Master's thesis

NTNU
Norwegian University of Science and Technology
Faculty of Economics and Management
Dept. of Industrial Economics and Technology Management

Dyrkolbotn, Sjur
Hove, Anders
Jørgensen, Carl Fredrik

Order execution in power futures markets for hedging power producers

Master's thesis in Industrial Economics and Technology Management
Supervisor: Stein-Erik Fleten
Co-supervisor: Vadim Gorski

June 2021

NTNU
Norwegian University of Science and Technology

Dyrkolbotn, Sjur
Hove, Anders
Jørgensen, Carl Fredrik

Order execution in power futures markets for hedging power producers

Master's thesis in Industrial Economics and Technology Management
Supervisor: Stein-Erik Fleten
Co-supervisor: Vadim Gorski
June 2021

Norwegian University of Science and Technology
Faculty of Economics and Management
Dept. of Industrial Economics and Technology Management

NTNU
Norwegian University of
Science and Technology

Problem description

Making hedging decisions is a classical problem faced by power producers. Power producers may hedge their cash flows to reduce price risk by selling power futures. A critical decision is then how much power should be exposed to the spot price and how much should be hedged using futures contracts.

Due to lack of liquidity in power futures markets, the power producer should not enter a significant position in futures contracts all at once as it will drain the market liquidity. Consequently, the power producers must decide how much and when to trade different power futures. In this regard, the trade-off between price impact and price risk is particularly relevant when trading in an illiquid market.

This problem is deconstructed into two sub-problems: the hedging problem and the order execution problem. The hedging problem considers how much of their production a producer should hedge and how the hedging volume should be allocated across futures contracts in the medium term. On the other hand, the order execution problem considers how trading volumes should be executed in the market, exploring how the power producer should trade power futures within daily trading periods. By drawing inspiration from the contemporary literature, a mathematical model which integrates the hedging problem and order execution problem is proposed in this thesis.

Preface

We want to express our gratitude towards those that have contributed along the way. The quality of this thesis is considerably higher, thanks to their valuable inputs. First and foremost, we are grateful to our supervisor, Professor Stein-Erik Fleten, for insightful discussions and valuable assistance throughout the semester. Further, we would like to thank Professor Nils Löhndorf and Vadim Gorski for their assistance and input during the model implementation phase. Additionally, we would like to thank them for granting us access to their general-purpose solver for large-scale stochastic optimization, QUASAR.

Trondheim, June 18th, 2021

Anders Hove

Anders Hove

Carl Fredrik Jørgensen

Carl Fredrik Jørgensen

Sjur Dyrkolbotn

Sjur Dyrkolbotn

Abstract

Due to the changing liquidity of power derivatives contracts during their lifecycle, efficient order execution strategies are highly relevant for power producers. The relevant aspect for a power producer aiming to hedge part of their production by trading power derivatives corresponds to how order execution decisions can be made to balance the trade-off between the preference for risk aversion and maximising revenues.

The existing literature emphasises the order execution problem in a vacuum without considering decision making outside of the trading period. Additionally, the existing literature predominantly considers order placement for blue-chip stocks to the author's best knowledge. This thesis fills a gap in the literature by evaluating order execution strategies in an illiquid market context.

The purpose of the thesis is to construct an order execution model that meets the requirements of a power producer. An integrated order execution model, referred to as the `Integrated Postponement model`, is proposed. The `Integrated Postponement model` introduces postponement optionality and dynamic trading volume allocation based on market liquidity, of which both aspects are novel contributions to the field. The `Integrated Postponement model` is composed of a multistage stochastic hedging model with a multistage stochastic order execution model.

The trading performance of the `Integrated Postponement model` has been evaluated using a backtesting framework that incorporates the limit order book microstructure. The `Integrated Postponement model` was seen to increase trading revenues by 89 bps compared to the best-performing benchmark. One can conclude from the results that the postponement optionality leads to a better trading performance by being more selective. Additionally, the introduction of dynamic trading volume allocation was seen to increase trading revenues by 39 bps.

Sammendrag

På grunn av den endrende likviditet til kraftderivater gjennom kontraktens livsyklus, er effektive strategier for ordreutførelse høyst relevante for strømprodusenter. Det relevante aspektet for en strømprodusent som ønsker å redusere risiko for deler av produksjonen sin ved å handle kraftderivater, handler om hvordan transaksjoner kan utføres slik at man balanserer preferansen for risikoaversjon og maksimering av inntekter.

Den eksisterende litteraturen innenfor fagfeltet fokuserer på ordreutførelse i et vakuum, uten å betrakte avgjørelser utenfor handelsperioden. Samtidig er den eksisterende litteraturen, til forfatternes viten, fokusert på handel av aksjer med høy markedsverdi. Denne avhandlingen fyller en mangel i litteraturen ved å evaluere strategier for ordreutførelse i et illikvid marked.

Formålet med denne avhandlingen er å lage en modell for ordreutførelse, som imøtekommer kravene til en strømprodusent. En integrert orderutførelsesmodell, som heretter vil