et al., 2014, Ehrenmann et al., 2019, Ocker and Jaenisch, 2020]. In particular, auction markets for hours far from delivery might help to increase liquidity by pooling orders, while continuous intraday markets starting close to delivery would be an optimal tool to integrate forecast errors for the output from variable renewables shortly before physical delivery. A similar design was recently introduced for the Spanish intraday market and it is planned for the Italian market as well. Alternatively, one could use a system of *frequent batch auctions* as proposed in

Budish et al. [2015], Deutsche Börse Group [2018] to combine the advantages of continuous trading and auctions.

Acknowledgements

The authors thank three anonymous referees for many insightful comments which helped to substantially improve the paper, *psaier.energies* for providing the German EPEX Spot Limit Order Book, and *Phinergy* for providing the submitted offers of the Italian intraday auctions in an easy accessible format.

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

... written in cooperation with Prof. Dr. David Wozabal, submitted, but not yet published.

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.

3.1 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.

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. [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, 2015, Kiesel and Paraschiv, 2017, Kremer et al., 2020a,b, Kulakov and Ziel, 2020],

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 [2021] 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 shear amount of order data [e.g. Glas et al., 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, 2020a], 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 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 3.2, we describe the relevant features of intraday power markets and discuss liquidity and the impact of VRES. Section 3.3 is dedicated to our trading policy. Section 3.4 describes the setting of our case study, while Section 3.5 discusses its results. Finally, Section 3.6 concludes the paper and discusses implications as well as avenues for further research.

3.2 Intraday Markets

In this section, we first describe the typical market design of continuous intraday power markets in Section 3.2.1, 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 3.2.2. Finally, we investigate market liquidity and its dependency on time to delivery in Section 3.2.3.

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, 2020]. 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 [2020]. 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, 2019]. 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*.

3.2.2 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, 2020].

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.

3.2.3 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 [2021].

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 11, 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

Figure 11: 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/).

falls to a low value close to delivery as also observed by Balardy [2022]. 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, 2020a]. The comparison of the two plots in panel 2 and 3 reveals 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.

3.3 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. [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

$$\varepsilon_t^s = f_t^{DA} - f_t^s, \tag{5}$$

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 12. 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

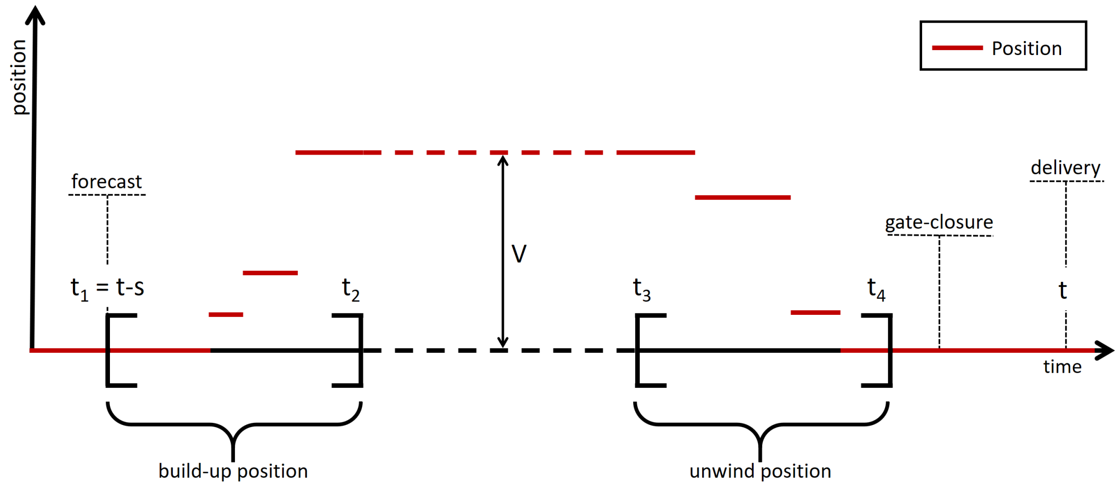

Figure 12: Schematic depiction of the sequence of events of the proposed trading strategy for the case where energy is bought.

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 $\Delta^\pm$ depending on the sign of the deviation. We thus define the traded quantity at time t_1 as

$$x_{t_1} = \begin{cases} V^+, & \text{if } \varepsilon_t^s > \Delta^+ \\ -V^-, & \text{if } \varepsilon_t^s < -\Delta^- \\ 0, & \text{otherwise,} \end{cases} \tag{6}$$

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 ±9 999€/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 δ€. 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 δ€. 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

$$\Pi_t = \sum_{\tau \in \mathcal{T}} V_\tau P_\tau + R_t \sum_{\tau \in \mathcal{T}} V_\tau - F \sum_{\tau \in \mathcal{T}} |V_\tau|, \tag{7}$$

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.

3.4 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 3.4.1 and Section 3.4.2, respectively. In Section 3.4.3, we discuss how we use the data to calibrate the parameters of our strategy.

3.4.1 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. We thus are able to accurately

	Hourly Contracts						Quarter Hourly Contracts					
	$s = 8$		$s = 5$		$s = 3$		$s = 8$		$s = 5$		$s = 3$	
$\lvert\varepsilon_t^s\rvert$	**%**	**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 8: 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.

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.

3.4.2 Weather Forecasts

In order to execute our strategies, we require the signals ε_t^s defined in (5), which are defined based on aggregated historical forecasts of solar and wind power production in Germany kindly provided by Meteologica [13]. 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.[14] Box plots of the forecast errors are provided in Figure 13. 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.

We investigate this aspect in Table 8, 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

[13] http://www.meteologica.com/

[14] https://transparency.entsoe.eu/

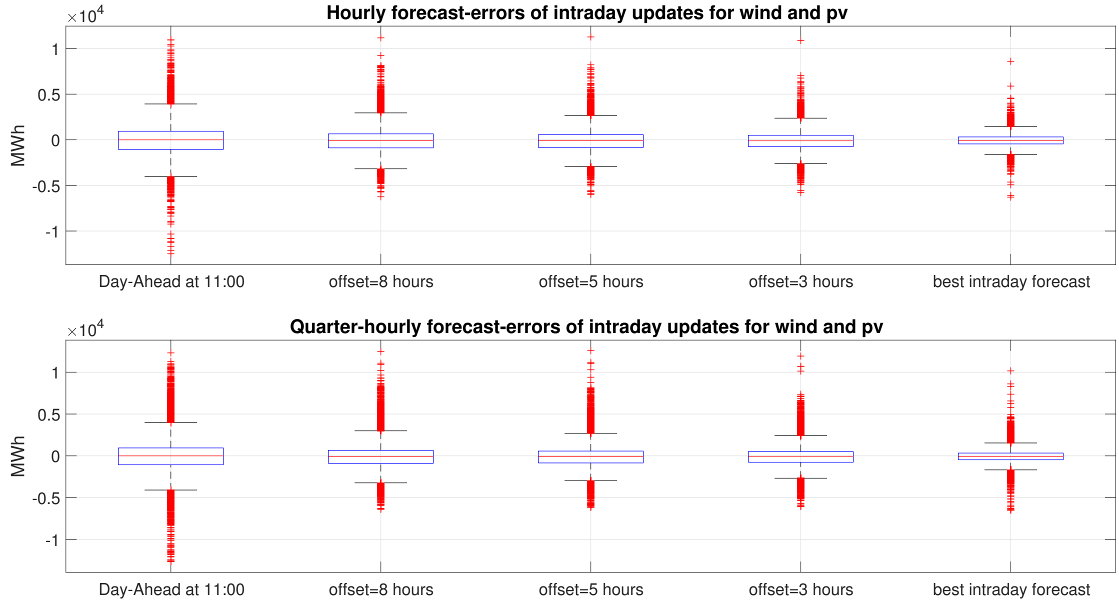

Figure 13: Forecast errors of intraday forecasts for hourly and quarter-hourly products traded on the German intraday power market between July 2017 and December 2018. The best forecast refers to the last forecast before delivery whose exact timing slightly varies with the product.

the signal, i.e., $|\varepsilon_t^s|$. In line with expectations and the results in Figure 13, 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.

3.4.3 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 3.3 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 3.4.2, 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

$$(\bar{\Delta}^\pm, \bar{V}^\pm) \in \arg\max \left\{ \sum_{d \in \mathcal{D}_1} \Pi_d(\Delta^\pm, V^\pm) : V^\pm \in \mathcal{V} \times \mathcal{V},\ \Delta^\pm \in \mathcal{L} \times \mathcal{L} \right\}, \tag{8}$$

where $\Pi_d(\Delta^\pm, V^\pm)$ is the sum of profits Π_t as defined in (7) for all products t that go into delivery on day d using the parameters $V^\pm$ and $\Delta^\pm$. For the calculation, we set the trading fees to 0.125€/MWh [EPEX, 2019] and use the quarter-hourly reBAP prices available from `https://www.regelleistung.net/` as balancing prices.

We note that the choice of $\Delta^\pm$ 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 $\Delta^\pm$

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 3.4.2. For example, if ε_t^8 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 $\Delta^{\pm}$ make the strategy react only to strong signals increasing the chance that forecast errors ε_t^0 have the same sign as ε_t^8 and are driving prices in the anticipated direction in the time window $[t_3, t_4]$. However, if $\Delta^{\pm}$ is chosen too large, then the strategy will rarely open a position decreasing overall profits. The optimization in (8) thus seeks to navigate this trade-off by choosing optimal parameters $\Delta^{\pm}$.

The second set of parameter chosen in (8) 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.

3.5 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 3.5.1, 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 3.5.2, 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.

3.5.1 Insample Results

In this section, we analyze the optimal parameter choice for $V^{\pm}$ and $\Delta^{\pm}$ 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 14 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 $\Delta^{\pm}$, 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 $\Delta^{\pm}$ and then start to fall again illustrating the trade off between frequent trading on weaker signals and infrequent trading on stronger signals.

The profits and the optimal parameter choices for the considered policies are listed in the first panel of Table 9. 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

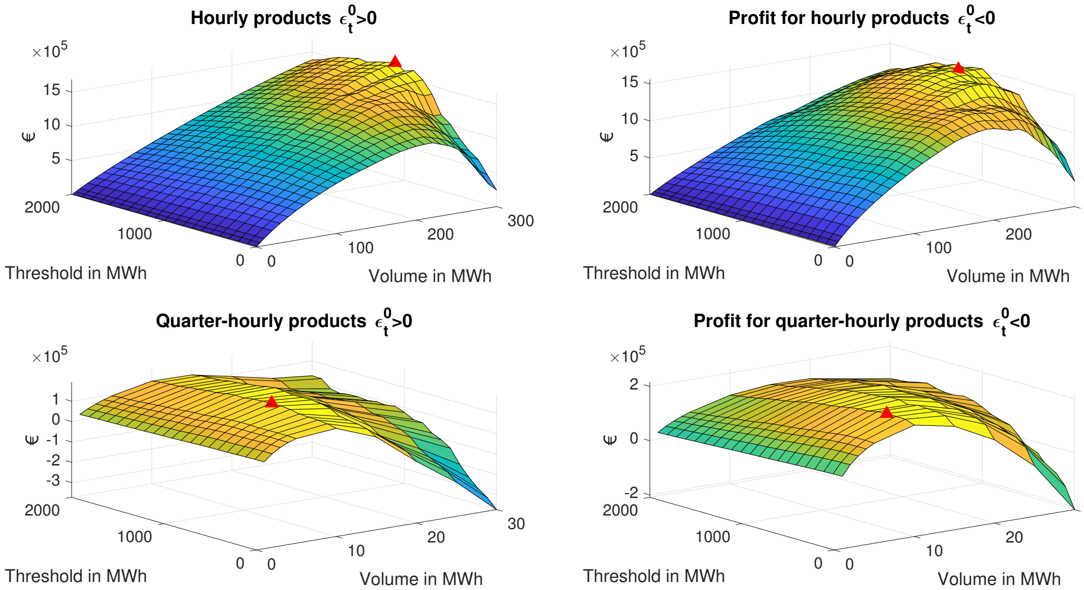

Figure 14: Optimal profits of the patient trader for real forecast errors for hourly products (above) and quarter-hourly products (below).

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 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 15. The plot exhibits many of the same characteristics as Figure 14 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 9, 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

			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^8)	Patient	QH	48438	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 9: Profits of insample strategies in € for hourly contracts (H) and quarter-hourly contracts (QH).

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 10 and indicate that, on average, the strategy requires a negative amount of capital with low positive maximal values. The profits displayed in Table 9 can therefore be realized with a small amount of risk capital and offer a high return on investment.

3.5.2 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 strat-

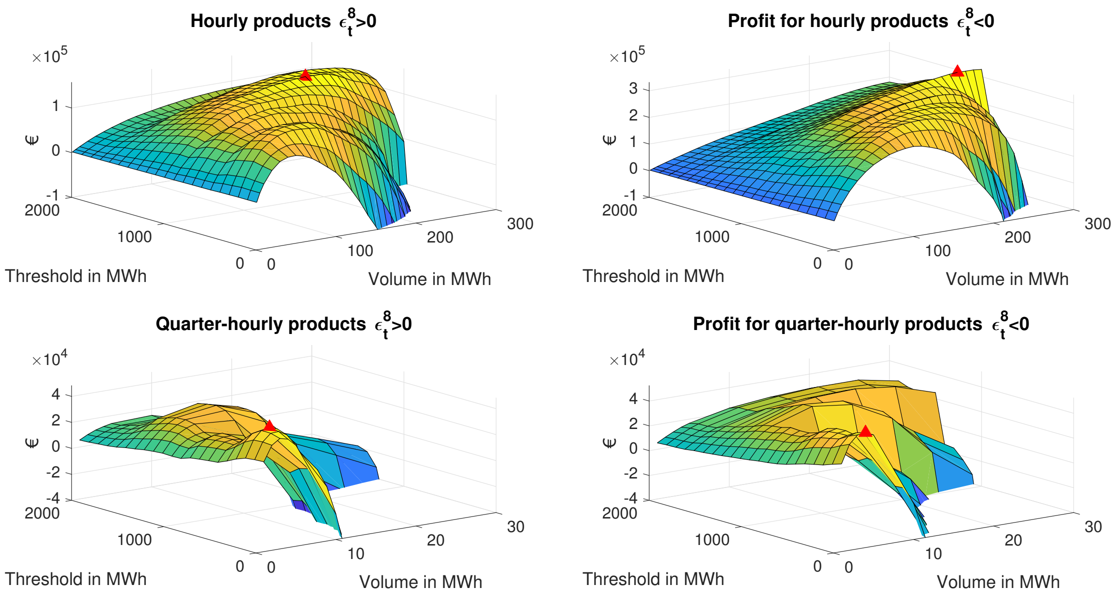

Figure 15: Optimal profits of the patient trader for forecast errors with an offset of 8 hours for hourly products (above) and quarter-hourly products (below).

egy and re-optimize the parameters $\Delta^{\pm}$ 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 $\Delta^{\pm}$ and $V^{\pm}$ by grid search as in (8). 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 16 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

			Mean	**Max**	**Min**	**Std**
Hour	**Patient**	ε_t^8	-22 163	5798	-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	5277	-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 10: Amount of net capital invested per day for the different trading strategies.

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 16, 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 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 16 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 17 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

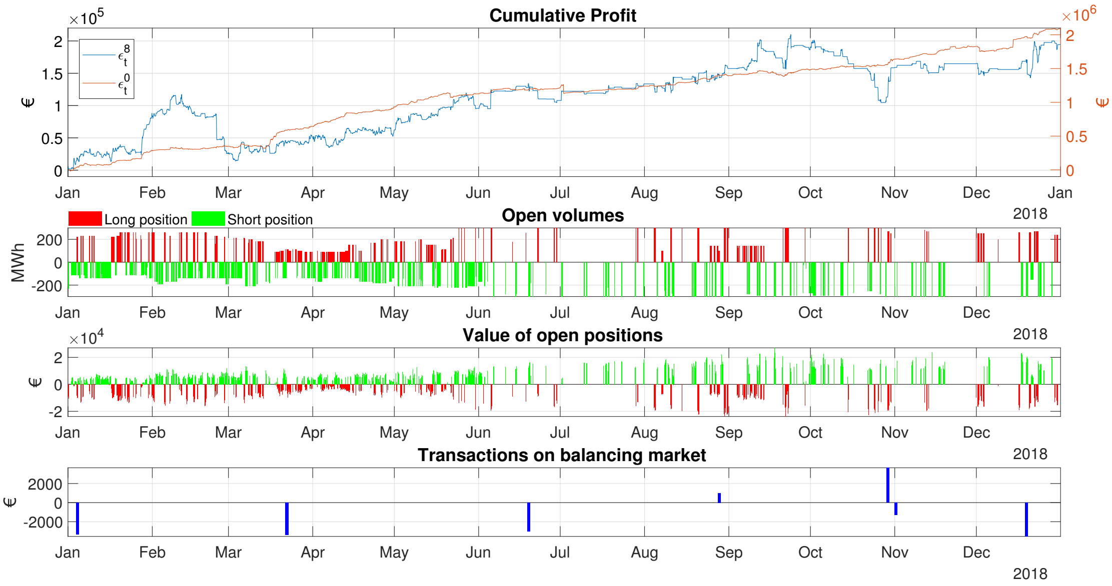

Figure 16: Cumulative profit of the optimal insample and out-of-sample strategy for hourly products in the first panel. Panels 2 and 3 display the opened volume and the financial value of the positions held by the out-of-sample strategy. Panel 4 shows daily payments on the balancing market for the out-of-sample strategy.

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

Table 11 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

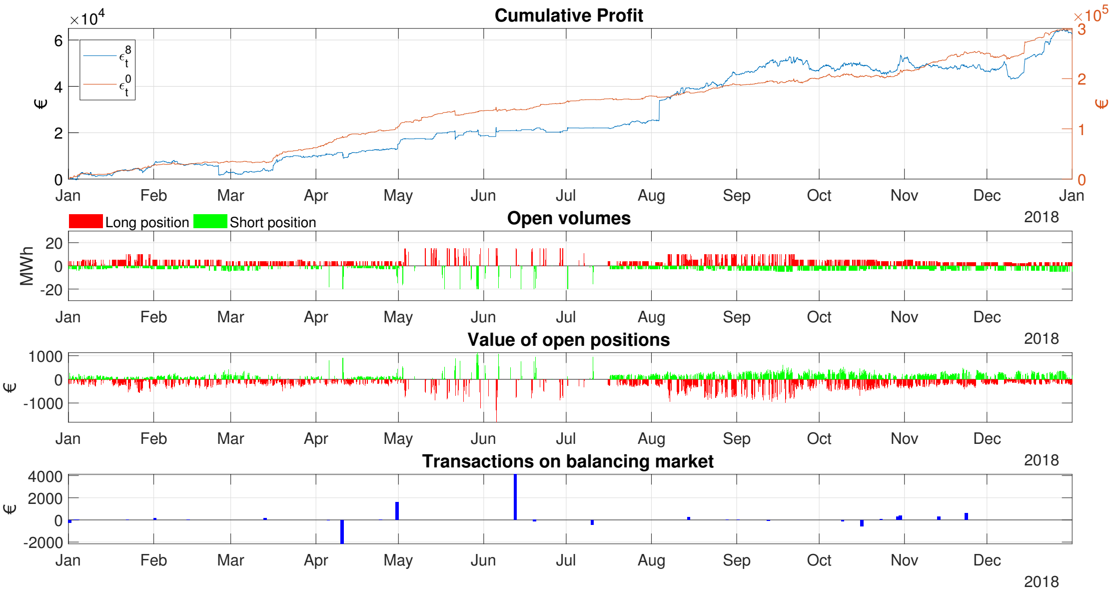

Figure 17: Cumulative profits, traded volumes, value of traded positions, and daily balancing payments for quarter-hourly products (see Figure 16 for a more detailed description of panels).

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.

3.6 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

	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	1824	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 11: Descriptive statistics for the profits of different strategies and the number of traded products and trades.

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 Giacomin who was enormously helpful in obtaining and understanding the forecasts.

4 Frequent Auctions for Intraday Electricity Markets

... written in cooperation with Prof. Dr. David Wozabal and Dr. Christoph Graf, submitted, but not yet published.

4.1 Introduction

In the last decades, electricity markets in most countries have seen fundamental changes due to the transition of the electricity sector from a vertically integrated, state controlled sector of the economy to a competitive industry. In addition, the electricity sector is the key to sustainable energy systems, facilitating a substantial increase in the use of renewable energy and consequently the phase-out of fossil fuels.

Electricity markets are typically organized as a sequence of future markets that trade products with ever shorter maturity and temporal resolution. Most market designs feature a day-ahead market that allows to trade electricity one day ahead of delivery and a market that gives firms the possibility to adjust their positions until shortly before physical delivery.

The latter markets — organized either as real-time or intraday markets — are of increasing importance because of the growing short-term uncertainty in production from renewable energy sources. Prominent examples for real-time markets include most US electricity markets [Milligan et al., 2016], while European short-term markets are organized as intraday markets.

In this paper, we focus on intraday markets. There are currently two prevailing designs for intraday trading [Ocker and Jaenisch, 2020]: continuous markets and repeated auctions. Both designs have their strengths and weaknesses: Auction based markets impose low entry barriers for participating firms and facilitate relatively high liquidity by pooling demand and supply. However, existing auction markets suffer from long lead-times making it difficult to trade the production of renewable energy sources and to quickly react to new information [Hannele Holttinen, 2013], mostly because there are only a handful of auctions with the last one closing several hours before delivery.

In contrast, continuous trading ensures a high level of *immediacy* as traders can instantaneously act on new information. Furthermore, since a trader can, in theory, accept orders for different products with known prices simultaneously, the continuous market makes it easier to trade complex profiles in asset backed trading. The main downsides of continuous trading are the higher complexity of trading, the lower liquidity [Kuppelwieser and Wozabal, 2021], which leads to low quality price signals that are often dominated by noise, and the incompatibility of order-book based trading with the physical realities of the electrical grid. Despite these downsides the recent trend is to discontinue auction based designs in favor of continuous trading [Ocker and Jaenisch, 2020].

Consequently, continuous intraday markets have recently attracted some attention in the academic literature. Weber [2010] analyzes the integration of wind energy considering different European market designs and finds that the intraday auctions in Spain are the most attractive in terms of trading volume. Balardy [2022] analyzes liquidity in terms of bid-ask-spreads and market depths. Baule and Naumann [2021] study the volatility of intraday markets as well as the drivers of price fluctuations on the German intraday market. Furthermore, there is a large literature on forecasting models and the identification of suitable covariates for intraday prices [e.g., Kiesel and Paraschiv, 2017, Uniejewski and Weron, 2018, Janke and Steinke, 2019, Narajewski and Ziel, 2020b, Marcjasz et al., 2020].

In this paper, we investigate a market design, which we believe is well suited to deal with the idiosyncrasies of electricity markets and represents a compromise between the extremes of continuous trading and *infrequent* auctions. In particular, we propose that orders should be batched in *frequent* auctions which are repeatedly conducted for every traded product until briefly before physical delivery of electricity starts. Since the intervals between auctions would be small, the proposed market format can be considered a hybrid between auctions and continuous trading. We argue that a frequency of 15 minutes to an hour strikes a good balance between immediacy and liquidity and thus yields lower transaction costs and more reliable price signals. Ideally, such a market has the potential to combine the advantages of both designs while avoiding most of the disadvantages.

Our proposal is motivated by general results in the finance literature and specific findings pertaining to the electricity markets. In the finance literature, which is reviewed in more detail in Section 4.3.1, there is a large body of literature that argues that the quality of price discovery benefits from either trading financial assets in auctions or at least complementing continuous trading by auctions [Schwartz, 2012]. Furthermore, there is evidence that, in particular thinly traded stocks with comparably less volume, benefit from the shift to auction based trading [e.g., Hu and Chan, 2005, Hu, 2006]. This provides an argument for the use of auctions in intraday electricity markets, which are characterized by rather low trading activity at times. Furthermore, in an independent stream of research Budish et al. [2015], Aquilina et al. [2021] argue that a switch to frequent auctions eliminates the excesses of high frequency trading and thereby increase welfare.

Additionally, there are several authors that discuss the specific advantages of auctions for electricity intraday markets. Neuhoff et al. [2016] studies the impact of the intraday auction in the German market that clears one day before delivery and find that the auction has a higher liquidity and lower volatility than continuous trading. Furthermore, the authors argue that auctions are better suited for smaller players that do not have the capacity to take part in a continuous intraday market.

Similarly to the market design proposed in this paper, Deutsche Börse Group [2018] proposes a model for frequent intraday auctions that takes into account transmission infrastructure to explicitly price scarce interconnector capacities, allows for more complex order types, and increases liquidity. The authors argue that these goals can be achieved in auctions, due to the increased time for clearing and the possibility to take into account orders at different locations across the network.

Ocker and Jaenisch [2020] discuss continuous trading and auction based intraday markets in the European context and identify liquidity, the resilience against the exercise of market power, and efficiency of the use of transmission capacity as the main advantages of auctions over continuous trading.

To get an idea how frequent auctions could impact market outcomes, we conduct a case study for the German market by creating a counter-factual for auction outcomes based on detailed order-book level data submitted to the EPEX continuous intraday market for the German market zone in the years 2017 and 2018. To this end, we construct hypothetical auction outcomes for a single auction per product as well as auctions with hourly and quarter-hourly frequencies. Our results show that the distribution of volume weighted prices remains virtually identical when switching to any of the proposed auction formats. When examining the traded volume, we are able to show theoretically that under certain conditions auctions clear less orders and therefore lead to lower traded volumes than continuous trading. This theoretical result is largely confirmed in our numerical experiments.

Despite the results on traded quantities, liquidity costs measured as costs of round trip trades as in Kuppelwieser and Wozabal [2021] are lower for auction based trading than for continuous trading. Based on these results, we argue that even though auctions tend to clear less volume, they are preferable in terms of liquidity cost.

Finally, we use a kernel regression based approach to investigate the signal to noise ratio of the price signals generated by the two market designs. We find that prices generated by frequent auctions are significantly less noisy than the prices resulting from continuous trading and are therefore expected to produce more reliable price signals that are more closely tied to changes in the fundamental value of the traded product.

The rest of the paper is organized as follows: In Section 4.2, we review the current European market design with a special emphasis on the German market, which we use in our case study. In Section 4.3, we review the trade-off between continuous trading and auctions as it is discussed in the finance literature and argue that frequent auctions have the potentials to combine the advantages of both designs. Section 4.4 details the computation of the counter-factual based on the order book data for the German market, while Section 4.5 discusses the numerical outcomes of the comparison between continuous trading and frequent auctions. Section 4.6 summarizes and concludes the paper.

4.2 The German Electricity Market Design

In this section, we review the current German design for short-term electricity markets. We first give a broad overview of all the markets and their integration in the larger European context in Section 4.2.1 and then discuss the specific market rules for intraday trading in Section 4.2.2.

4.2.1 Overview of German Short-Term Electricity Markets

The German short-term market for electricity is embedded in the wider zonal European market design, which is organized as a cascade of forward markets with the day-ahead market being especially important. Specifically, the day-ahead market determines schedules for European cross-border flows via the European single day-ahead coupling and consists of several bidding zones in which Germany represents a single zone.[15] Trading with other zones is defined through "flow-based" market coupling.[16] The market is organized as a double auction that yields locational marginal prices and trades products for delivery in every hour of the following day.

The main philosophy of the German market design is that market participants should deliver on their day-ahead promises in real-time at the firm level [Cramton, 2017]. However, because demand and non-dispatchable supply are uncertain, market participants have the opportunity to minimize their real-time imbalances at the intraday market that opens shortly after the day-ahead market has cleared.

In Europe, there are currently two competing types of intraday trading systems: auction markets and continuous intraday trading. In 2015, the EU committed to 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]. In Germany the intraday market is hosted by the EPEX, the largest electricity exchange in Europe [see Viehmann, 2017, for a detailed description] and is organized as a continuous trading market.

[15] See https://www.entsoe.eu/network_codes/cacm/implementation/sdac.

[16] Because the European day-ahead market is a zonal market and not a nodal market as it is the case in the United States, the flow-based market coupling intends to approximate trade flows between zones according to the impedance of the network.

Similar to day-ahead markets, German intraday markets are also coupled with other intraday markets across Europe via the *single intraday coupling* (SIDC). Orders of each bidding zone are collected in local limit order books, and cross-border capacities are used to build shared limit order books that are used to match orders from different zones.[17] As of today, cross-border flows in SIDC are agnostic towards Kirchhoff's and Ohm's laws and are computed according to shortest paths between nodes of the network.[18] Consequently, the power flows computed in SIDC might not match actual physical flows leading to potentially costly re-dispatch by the transmission system operator (TSO).

Unlike the electricity markets in the United States, the German market design can be described as trader-centric, because day-ahead and intraday market offerings are not tied to physical units but to traders. In a separate process market participants with dispatchable units communicate with the TSO about which units at which locations are scheduled. More generally, the European take on electricity market design aims to decouple the trading reality from the physical reality. Because the latter has to be taken care of in order to avoid system failure, complex, spatially heterogeneous, and partially intransparent processes are operated to ensure that electricity supply can be maintained even under circumstances where the trading decisions produce physically infeasible market results.

Germany requires market participants to firmly report schedules at 2:30 pm one day ahead of delivery.[19] These schedules must be within the range of the maximum physical withdrawal or injection capacity which limits arbitrage trading between the day-ahead and intraday markets. Furthermore, this framework effectively bans speculative traders that do not have a natural short or long position from participating in the day-ahead market. Nevertheless, the price differences between day-ahead market and intraday-market is typically small on average. 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-spread between the day-ahead market and the balancing market is often significant. Balancing prices are determined by distributing costs of reserve call offs to those market participants that caused the imbalance. The German regulatory framework forbids to arbitrage between the day-ahead or the intraday market and the balancing market by creating imbalances on purpose.[20] In practice, this is less strictly enforced for producers of renewable electricity and the demand side

[17]See https://www.entsoe.eu/network_codes/cacm/implementation/sidc/.

[18]See https://www.nemo-committee.eu/assets/files/public-description-of-the-continues-trading-matchi pdf for a more detailed description of the matching algorithm.

[19]See Stromnetzzugangsverordnung (StromNZV), §5(1), https://www.gesetze-im-internet.de/stromnzv/__5.html.

[20]See https://www.amprion.net/Strommarkt/Bilanzkreise/Bilanzkreisvertrag/.

because of the difficulties in precisely predicting production and load.

Capacity for reserve call offs is procured in separate auctions which are held every day for the respective next day before the day-ahead market is cleared. TSOs then *call off* positive or negative reserve energy in real time to ensure that planned power flows are physically feasible. The revenue stream of a supplier eligible to participate in the balancing market thus consists of (i) selling energy in the day-ahead and intraday markets, (ii) arbitraging between day-ahead market and intraday market, and (iii) providing the flexibility to adjust prior schedule commitments in real-time that are needed to balance the system if other market participants are not able to stick to their commitments. Therefore all these markets are implicitly coupled with each other.

Although in the day-ahead market transmission constraints are not explicitly accounted for *within* Germany, these can still be relevant and the TSO might have to intervene by redisaptching some power plants to ensure feasible flows within the German zone. This invites what is called the INC/DEC game [Graf et al., 2020], where market participants with dispatchable capacity may be able to profit from the discrepancy between the clearing result from, for example, the day-ahead market and the actual real-time demand for their supply units.[21] Using this strategy market participants can endogenously create a demand for re-dispatch by submitting — from a system perspective — unfavourable schedules. To limit the profitability and therefore the attractiveness of this strategy, Germany uses a cost-based re-dispatch [Hirth and Schlecht, 2019].

4.2.2 The German Continuous Intraday Market for Electricity

The German continuous intraday market opens shortly after the clearing of the day-ahead market and is organized as an order book based continuous trading market that features hourly, half-hourly, 15-min, as well as block products. The intraday market remains open until 5 minutes before the delivery of the respective product starts. However, the German market zone is split up into four zones, one per TSO, 25 minutes before gate closure, even if no actual congestion is present.

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

[21] INC stands for incremental energy and DEC for decremental energy. The game is played by market participants that choose to schedule their units in the day-ahead market such that their re-dispatch revenues are maximized. Because markets closer to real-time are typically less competitive, especially if they are local, this can be a profitable strategy and is possible if relevant constraints that matter in real-time are ignored in the day-ahead market-clearing.

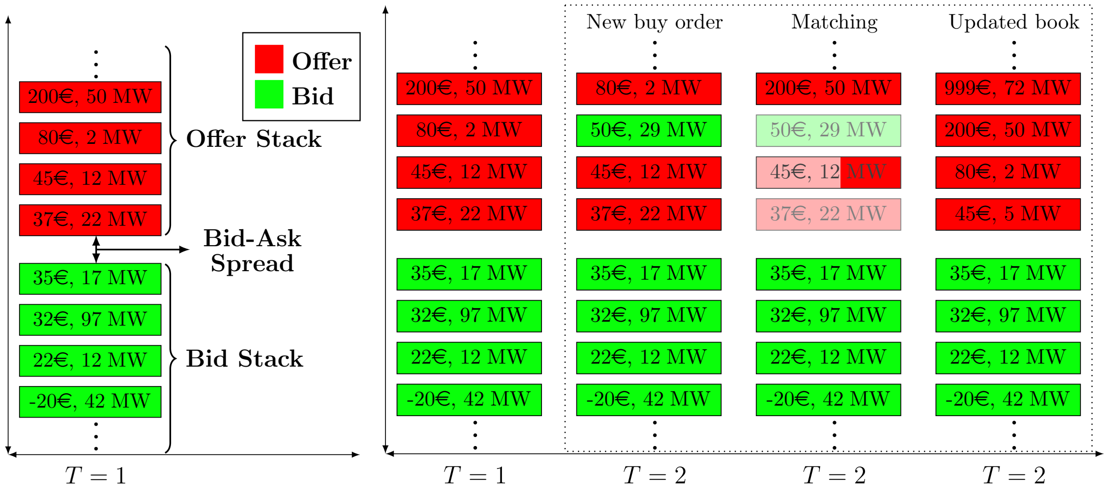

Figure 18: Example of the state of the order book at $T = 1$ in the left panel. Clearing of a new bid in the right panel.

is kept in the LOB 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 the remaining quantity.

Market participants can add the usual order qualifiers such as *immediate-or-cancel* (IOC) or *fill-or-kill* (FOK) [EPEX, 2019].[22] 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 basic mechanism of continuous trading is illustrated in Figure 18 by a concrete example: In the left panel, the state of the order book at $T = 1$ is displayed with the orders sorted according to their limit price. 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*.

The dynamics of the order book are exemplified in the right panel of Figure

[22]IOC is either executed immediately, or, if the order cannot be matched, deleted without entry in the order book. FOK is either executed immediately and with its full quantity or, if the order cannot be matched with its entire quantity, deleted without entry in the order book. Note that the difference between IOC and FOK is that an IOC order can be partially cleared. See `https://www.nemo-committee.eu/assets/files/public-description-of-the-continues-trading-matching-algorithm-.pdf`.

18: A bid with a price higher than the lowest ask is added to the book at $T = 2$ and then cleared against the cheapest possible offers until either the whole order is fulfilled (as is the case in the figure) or there are no offers with lower prices left. Note that the clearing is instantaneous, i.e., columns $2 - 3$ of the figure are purely illustrative and do not correspond to market states that can be observed by traders.

4.3 Frequent Auctions for Intraday Power Markets

In this section, we discuss the literature on the trade-off between continuous and auction based trading in Section 4.3.1 and then proceed to discuss how the specifics of the electricity market influences the choice between the two market designs in Section 4.3.2.

4.3.1 Continuous Trading versus Auctions in the Finance Literature

The finance literature discusses the choice between auction based trading and continuous markets as a trade-off between liquidity and quality of the price signal versus the possibility to react quickly to new information. Continuous trading is at one extreme of this trade-off, offering maximal immediacy. Having a single auction is the other extreme which would be, by definition, welfare optimal if there is no information flow in the time the auction market is open.

Although continuous trading is the prevailing market design in all major markets trading shares, futures, options, and other financial products, there are prominent critical voices in the finance community. Most of these authors advocate to replace or complement continuous trading by auctions in order to either improve liquidity and the quality of the price signal [e.g., Schwartz, 2012] or to avoid the excesses of high frequency trading that effectively imposes a fee on trading and thereby lead to welfare losses [e.g., Budish et al., 2015].

Kregel [2001] provides an account of the historical roots of continuous trading and points out that in early stock exchanges trading was organized in auctions. However, since, without computer technology, auctions had to be performed sequentially in order to give every trader the chance to participate in all auctions, auctioning became impractical as the number of listed companies increased. To resolve this problem, trading floors where brokers could continuously strike bilateral deals evolved. This form of trading eventually lead to continuous trading as it is practiced today. Paradoxically, nowadays the use of modern information technology, whose absence necessitated this development, is to blame for most of the adverse effects of continuous trading.

One of the advantages of continuous trading is its immediacy, i.e., the ability of market participants to instantaneously trade on information whenever it arrives.

Correspondingly, in a classic paper, Brennan and Cao [1996] show that, if the timing of information is not predictable and specific assumptions on a *demand for immediacy* and risk aversion of market participants hold, continuous trading yields a Pareto efficient equilibrium that is preferable to auctions.

However, in a survey of equity traders Economides and Schwartz [2001] find that there is no fundamental economic reason to execute trades within seconds or minutes. Based on these findings, Steil [2001] argues that the demand for immediacy is endogenous to continuous trading and would almost entirely vanish if trading would take place in auctions. In particular, traders want to react fast to either profit from being the first to react to new information or to avoid *front-running* of other market participants. When trading is organized via auctions, the notion of being first looses its meaning and front-running cannot occur. Furthermore, Economides and Schwartz [2001] find that the lack of liquidity in continuous trading might actually reduce immediacy for larger positions, since traders are forced to trade *patiently* over longer time periods to avoid an excessive price response to their trades.

Recently, Budish et al. [2015], Aquilina et al. [2021] take a slightly different angle in criticizing continuous trading when study the impact of high frequency trading on market outcomes. The authors observe that treating time as continuous in trading systems that serially process orders opens the door for *latency arbitrage* where high frequency traders compete on speed to capitalize on new information that signals a change in the fundamental value of traded assets by *sniping stale orders* of other market participants that are slower to react. The authors argue that the resulting arms race for speed is socially wasteful and effectively introduces a fee on trading that reduces liquidity.

In particular, Aquilina et al. [2021] find that frequent batch auctions would reduce the cost of liquidity by 17% and that a remarkably large portion of overall trading volume (about 20%) can be attributed to latency-arbitrage races. Similarly, Wah and Wellman [2013] find in a simulation study that replacing continuous markets with periodic call markets eliminates latency arbitrage opportunities and achieves substantial efficiency gains.

There is a strand of literature that explores the trade offs between continuous trading and repeated auctions. An early paper in this direction is Garbade and Silber [1979] who investigate how the frequency of trading influences liquidity risk and identify two opposing effects: Firstly, longer auction periods help to collect more participants in an auction and therefore reduce the noise in the price signal and therefore liquidity risk. Secondly, they identify the drift in the equilibrium price as driving the volatility risk up as auction intervals become longer and the shocks in equilibrium prices increase between to consecutive auctions. Hauser et al. [2001] extend this argument and find that more liquid stocks tend to have less return volatility when continuously traded whereas discrete trading with longer intervals is preferable for thinly traded stocks.

Du and Zhu [2017] propose a model of sequential double auctions, which allows to capture the trade-off between the welfare increasing ability to react quickly to information changes and welfare decreasing bid shading both of which increases in the frequency of trading. In the model, frequent trading reduces liquidity (and thus welfare) in every single auction, but increases welfare by reducing externally assumed holding costs of agents.

Complementing this theoretical research, there is a large body of empirical literature that informs the discussion based on observational data from stock markets. Pagano and Schwartz [2003] show that complementing continuous trading with call auctions lowered execution costs and improved price discovery in the Paris stock exchange. Comerton-Forde et al. [2007] find similar results for the Singapore stock exchange.

On the other hand, Muscarella and Piwowar [2001] find that for assets traded on the Paris stock exchange, the traded volume increases when stocks are shifted from auction based trading to continuous trading and decreases when the shift goes in the other direction.

Twu and Wang [2018] show in a case study of the Taiwan stock exchange that decreasing the interval between consecutive auctions improves overall market quality. Contrary to these findings, Hu and Chan [2005], Hu [2006] find that shorter intervals correspond to a worse signal-to-noise ratio in prices at the exchange.

Lauterbach [2001] shows that for most stocks on the Tel-Aviv stock exchange liquidity and the quality of the price signal improved after being shifted to continuous trading but also identifies some exceptions of stocks that are thinly traded and for which continuous trading does not work well. Finally, Chelley-Steeley [2008, 2009] finds that market quality on the London stock exchange improves with the introduction of a closing call auction and that this improvements are especially pronounced for the least actively traded securities.

Overall, the emerging picture is ambiguous. However, it seems fair to conclude that thinly traded stocks tend to profit from auction based formats. This suggests that intraday markets for electricity, which, except for a brief period shortly before delivery, are characterized by rather low activity could profit from an auction based design.

4.3.2 Auctions versus Continuous Trading in Electricity Markets

In this section, we describe key differences between electricity markets and financial markets, discuss the applicability of continuous trading in electricity markets, and highlight where auctions might have advantages or disadvantages.

A continuous market design allows to trade immediately and all current orders are visible to each trader (*transparency*). However, since intraday markets for electricity are, compared to most financial markets, rather thinly traded and many market participants do not have the ability to react fast to new information, the loss in immediacy is likely to be marginal when switching to a market design with frequent auctions.

In an auction the submission time of orders is irrelevant. Therefore auctions help to create a level playing field without advantages for those who are able to act quickly and therefore prevent front-running and other costly excesses of high frequency trading as they are observed in financial markets. Furthermore, the technical complications and high fixed cost of operating a trading desk that participates in continuous trading at competitive speeds might hinder market entry of some firms. Hence, an auction-based market might facilitate more participation and ultimately a more liquid market [Ocker and Jaenisch, 2020].

Another difference between the two market designs is price variance and the quality of the price signal, i.e., the information about fundamental values that is contained in the price. While an auction based market design with too few auctions disseminates information slowly and might lead to large price shifts between auctions [see Garbade and Silber, 1979], a design with frequent auctions can be used to address these issues. Furthermore, by pooling orders in auctions extreme bids and offers are likely to be infra-marginal and therefore do rarely directly influence the clearing price. Because individual orders in an auction are not observable to market participants, the decision on how to trade would likely be more influenced by observable shocks such as changes in forecast for demand or renewable production than by the actions of other market participants. Taken together this shifts the focus of traders away from an introspective view on the markets towards fundamental factors, which potentially reduces trading-induced noise in the price signal.

One of the advantages of continuous trading is that firms can more easily trade complex profiles involving more than one product. Consider the example of an electricity storage that wants to buy electricity in one period, store it, and then sell it at a later point in time. In continuous trading the storage can accept existing limit orders for two products at the same time for known prices and thus minimize the risk that only one leg of the transaction is executed. Contrast this to the same situation in an auction setting, where it could very well happen that the bids of the trader are only accepted in some of the auctions leaving the storage with an open position. This situation can be circumvented by submitting perfectly inelastic supply or demand functions with the effect that the storage owner may not receive the arbitrage value she had hoped for [see also the discussion in Löhndorf and Wozabal, 2022].

Next we discuss to what extent the market design can deal with technical properties of the underlying electricity network and how this might affect the choice of the trading system. In particular, we remark that electricity markets are special because generators, consumers, and traders interact through a rather unforgiving physical system whose failure induces significant negative external effects on societies as a whole [Kirschen, 2021]. That is the reason why effectively *all* electricity markets have a transmission system operator as a centralized authority that has the objective to secure continuous electricity supply for as many customers as possible.

In electricity markets, network constraints that result from the aforementioned physical limitations introduce complications in market-clearing because the resulting power flows have to be computed for every trade. Flow-based market coupling reflects congestion in the network and assigns prices to transmission capacities accordingly and therefore would be the first-best solution. However, the computational performance of the ensuing calculations depend non-linearly on the size of the network and can cause latency in continuous trading, which would slow down trading and decrease the advantage of immediacy. Hence, if power flows would be accurately computed while clearing the market, this would effectively act as a *speed limit* for continuous trading.

These leaves two options to price transmission capacities in continuous trading: (i) exogenously defined locational markets, which obviously is inefficient because cross-border trading is restricted even if market conditions would allow for it, (ii) transmission capacity is allocated for free on a first-come first-served basis [Ehrenmann et al., 2019]. The latter approach aggravates the potential race for speed in continuous markets. Summarizing, it is highly unlikely that a sequence of bilateral trades through continuous trading is able to solve the complementary goods problems for energy and transmission capacity [Mansur and White, 2012, Ehrenmann et al., 2019].

In contrast, flow based market-coupling can easily be accounted for in an auction that centrally maximizes welfare subject to transmission constraints. In this setting, locational prices can be easily constructed from the duals of the energy balance constraint and the duals of the transmission constraints [see, e.g., Graf and Wolak, 2020]. The advantage is that the price of congestion will be defined endogenously based on orders. At this point it is important to highlight that in meshed electricity networks with loops, transferring 1 MW of electricity from one node to another affects every single flow in the network. Consequently, clearing the market jointly with many different offers and bids at different locations will increase the efficiency of the market outcome [Ehrenmann et al., 2019]. Even more important in terms of computational complexity are constraints that link several time instances [Neuhoff et al., 2016]. Both attributes favor auctions because computation time is less critical compared to continuous trading markets [see also

Deutsche Börse Group, 2018].

We also want to emphasize that because bids and offers are cleared sequentially in continuous trading, liquidity is critical when blocks of energy are traded. More precisely, in continuous trading a block-offer will only be sold conditional on the presence of a market participant willing to buy a given block bid. In practice this constraint can lead to severe liquidity crunches. In contrast, in auctions block-offers can easily be traded against multiple simple demand bids, if auctions for several products are cleared simultaneously [Ehrenmann et al., 2019].

Finally, we remark that with continuous trading dispatchable units that participate in the reserve markets may use schedule changes ordered by the transmission system operators to front-run the continuous intraday market. In particular, if market operators with dispatchable units obtain the information that one of their units' output will be changed only a fraction of a second before the public, this information can be monetized on the continuous intraday market. If an auction would take place, e.g., every five minutes, the timing of when the information of a schedule change will be released is less critical and all market participants are in the same position to react to this information update. Hirth and Mühlenpfordt [2021] empirically study this aspects for the German intraday market where balancing orders are not published in real-time. The authors find a statistically robust correlation between reserve call offs and intraday price changes.

4.4 Counter-Factual Frequent Auction Design

To assess how a transition to frequent auctions affects intraday electricity prices, we construct counter-factual market results based on observed order book data from the German continuous market. We conduct a ceterus paribus analysis where the market design changes, but the orders of the market participants stay the same, i.e., we use the historical orders submitted to the continuous market as hypothetical auction bids. For the purpose of this paper, we consider hourly and quarter-hourly products traded at the EPEX and disregard the half-hourly products in order to keep our numerical study manageable.

In Figure 19, we visualize a stylized sequential auction design for one product in a single-zone market. The main idea is that a uniform price auction is run repeatedly with the last auction shortly before delivery and the first auction several hours before that. Auctions should be frequent enough to allow for a timely reaction to new information but leave enough time between auction clearings in order to ensure sufficient participation – and thereby liquidity – in every single auction. We therefore envision hourly or quarter-hourly auctions that are run simultaneously for all traded products. We believe that the flow of fundamental information pertaining to electricity supply and demand is such that this design does not represent a relevant limitation of immediacy in trading. For the sake of

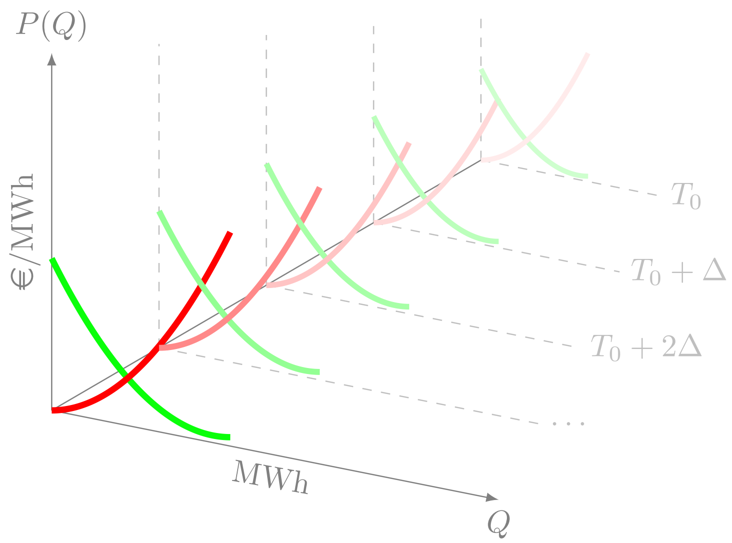

Figure 19: Stylized Sequential Batch Auctions for one Product

simplicity, we propose the auctions to be evenly-spaced in time. After an auction clears, the next one immediately opens and bids and offers are collected by the market operator until gate closure, which is the time the auction is cleared.

To construct auctions, we fix a product, i.e., a specific hour or quarter-hour in which electricity is delivered and an auction interval $\Delta \in \{15, 60\}$. An order is characterized by the vector $(q, p, t_0, t_1)^\top$ consisting of a quantity $q \in \mathbb{R}^+$, a limit price p, a start validity date t_0, and an end-validity date t_1. For $\Delta = 15$ the first auction opens at $T_0 = 16{:}00$ for quarter-hourly products and at $T_0 = 15{:}00$ for hourly products while for $\Delta = 60$, the first auction opens in such a way that we obtain the last clearing 30 minutes before delivery.[23] For every auction k, we then collect all offers $\mathcal{O}_k$ and bids $\mathcal{B}_k$ that are submitted to the continuous market and are valid in the period $(T_0 + (k-1)\Delta, T_0 + k\Delta]$, i.e., whose end of validity date is larger than $(T_0 + (k-1)\Delta)$ and which have not been cleared in a previous auction. For the sake of simplicity, we exclude block orders from the analysis.

For every auction k, the cleared quantity is defined as the maximizer q^* of the following welfare optimization problem

$$\max_q \int_0^q D_k^{-1}(x) - S_k^{-1}(x)\,\mathrm{d}x, \tag{9}$$

[23]We note that in a practical implementation it would be desirable to trade even closer to delivery, say 5 minutes. However, since 30 minutes before delivery the Germany wide order book is replaced by 4 order books for smaller zones, we stick to 30 minutes in our experiment.

where D_k^{-1} and S_k^{-1} are the inverse demand and supply functions. If D_k^{-1} and S_k^{-1} would be continuous and strictly monotonously decreasing and increasing, respectively, q^* would be uniquely defined as the point where the two functions intersect, i.e., $S_k^{-1}(q^*) = D_k^{-1}(q^*)$.

In the single zone case and with no product dependencies the market-clearing thus simply involves constructing D_k^{-1} and S_k^{-1} by sorting offers in ascending order and bids in descending orders and finding their intersection which yields the market-clearing price as well as market-clearing quantities. To this end, we define the following step functions as bid and offer curves

$$D_k^{-1}(q) = \max\left\{p : B \subseteq \mathcal{B}_k \text{ with } p^b \geq p,\ \forall b = (q^b, p^b, t_0^b, t_1^b) \in B,\ \sum_{b \in B} q^b \geq q\right\} \tag{10}$$

$$S_k^{-1}(q) = \min\left\{p : O \subseteq \mathcal{O}_k \text{ with } p^o \leq p,\ \forall o = (q^o, p^o, t_0^o, t_1^o) \in O,\ \sum_{o \in O} q^o \geq q\right\}. \tag{11}$$

When dealing with these piecewise constant functions, there is some ambiguity about the clearing price as defined in (9) and, in some cases, also the cleared quantity that maximizes welfare. We resolve this ambiguity in the same way, as it is done in the current design of the Italian intraday auctions [Caramanis and Inc., 2002, Graf and Wolak, 2020].

In particular, if the piece-wise constant functions (10) and (11) intersect either at one point q^* or on a vertical segment, then the cleared quantity is uniquely defined. In both cases the price is a function of the price of the last bids and offers that are cleared, i.e., the highest bid and lowest offer price. In the former case, we choose the price of the order (either bid or offer) that is not fully cleared and in the latter case the price is defined as the price of the last accepted offer.

In case that the functions intersect at a horizontal part of the functions D_k^{-1} and S_k^{-1}, the price is uniquely defined and the quantity is chosen maximally, i.e., such that the smaller of the last orders is fully executed. We denote the resulting prices and quantities for the k-th auction for given product by p_k^A and q_k^A.

In case an offer or bid is not cleared or only partially cleared but has an end validity date that extends into the next auction period, we place the remaining quantity as an order in the next auction. In particular, for the k-th auction, we include orders with an end-validity date after $T_0 + (k-1)\Delta$ as active orders. Orders with longer validity times can therefore be considered in more auctions until their available volume is fully cleared. We furthermore propose that market participants should be able to cancel or modify their offers or bids until shortly before the auction is cleared and accordingly match the order changes observed in continuous trading in our hypothetical auctions.

Clearly, while the obtained prices mimic the results of a hypothetical frequent auction market, the analysis has several limitations, which are discussed in the following:

1. The experiment suffers from the issue that we use orders that were submitted to a continuous market to construct counter-factual auction results. This is of course suboptimal, since it can be expected that market participants would change their bidding behavior if the market design would change to frequent auctions.

   However, under the assumption that limit orders in the continuous market reflect the marginal willingness to pay/sell and the same holds true for the auction, the difference in bidding behaviour is limited to differences caused by technical rules of trading. Hence, in an efficient and competitive market where players do not have the possibility to bid strategically and game the system, bidding behavior should largely be the same.

2. There are several data imperfections in the limit order book data as supplied by the EPEX:

   (a) A subset of orders in the order book for 2017 and 2018 is cleared via XBID with orders from other countries. Although we do not have access to the order books for other European countries, we can identify the orders in question. To be able to reproduce the historical clearings, we add virtual orders of the size of the cleared German orders on the respective other market side. We set the limit price to the price of the observed clearing, the quantity to the cleared quantity, and the end-validity date to the date of the clearing with the foreign order. We also use these orders when constructing the counter-factual auction outcomes.

   However, since this process of *order book completion* is clearly not perfect and we might be additionally missing orders from abroad which are not cleared in continuous trading, cross border trading has the potential to distort our results.

   (b) In the data provided by the EPEX, the end-validity time of cleared orders is overwritten with the clearing time in case an order is cleared. As a result, information is lost and some of the orders we consider would have been valid for longer than is visible in the available data.

   (c) Market orders are not explicitly identified and limit prices for these orders are not set to the maximum/minimum possible price of EUR

$\pm 9,999$/MWh, but to the last bid/ask price the order is executed against.

3. Furthermore, we ignore block-orders. Considering these orders would introduce a dependence between auctions for different products and complicate our experiment. While this is very much possible in an auction based framework, which is much less sensitive to increased computation times resulting from the ensuing complications than continuous trading, we avoid coupling auctions to keep the discussion simple.

4.5 Numerical Results

In this section, we discuss the comparison of the outcomes for the continuous market with the counter-factual auctions. We consider quarter-hourly and hourly products with auctions being cleared every 15min and 60min and in one single auction, which clears all the observed orders for a product at once. The latter results are thought of as an extreme opposite of continuous trading and serve as a reference.

In our experiments, we consider the full German order book of 2017 and 2018 containing roughly 71 mio orders for the hourly products and 157 mio orders for the quarter-hourly products. All computations, except visualizations, including the clearing algorithms for both market regimes and the computation of liquidity costs are implemented in JAVA.

4.5.1 Prices

In this section, we compare the distribution of prices generated by the two clearing methods. To that end, we define the volume-weighted price of a product in continuous trading as

$$\bar{p}^C = \frac{\sum_{t \in \mathcal{T}} p_t^C q_t^C}{\sum_{t \in \mathcal{T}} q_t^C}, \tag{12}$$

where $\mathcal{T}$ are the time points where quantity $q_t^C > 0$ of the product is cleared. Analogously, we define the volume-weighted price of the same product in the auction based clearing

$$\bar{p}^A = \frac{\sum_k p_k^A q_k^A}{\sum_k q_k^A}, \tag{13}$$

where p_k^A q_k^A are the prices and quantities in auction k, respectively.

Statistics for the volume weighted prices across all hourly and quarter hourly products for continuous trading and the three auction formats are reported in Table 12. As can be seen, the distributions of the volume weighted prices are

	Method	Mean	Min	Max	Median	Std	1%	5%	95%	99%
Hourly	**S**	39.56	-112.00	232.80	38.60	19.26	-3.86	10.00	70.60	88.80
	H	39.50	-106.40	206.34	38.54	19.34	-4.43	9.90	70.55	89.15
	QH	39.45	-103.57	203.90	38.49	19.33	-4.74	9.69	70.59	88.83
	C	39.45	-116.68	204.16	38.48	19.43	-5.30	9.58	70.65	89.30
QHourly	**S**	39.37	-200.60	278.00	39.00	21.49	-14.92	6.10	72.30	92.92
	H	39.29	-190.19	352.16	38.91	21.60	-14.69	5.99	72.31	93.21
	QH	39.24	-184.80	355.63	38.88	21.68	-15.08	6.00	72.45	93.66
	C	39.20	-192.94	330.84	38.81	21.84	-16.15	5.70	72.41	93.90

Table 12: Comparison of the volume weighted price across all products. S, H, and QH represent the single auction, auctioning every 60 minutes, and auctioning every 15 minutes, respectively, while the rows starting with C contain the results for continuous trading. The upper panel compares results for hourly products while the lower panel reports results for quarter-hourly products.

quite similar, which implies that no fundamental disruptions in price levels are to be expected when transitioning from continuous trading to frequent intraday auctions. The fact that not only the average volume weighted prices but also the 1%, 5%, 95%, and 99% quantiles are practically identical is quite surprising and indicates that, averaged over products, not even the most extreme prices change substantially with the introduction of frequent auctions.

We conclude that the effect of a transition to frequent auction on average price levels and distributions are negligible for both product types and across all proposed methods of clearing.

4.5.2 Traded Volume

Next we analyze the cleared volumes for both market designs. Since auctions have the reputation to produce more reliable price signals and lower liquidity cost for participants, on a first glance, it seems intuitive that the quantity cleared in an auction would exceed the quantity cleared in continuous trading [see also Deutsche Börse Group, 2018].

However, ceterus paribus, this intuition turns out to be wrong as we show below. To give an idea why this is the case, consider the following small example. Suppose for a certain product, the following orders are recorded in the limit order book:

Order Number	**Start Validity**	**Direction**	**Limit**
1	1	buy	50
2	2	sell	50
3	3	buy	100
4	4	sell	75

where all orders have an end-validity date of 4 and a size of 1 MWh.

Clearly, in continuous trading, order 1 would be cleared with order 2 and order 3 with order 4, leading to an overall traded quantity of 2 MWh. If instead the orders would be cleared in a single auction at time $t = 4$, the clearing price could be anywhere between 50 and 100 and only orders 2 and 3 would be cleared, which leads to a cleared quantity of only 1 MWh. However, note that the auction, by design, achieves the maximum welfare gain of 50, while the welfare gain from continuous trading is only 25. This example shows that even if all orders are collected in a single auction the volume that is cleared in continuous trading may be larger.

Note that the situation in the above example is not at all pathological or exceptional but rather the rule in a situation where continuous trading is replaced with a single auction at the end of the trading window. In fact, we can show the following proposition.

Proposition 1. *Given a set of limit orders in an order book posted in the time interval $[\underline{t}, \overline{t}]$ with end-validity date $\overline{t}$. If the orders are cleared according to the rules of continuous trading, the cleared quantity is always greater than the quantity cleared if the orders where bids in a uniform price double auction.*

Proof. Assume without loss of generality that the size of all orders equals the minimum order quantity. Note that any sequence in which orders arrive will always lead to a state of the order book where either all bids or all offers that are cleared in the auction are also cleared in continuous trading. If this was not the case, at least one bid and one offer which would have been cleared in the auction would remain in the book. By definition the price of the offer is smaller than or equal the auction price and the price of the bid is larger. Hence, the orders should have been cleared in continuous trading, which leads to a contradiction, establishing the claim. □

Observe that continuous trading yields the same outcome as the auction if the orders arrive in pairs sorted according to the marginal welfare gain of matching them, i.e., at $\underline{t}$ the highest bid and the lowest ask are added to the order book, then a little while later the second highest bid and the second lowest ask and so on. In summary, one could say that the fact that the auction leads to a welfare optimal clearing in many cases forces outcomes with less cleared volume.

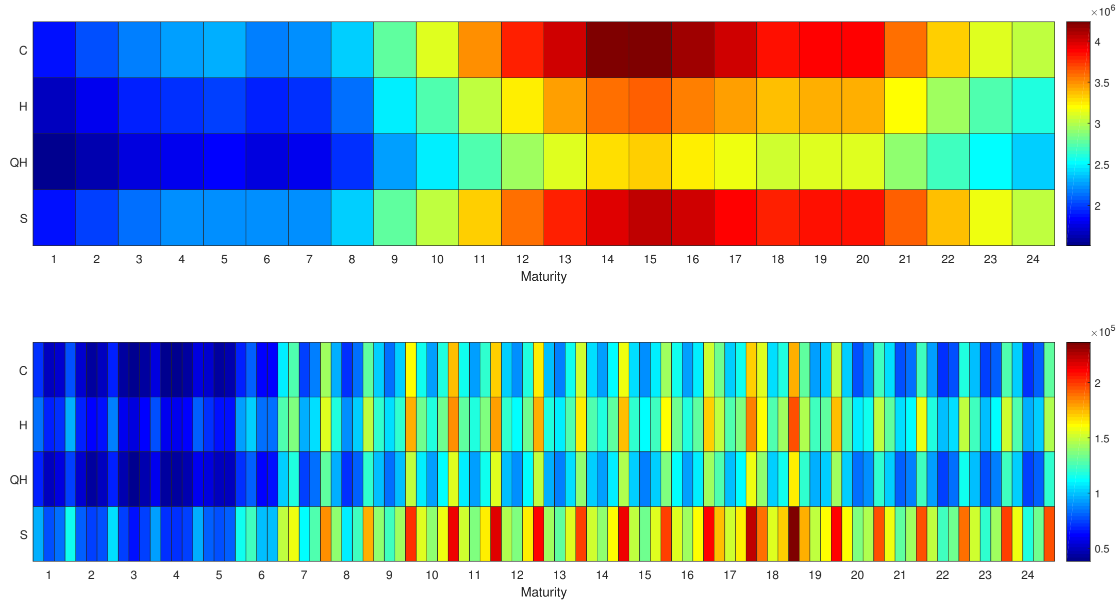

Figure 20: Comparison of total clearing volumes. The upper heatmap compares results for hourly products, while the lower panel shows the results for the quarter-hourly products.

Figure 20 displays heatmaps showing the aggregate traded volumes for all hourly and quarter-hourly products throughout the investigated period and for all market designs. As expected, the figure shows a clear pattern with less trading in the night and morning hours and a volume peak in the afternoon. While Proposition 1 does not do justice to the complexity of repeated auctions, the described effect remains valid and dominates in the numerical examples for hourly products: continuous trading clears slightly more volume than even the single auction, especially in the early afternoon, and significantly more volume than the more frequent auctions. This is all the more the case, since the orders we use to compute counter-factual auction outcomes are generated for continuous trading and consequently the limit prices for many orders are chosen such that they are adapted to the price level at the time of submission and therefore are more or less immediately executed in continuous trading. An analysis for the number of cleared orders instead of the cleared volumes yields similar results.

For quarter-hourly products the situation is less clear and the effect discussed in Proposition 1 is offset by the fact that we extend the end-validity date of the orders to the end of the auction period and thereby artificially increase the number of active orders relative to continuous trading. Also note that for quarter-hourly products the liquidity is generally lower, which leads to a rather wide bid-ask spread and volatile best-bid and best-ask prices, which in turn results in many order that are cleared in auctions but not in continuous trading.

Clearly, the effect of the extended end-validity date plays a larger role for the single and the hourly auctions than for the quarter-hourly auctions as is also visible from the results in Figure 20, which demonstrate that the two former auctions clear more volume than continuous trading while the latter clears less. However, since the described effect is an artefact of how we compute counter-factual auction results, we conclude that by and large volumes in continuous trading tend to be higher than in auctions.

Hence, if the goal of a market design is to generate trading volume, then continuous trading can be considered superior to auctions. However, traded volume is usually not viewed as an end in itself but rather a means to decrease liquidity costs and produce more reliable price discovery.

We will show in the following sections that a transition from continuous trading to frequent auctions would likely positively impact these aspects, even if less volume is cleared. For this reason and since auction trading is easier to handle, especially for smaller market participants who might be overburdened by the complications of a continuous market, we expect that participation and therefore traded volumes would increase relative to our benchmark if an auction based trading would be established.

4.5.3 Liquidity Costs

Liquidity is typically a vague term as, e.g., argued in Schwartz et al. [2020]. In order to compare liquidity costs between auctions and continuous trading, we use an approach proposed in Kuppelwieser and Wozabal [2021] who employ a cost of a round-trip (CRT) measure as a surrogate for liquidity cost. The CRT is calculated as the hypothetical cost of buying and immediately selling a certain quantity of electricity and can be calculated for auctions as well as for continuous trading.

In particular, for the continuous market, the buy and sell sides of the LOB are sorted at each point in time t by price to obtain $\dots < P^t_{-2} < P^t_{-1} < P^t_{-0} < P^t_0 < P^t_1 < P^t_2 < \cdots$, where P^t_{-0} is the highest bid-price and P^t_0 is the lowest ask-price. The corresponding order quantities are denoted by Q^t_i. For a fixed $V \in \mathbb{R}^+$, we define how much of order i would be cleared when placing a market order of size V MWh by

$$\bar{Q}^t_i(V) = \min\left(\max\left(V - \sum_{k=0}^{i-1} Q^t_k, 0\right), Q^t_i\right), \ \bar{Q}^t_{-i}(V) = \min\left(\max\left(V - \sum_{k=-i+1}^{-0} Q^t_k, 0\right), Q^t_{-i}\right).$$

This allows us to define the CRT measure at time t for a volume V as

$$CRT_t(V) = \underbrace{\frac{1}{V}\sum_k P_k^t \bar{Q}_k^t(V)}_{\text{average cost}} - \underbrace{\frac{1}{V}\sum_k P_{-k}^t \bar{Q}_{-k}^t(V)}_{\text{average revenue}}. \tag{14}$$

In order to deal with spikes in the continuous CRT, we transform CRT_t into a discrete measure as proposed in Kuppelwieser and Wozabal [2021]. In particular, we define CRT_τ at a time point τ by averaging over 15 minutes of CRT_t before τ, i.e.,

$$CRT_\tau(V) = \frac{1}{15}\int_{\tau-15}^{\tau} CRT_t(V)\, dt = \frac{1}{15}\sum_{i=2}^{N} \frac{CRT_{t_i}(V) + CRT_{t_{i-1}}(V)}{2}(t_i - t_{i-1}),$$

where $t_1, \dots, t_N$ are the N points in time where the LOB changes in the 15-minute time interval $[\tau - 15\text{min}, \tau]$. The value $CRT_\tau(V)$ thus captures the *average* CRT a trader has to pay for selling or buying a fixed quantity V during the time interval $[\tau - 15\text{min}, \tau]$.

Trading in the German continuous intraday market mostly occurs a few hours before physical delivery. Therefore CRTs of early time periods with little trading are rather high. In order to avoid an upwards distortion of computed CRTs for continuous trading we therefore consider the *trading volume weighted CRT* as defined in Kuppelwieser and Wozabal [2021]

$$CRT_{V,h}^C = \sum_\tau \frac{CRT_{V,h,\tau} Q_{h,\tau}}{\sum_\tau Q_{h,\tau}},$$

where $Q_{h,\tau}$ is the traded volume for product h in the time period $[\tau - 15\text{min}, \tau]$ and index τ runs over a grid of time points with 15 minute spacing. By taking a weighted average and using traded volumes as weights, we avoid the statistic to be dominated by outliers.

In order to calculate the CRT for auctions, we add a bid of size V as a market order to the orders of a specific auction k trading a product h and then clear the market obtaining a price $P_{V,h,k}^B$. We then repeat this process using an offer of size V, which yields a price $P_{V,h,k}^S$. Finally, we calculate the cost-of-round-trip measure for a fixed value V for the auction market k as

$$CRT_{V,h,k}^A(V) = P_{V,h,k}^B - P_{V,h,k}^S.$$

Similar to the CRT for continuous trading, we use cleared volumes as weights in order to calculate the cost of roundtrip $CRT_{V,h}^A$ for the auctions by taking a volume weighted average over all CRTs in auctions that trade product h.

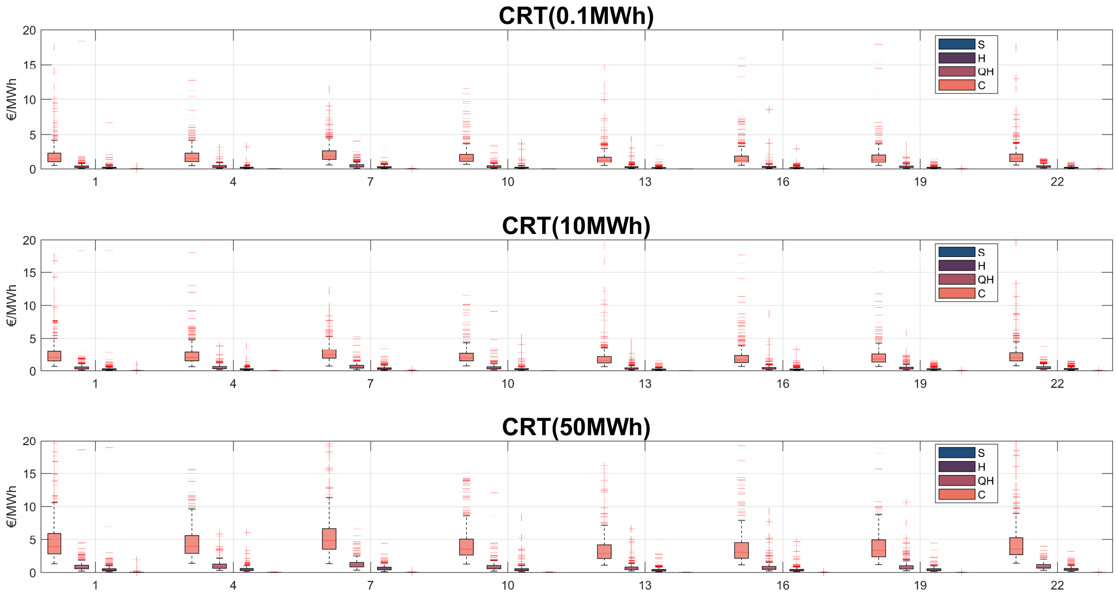

Figure 21: CRT-costs for selected hourly products for all market designs. The labels on the x-axis signify the number of the hourly product, i.e., 1 corresponds to delivery between 0:00 and 1:00, 4 to delivery between 3:00 and 4:00 and so on.

The results of the comparison between auctions and continuous clearing are shown in Figure 21 for hourly products and in Figure 22 for quarter hourly products for selected products (the results for the rest of the products are structurally identical and are therefore not displayed to keep the presentation tractable). The results demonstrate that CRT costs of continuous trading are significantly higher than that of auctions for hourly products. The analysis shows that especially for $\Delta = 60$, liquidity costs decrease dramatically to levels very close to the benchmark of a single auction. This is true uniformly for all the volumes and analyzed hours. The liquidity costs for 50MWh in continuous trading on average amount to 5.33€/MWh, while the same figure for auction trading is only 0.97€/MWh for quarter hourly auctions and 0.5€/MWh for hourly auctions. For quarter-hourly products the effect is even more pronounced with average liquidity costs of 41.23€/MWh for 15MWh in continuous trading versus 19.84€/MWh for quarter hourly auctions and 5.26€/MWh for hourly auctions.

The lower liquidity cost for auctions based trading would likely increase market participation and traded volumes and therefore further decrease liquidity costs. Hence, lower liquidity cost can be considered one of the main advantages of auction based trading. Especially for products which are less traded such as quarter hourly products.

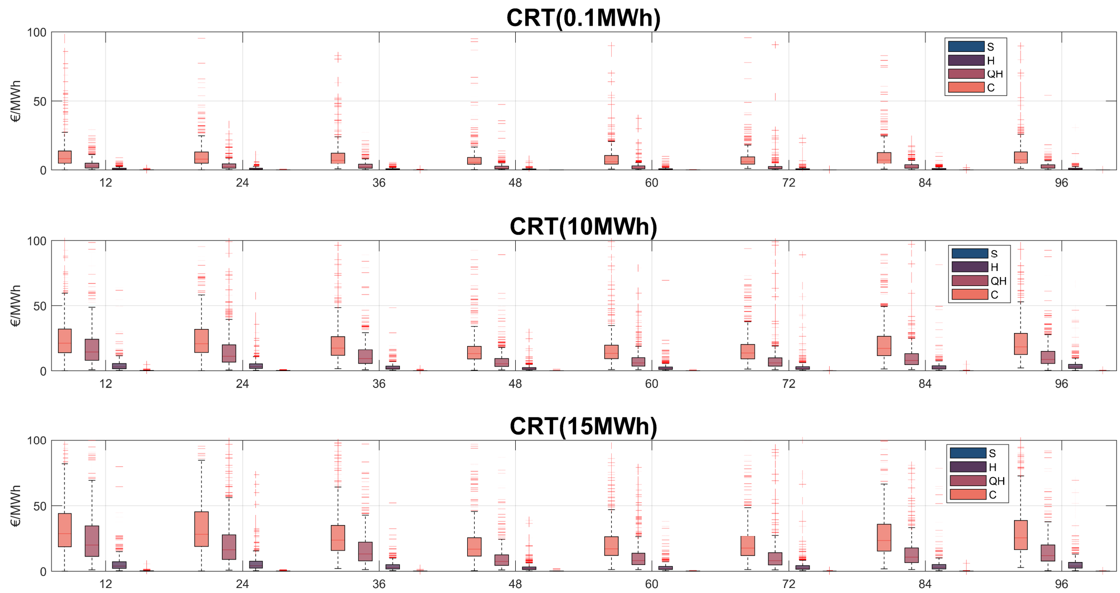

Figure 22: CRT-costs for selected quarter-hourly products for all market designs. The numbers on the x-axis signify the number of the quarter hourly product, with product one delivering between 0:00 and 0:15, product 4 between 0:45 and 1:00 and so on.

4.5.4 Noise versus Signal

The intra-product price path, that is, the price process for a single product between gate-opening and gate-closure during each trading session, can be volatile. Price changes during this period either occur due to arrival of new information, changing the fundamental value of the traded product, or due to noise induced by the trading process itself. While the former is a desirable feature of information processing by the market, the latter distorts this information. Hence, any good market design has to navigate the trade-off between these effects with the aim to maximize the information contained in prices while keeping noise as low as possible.

In order to study this issue for intraday markets prices, we propose to disentangle the price curve into signal and noise. While the constructed signal aims to capture the real underlying values of the product, the noise is defined as the transient deviation of the prices from these values.

To this end, we employ a locally linear regression framework [see, e.g., Hollander et al., 2014] with a normal kernel to decompose intra-product price paths into signal and noise. The principal idea of this framework is to perform linear regression locally at x weighting observations less, the further they are away from x. We focus on the price processes between gate closure and five hours before gate-closure because trading activity in more distant hours from gate closure is typically very low. We round the transaction times of the continuous market as well as the counter-factual market-clearing prices onto a minute by minute grid with $N = 300$ ticks for the considered five hour period. If there is more than one transaction in a given minute before gate-closure, we compute the quantity weighted average price and if there is no transaction, we fill the missing values with data from the previous minute.

The derived price path serves as input to the local linear regressions that we perform for each product. More precisely, for a vector of prices $y = [y_1, \ldots, y_{300}]$ and a vector of standardized time steps $x = [1, \ldots, 300]$, we compute the estimator

$$\begin{bmatrix} \hat{\beta}_0(x) \\ \hat{\beta}(x) \end{bmatrix} = \underset{\beta_0, \beta}{\operatorname{argmin}} \sum_{i=1}^{N} \mathcal{K}_h(x - x_i)(y_i - \beta_0 - (x - x_i)'\beta)^2, \tag{15}$$

where $\mathcal{K}_h(x) := e^{-x^2/2h}$ is the Gaussian kernel with bandwidth parameter h. The kernel is used as a weighting function which controls the amount of smoothing with larger values of h resulting in smoother functions. For our experiments we use $h = 5$ and identify the prediction from the regression (15) as the signal and the residuals as the noise.

The approach is visualized in Figure 23. At the top panel, we display bid and ask values, executed transactions, as well as counter-factual auction market-clearing prices for the product with delivery between 11:00 and 12:00 on the 2018-10-11. The solid red line and solid blue line at the panel in the center show the prices for continuous trading and frequent auctions. At the panel in the middle in Figure 23, the dashed red line and dashed blue line represent the extracted signal. At the bottom panel, the noise defined as the difference between the solid and the dashed line is depicted.

In order to compare the signal-noise trade-off of frequent auctions with continuous trading, we compute the sum of the absolute deviations between price and signal on the grid for all products. The fan charts in Figure 24 summarize the results comparing different market designs for the hourly market (left column) and the quarter-hourly market (right column). We find that for both product categories, the noise is the largest for continuous trading (60-min product average noise: 86.24, 15-min product average noise: 142.22), followed by auction-clearing every 15 minutes (60-min product average noise: 40.76, 15-min product average noise: 78.11) and auction-clearing every hour (60-min product average noise: 16.07, 15-min product average noise: 34.48).

The qualitative interpretation of the results, i.e., that signal quality is lower for continuous trading and that continuous trading thus leads to noisier price paths compared to auctions, remains unchanged for bandwidths of $h = 10$ or $h = 15$.

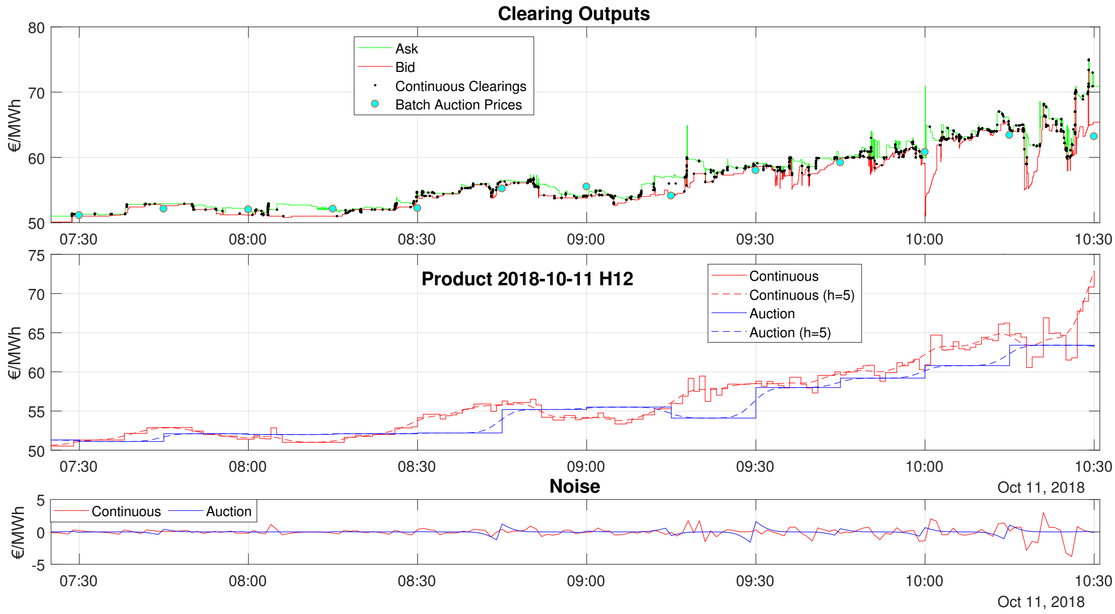

Figure 23: Decomposing Price Paths into signal and noise for the hourly product from 11:00 to 12:00 on the 2018-10-11.

The above results are also robust to replacing the locally weighted *linear* regression estimator by the classic *locally constant* Nadaraya-Watson kernel regression estimator.

Interpreting the results in the classical framework by Garbade and Silber [1979], we conclude that the chosen auction design decreases the noise from continuous trading while at the same time avoiding high price shocks induced by overly long auction intervals.

4.6 Conclusion & Policy Implications

In this paper, we investigate the use of *frequent* auctions for intraday electricity markets. The proposed design can be seen as a compromise between continuous trading, the de-facto standard in European intraday markets, and *infrequent* auctions as they have been used for some time in countries such as Spain, Italy, or Greece. Ideally, frequent auctions decrease liquidity costs and improve the quality of the price signal relative to continuous trading, while improving the ability of electricity traders to react quickly to new information over infrequent auctions.

We construct counter-factual auction outcomes from historical limit orders submitted to the German continuous market to investigate how prices, traded quantities, liquidity costs, and the signal to noise ratio of the price process would change

Figure 24: Fan chart of the noise with extreme percentiles of 1% and 99%.

if the market design was to transition to frequent auctions. We summarize the results of the resulting numerical comparison Figure 25.

While average and volume weighted prices do not significantly differ between the two market designs, the same does not hold true for traded quantities for which the outcomes are ambiguous with uniformly lower volumes in auctions for hourly products and higher volumes for most auctions of quarter-hourly products. However, the latter effect is an artefact of how we construct the counter-factual auction outcomes so that we conclude that overall auctions lead to less traded volume, which is also supported by theoretical considerations in Proposition 1.

Despite the decreased volumes, the cost-of-roundtrip markedly decreases for all considered auctions, i.e., liquidity costs are higher for continuous trading. Finally, looking at the measurement of noise in the price signal, we observe that auction based trading produces more robust and reliable price signals that are less affected by transient short-term shocks that are not connected to fundamental changes of the value of the traded product.

Overall and in accordance with Deutsche Börse Group [2018] and Ehrenmann et al. [2019], Ocker and Jaenisch [2020], we conclude that the proposed auction design to a large extend captures the advantages of both continuous trading and auctions. In particular, the fact the auctions are frequent and the last auction takes place close to delivery mitigates the disadvantages of previously existing auction systems which had only few auctions with the last one closing several hours before

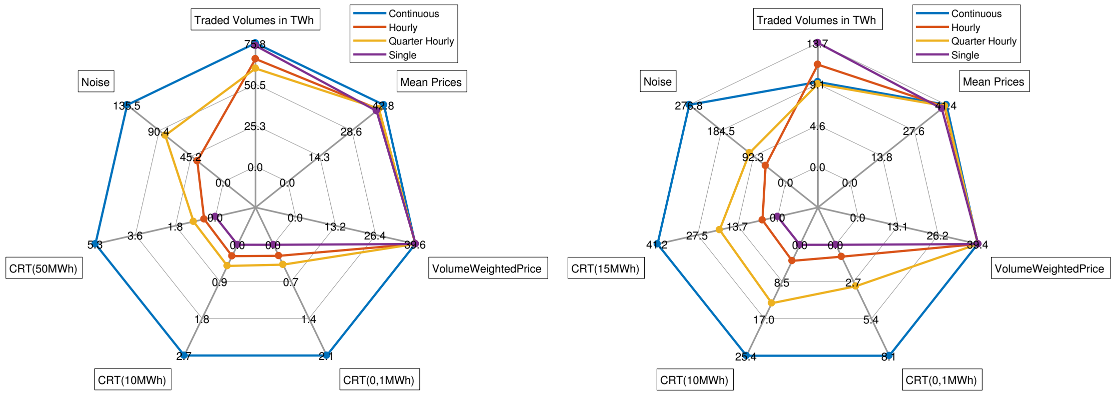

Figure 25: Overview of the results for prices, traded quantities, noise, and liquidity costs for both hourly and quarter-hourly products.

physical delivery.

Furthermore, auctions give the clearing authority the possibility to account for the realities of the restrictions imposed by the physical properties of electricity grids: An auction based intraday market design would allow for more time to clear the market, making it possible to price cross-border capacities implicitly and making it easier to integrate block orders. Overall, if the market design is chosen carefully, we thus believe that prices would more accurately reflect true scarcity rents and would thus be more closely connected to prices on the balancing market.

Lastly, on an administrative level, auctions are simpler allowing more participants to take part in trading. This is especially true, since day-ahead markets, which most electricity firms are familiar with, are also organized as auctions. Furthermore, since orders are batched over longer time periods, there are no speed advantages for advanced traders and therefore no possibility for latency arbitrage in an auction based market.

5 Conclusion

In this dissertation, I used detailed order-book data to analyse liquidity costs of European intraday markets and found a profitable passive trading strategy using updated weather forecasts. Finally, a new market design as a combination of continuous and auction markets is proposed.

5.1 Results and Contribution to the Field

A new liquidity measure that captures multiple dimensions of the continuous and auction intraday markets was defined. This method allows to estimate the cost per MWh of a fixed volume for a round-trip, i.e. buying and selling the fixed amount virtually at the same time. Previous work used the bid-ask-spread Balardy [2022] or trading volumes Hagemann and Weber [2013] to measure the liquidity of the German intraday market. In this dissertation, I introduced the cost of round-trip (CRT) measure, which comes with the advantage of quantifying the cost of the bid-ask-spread and the cost for trading larger volumes both at the same time. The calculation of this measure is only possible if the whole order-book is available.

The analysis of the cost per round-trip shows a different liquidity distribution along trading time when comparing the Italian intraday auction with the German continuous intraday market. Auction markets are more liquid than continuous trading shortly after the day-ahead marked closes, and liquidity of the continuous trading is higher shortly before physical delivery. However, the big advantage of the continuous market is the shorter lead-time, allowing market participants to trade forecast-updates of renewables until five minutes before physical delivery. This allows for integrating the latest production forecasts into the power schedule, which are more precise including real-time measurements. Intraday auction markets close a few hours before physical delivery and do not allow to trade these latest forecast-updates. Although, the big advantage of the auction market is the pooling of orders, which decreases liquidity costs.

As a second step, a trading strategy for the German continuous intraday market was developed. The key of the strategy was to use forecast errors in combination with a passive trading strategy that avoids the enhanced bid-ask-spread at position opening and closing. This strategy is based on the information of updated-forecasts for photovoltaic and wind production, which is available at the opening of the market. However, liquidity of the continuous intraday market is relatively low at that time and a market order would lead to large trading costs due to the large bid-ask-spread. As a solution, I placed a limit order on the top of the order-stack as passive order and updated the price to remain the best available order until an active trader picks the order. A limit order is also used to close the position before the market closes. If the placed order was not executed, a market order is placed shortly before the market closes to avoid balancing costs. Possible balancing costs are also considered in the analysis.

This trading strategy is implemented by setting the following parameters: (i) maximum volume to open risking to suffer from liquidity costs, (ii) decision to use the first available forecast-update or to wait for a better one risking that other traders take advantage of the new information, and (iii) timing for opening and closing the position, and the time to close the market order risking to close at unfavorable prices. A grid-search method was used to find a trading strategy on an insample set, which was then evaluated on an outsample set. The developed strategies are profitable, with larger profits when replacing the latest intraday forecast with the exact production. Hence, accurate and early available production forecasts would allow for larger profits, leading to an arms race in weather forecasting.

Finally, a new design for the European intraday market was proposed. This design combines the advantages of both prevailing market designs and consists of sequential batch auctions with intervals of 15 or 60 minutes. As a consequence, this allows to trade until 30 minutes before physical delivery. The design would increase liquidity by pooling all orders between two clearings and allowing to integrate renewables with a short lead-time. The auction clearing of the new market design was implemented using submitted orders of the German continuous intraday market. A selection of the realized outcomes of the continuous market was used to analyze the counter-factual outcomes of the batch auctions. The new market design shows higher liquidity and less extreme prices, but lower trading volumes compared to the operative continuous intraday market.

5.2 Further Research Prospects

European institutions decided to use the continuous market design for the intraday market, which offers the possibility to add auction markets. Historically, the intraday markets in Spain, Portugal and Italy were exclusively managed via auctions. In recent times, these countries have introduced a continuous market, which currently exists alongside with a reduced amount of auction markets.

Exploring the impact of introducing a continuous market in these three countries could be an interesting question for further research. In these countries, trading on auction markets is the common practice. Therefore, it might be interesting to investigate whether traders continue to trade on the remaining auction markets or whether they accept the introduced continuous market.

An analysis of the impact on liquidity assuming the hybrid market scenario in Spain, Portugal and Italy could also be a valuable idea for further research. To this end, the impact of the remaining auction markets on continuous market liquidity could be investigated. Also, the behaviour of liquidity before and after auctions could be an interesting research question to address.

The European power market consists of sequential power markets and they are all interconnected. Exploring the potential of novel trading strategies operating in the intraday, day-ahead, and balancing markets could therefore be a possible avenue for future research. Also, future investigation of trades between the four quarter-hourly products and the corresponding hourly product would also be interesting.

The limit order books of the years 2015–2018 were available for this dissertation and one can observe an enhanced development of the market. The lead-time decreased in this period from 45 to 30 minutes (5 minutes within the same control area) before physical delivery. Also, 30 minutes products were introduced. The market is strongly growing and liquidity shows a continuously increasing trend.

References

Matteo Aquilina, Eric Budish, and Peter O'Neill. Quantifying the High-Frequency Trading "Arms Race". *The Quarterly Journal of Economics*, forthcoming, 2021.

ARERA. Memoria. https://www.arera.it/allegati/docs/19/229-19.pdf, 2019. [Online; accessed 10-December-2021].

C. Balardy. An empirical analysis of the bid-ask spread in the continuous intraday trading of the german power market. *The Energy Journal*, 43(3), 2022.

R. Baule and M. Naumann. Volatility and dispersion of hourly electricity contracts on the german continuous intraday market. *Energies*, 14(22):7531, 2021.

Julia Bellenbaum, Michael Bucksteeg, Thomas Kallabis, Christian Pape, and Christoph Weber. Intra-day cross-zonal capacity pricing. *Study on behalf of Ofgem*, 2014.

Gilles Bertrand and Anthony 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.

M.J. Brennan and H.H. Cao. Information, trade, and derivative securities. *The Review of Financial Studies*, 9(1):163–208, 1996.

Eric Budish, Peter Cramton, and John Shim. The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. *The Quarterly Journal of Economics*, 130(4):1547–1621, 2015.

Bundesnetzagentur. Beschluss bk6-12-024. https://www.regelleistung.net/ext/static/rebap, 2012. [Online; accessed 21-Febrary-2021].

Bruno Burger. Net Public Electricity Generation in Germany in 2018. Report, Fraunhofer Institute for Solar Energy Systems, ISE, January 2019. URL https://www.ise.fraunhofer.de/content/dam/ise/en/documents/News/Stromerzeugung_2018_2_en.pdf.

Tabors Caramanis and Associates Inc. UPPO Auction Module User Manual. https://www.mercatoelettrico.org/En/MenuBiblioteca/Documenti/20100429MarketSplitting.pdf, 2002. [Online; accessed 09-September-2021].

P. Chelley-Steeley. Price synchronicity: The closing call auction and the london stock market. *Journal of International Financial Markets, Institutions and Money*, 19(5):777–791, 2009.

P.L. Chelley-Steeley. Market quality changes in the london stock market. *Journal of Banking & Finance*, 32(10):2248–2253, 2008.

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018.

C. Comerton-Forde, S.T. Lau, and T. McInish. Opening and closing behavior following the introduction of call auctions in singapore. *Pacific-Basin Finance Journal*, 15(1):18–35, 2007.

NEMO Committee. Euphemia public description. https://www.mercatoelettrico.org/it/MenuBiblioteca/Documenti/20181212EuphemiaPublicDescription.pdf, 2019. [Online; accessed 25-November-2021].

P. Cramton. Electricity market design. *Oxford Review of Economic Policy*, 33(4), 2017.

Deutsche Börse Group. Continuous auction market model: A proposal for the future european intraday power market. Technical report, Deutsche Börse AG, 10 2018.

S. Du and H. Zhu. What is the optimal trading frequency in financial markets? *The Review of Economic Studies*, 84(4):1606–1651, 2017.

ECA. Final report. https://extranet.acer.europa.eu/en/Electricity/Market%20monitoring/Documents_Public/ECA%20Report%20on%20European%20Electricity%20Forward%20Markets.pdf, 2015. [Online; accessed 25-November-2021].

Nicholas Economides and Robert A Schwartz. Equity trading practices and market structure: Assessing asset managers' demand for immediacy. In *The Electronic Call Auction: Market Mechanism and Trading*, pages 169–205. Springer, 2001.

Andreas Ehrenmann, Pierre Henneaux, Gerd Küpper, Johan Bruce, and Linda Schumacher Björn Klasman. The future electricity intraday market design. Technical report, Study by TRACTEBEL Advisory and Advanced Analytics and SWECO for the European Commission, 2019. URL https://op.europa.eu/en/publication-detail/-/publication/f85fbc70-4f81-11e9-a8ed-01aa75ed71a1/language-en. [Online; accessed 17-October-2021].

E. Ela, M. Milligan, A. Bloom, A. Botterud, A. Townsend, and T. Levin. Evolution of wholesale electricity market design with increasing levels of renewable generation. Technical report, National Renewable Energy Lab. (NREL), 2014.

ENTSO-E. STATISTICAL FACTSHEET 2018. Report, ENTSO-E, June 2019. URL https://eepublicdownloads.entsoe.eu/clean-documents/Publications/Statistics/Factsheet/entsoe_sfs2018_web.pdf.

EPEX. Epex spot operational rules. https://www.epexspot.com/en/download#rules-fees-processes, 2019. [Online; accessed 16-December-2019].

EPEX. Epex spot annual report 2019. https://www.epexspot.com/sites/default/files/download_center_files/Epex-spot-2019_200703_Planche.pdf, 2020. [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].

Shu Fan and Rob J Hyndman. Forecasting electricity demand in australian national electricity market. In *2012 IEEE Power and Energy Society General Meeting*, pages 1–4. IEEE, 2012.

K.D. Garbade and W.L. Silber. Structural organization of secondary markets: Clearing frequency, dealer activity and liquidity risk. *The Journal of Finance*, 34(3):577–593, 1979.

Ernesto Garnier and Reinhard Madlener. Balancing forecast errors in continuous-trade intraday markets. *Energy Systems*, 6:361–388, 2015.

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.

GME. Documento di consultazione 05/2016. https://www.mercatoelettrico.org/it/MenuBiblioteca/Documenti/20160606_DCO_Nuovi_MI.pdf, 2016. [Online; accessed 07-December-2019].

GME. GME Bilancio d'Esercizio 2018. Report, Gestore dei Mercati Energetici S.p.A., 2019. URL https://www.mercatoelettrico.org/En/MenuBiblioteca/Documenti/20190618BilancioGME2018.pdf.

Shadi Goodarzi, H. Niles Perera, and Derek Bunn. The impact of renewable energy forecast errors on imbalance volumes and electricity spot prices. *Energy Policy*, 134:110827, 2019.

C. Graf and D. Wozabal. Measuring competitiveness of the EPEX spot market for electricity. *Energy Policy*, 62:948—-958, 2013.

Christoph Graf and Frank A. Wolak. Measuring the Ability to Exercise Unilateral Market Power in Locational-Pricing Markets: An Application to the Italian Electricity Market. Working paper, 2020.

Christoph Graf, Federico Quaglia, and Frank A. Wolak. Simplified electricity market models with significant intermittent renewable capacity: Evidence from italy. NBER Working Papers 27262, National Bureau of Economic Research, 2020. URL http://www.nber.org/papers/w27262.

S. Hagemann and C. Weber. Trading volumes in intraday markets: Theoretical reference model and empirical observations in selected european markets. EWL Working Paper 03/15, University of Duisburg-Essen, Chair for Management Science and Energy Economics, 2015.

Simon Hagemann. Price determinants in the german intraday market for electricity: An empirical analysis. EWL Working Paper 18/13, Essen, 2013. URL http://hdl.handle.net/10419/103294.

Simon Hagemann and Christoph Weber. An empirical analysis of liquidity and its determinants in the german intraday market for electricity. *SSRN Electronic Journal*, 01 2013.

Jari Miettinen & Samuli Sillanpää Hannele Holttinen. Wind power forecasting accuracy and uncertainty in finland. Discussion papers of diw berlin, VTT Technology 95, 2013.

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.

S. Hauser, A. Levy, and U. Yaari. Trading frequency and the efficiency of price discovery in a non-dealer market. *The European Journal of Finance*, 7(3):187–197, 2001.

Lion Hirth and Jonathan Mühlenpfordt. Handel auf Basis des Regelleistungs-Abrufs. Technical report, Neon – Neue Energieökonomik, 2021. URL https://neon.energy/Neon_2021_Intraday_Regelleistung.pdf.

Lion Hirth and Ingmar Schlecht. Redispatch Markets in Zonal Electricity Markets: Inc-Dec Gaming as a Consequence of Inconsistent Power Market Design (not Market Power). EconStor Preprints 194292, ZBW – Leibniz Information Centre for Economics, 2019.

Myles Hollander, Douglas A. Wolfe, and Eric Chicken. *Nonparametric statistical methods*. 2014.

S.-Y. Hu. A simple estimate of noise and its determinant in a call auction market. *International Review of Financial Analysis*, 15(4-5):348–362, 2006.

S.-Y. Hu and C. Chan. Trading frequency and noise. *Applied Financial Economics Letters*, 1(4):243–247, 2005.

Paul Irvine, George J. Benston, and Eugene Kandel. Liquidity beyond the inside spread: Measuring and using information in the limit order book. *SSRN Electronic Journal*, 05 2000.

Tim Janke and Florian Steinke. Forecasting the price distribution of continuous intraday electricity trading. *Energies*, 12(22), 2019.

Roberta Jannuzzi. Storia di enel dal 1962 ai nostri giorni. https://web.archive.org/web/20150223164737/http://www.radioradicale.it/storia-di-enel-dal-1962-ai-nostri-giorni, 2021. [Online; accessed 10-December-2021].

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.

Christopher Kath and Florian Ziel. Optimal order execution in intraday markets: Minimizing costs in trade trajectories, 2020.

Rüdiger Kiesel and Florentina Paraschiv. Econometric analysis of 15-minute intraday electricity prices. *Energy Economics*, 64:77 – 90, 2017. ISSN 0140-9883.

Daniel S. Kirschen. Risk Exposure, Externalization and Allocation in Unbundled Power Systems. *IEEE Transactions on Power Systems*, 36(3):1879–1889, 2021.

R. Kovacevic and D. Wozabal. A semiparametric model for electricity spot prices. *IIE Transactions*, 46:344–356, 2014.

Jan Kregel. The call market: Historical artifact or market architecture of the future. In *The Electronic Call Auction: Market Mechanism and Trading*, pages 27–42. Springer, 2001.

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.

Sergei Kulakov and Florian Ziel. The impact of renewable energy forecasts on intraday electricity prices. *Economics of Energy & Environmental Policy*, 10 (1), 2020.

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. *Energy Policy*, 154, 2021.

Albert S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6): pp. 1315–1335, 1985. ISSN 00129682. URL http://www.jstor.org/stable/1913210.

B. Lauterbach. A note on trading mechanism and securities' value: The analysis of rejects from continuous trade. *Journal of Banking & Finance*, 25(2):419–430, 2001.

Hannes Leeb and Benedikt M. Pötscher. Model selection and inference: Facts and fiction. *Econometric Theory*, 21(1):21–59, 2005.

S.C. Littlechild. Competition and regulation in the british electricity industry. *Utilities Policy*, 2(4):270–275, 1992. ISSN 0957-1787. doi: https://doi.org/10.1016/0957-1787(92)90002-Z. URL https://www.sciencedirect.com/science/article/pii/095717879290002Z. Special Issue Clearing the utility markets.

N. Löhndorf and D. Wozabal. The Value of Coordination in Multimarket Bidding of Grid Energy Storage. *Operations Research*, to appear, 2022.

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.

Erin T. Mansur and Matthew W. White. Market Organization and Efficiency in Electricity Markets. Technical report, 2012.

G. Marcjasz, B. Uniejewski, and R. Weron. Beating the naïve—combining lasso with naïve intraday electricity price forecasts. *Energies*, 13:1667, 04 2020.

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.

C.J. Muscarella and M.S. Piwowar. Market microstructure and securities values:: Evidence from the paris bourse. *Journal of Financial Markets*, 4(3):209–229, 2001.

M. Narajewski and F. Ziel. Econometric modelling and forecasting of intraday electricity prices. *Journal of Commodity Markets*, 19:100107, 2020a.

Michał Narajewski and Florian Ziel. Ensemble forecasting for intraday electricity prices: Simulating trajectories. *Applied Energy*, 279:115801, Dec 2020b. ISSN 0306-2619.

Karsten Neuhoff, Nolan Ritter, Aymen Salah-Abou-El-Enien, and Philippe Vassilopoulos. Intraday markets for power: Discretizing the continuous trading? Discussion Papers of DIW Berlin 1544, DIW Berlin, German Institute for Economic Research, 2016.

Fabian Ocker and Vincent Jaenisch. The way towards european electricity intraday auctions – status quo and future developments. *Energy Policy*, 145:111731, 2020.

I. Oksuz and U. Ugurlu. Neural network based model comparison for intraday electricity price forecasting. *Energies*, 12(23), 2019.

M.S. Pagano and R.A. Schwartz. A closing call's impact on market quality at euronext paris. *Journal of Financial Economics*, 68(3):439–484, 2003.

Christian Pape, Simon Hagemann, and Christoph Weber. Are fundamentals enough? explaining price variations in the german day-ahead and intraday power market. *Energy Economics*, 54, 01 2016.

Michael 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.

R.A. Schwartz. *The Electronic Call Auction: Market Mechanism and Trading: Building a Better Stock Market*. The New York University Salomon Center Series on Financial Markets and Institutions. Springer US, 2012. ISBN 9781461516972.

Robert A. Schwartz, Reto Francioni, and Peter Weber. Market Liquidity: An Elusive Variable. *The Journal of Portfolio Management*, 2020.

A. Shapiro, D. Dentcheva, and A. Ruszczyśki. *Lectures on Stochastic Programming: Modeling and Theory*. MOS-SIAM series on optimization. Siam, 2009.

StataCorpLLC. Stata lasso reference manual release 16. https://www.stata.com/manuals/lasso.pdf, 2019. [Online; accessed 16-December-2019].

Benn Steil. Call market trading: history, economics, and regulation. In *The Electronic Call Auction: Market Mechanism and Trading*, pages 43–53. Springer, 2001.

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.

M. Twu and J. Wang. Call auction frequency and market quality: Evidence from the taiwan stock exchange. *Journal of Asian Economics*, 57:53–62, 2018.

Bartosz Uniejewski and Rafał Weron. Efficient forecasting of electricity spot prices with expert and lasso models. *Energies*, 11(8), 2018. ISSN 1996-1073.

J. Viehmann. State of the german short-term power market. *Zeitschrift für Energiewirtschaft*, 41(2):87–103, 2017.

N.G. von Luckner, Á Cartea, S. Jaimungal, and R. Kiesel. Optimal market maker pricing in the german intraday power market. Working paper, 2017.

Elaine Wah and Michael P. Wellman. Latency arbitrage, market fragmentation, and efficiency: A two-market model. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, EC '13, pages 855–872, New York, NY, USA, 2013. Association for Computing Machinery.

Christoph Weber. Adequate intraday market design to enable the integration of wind energy into the european power systems. *Energy Policy*, 38(7):3155–3163, 2010.

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.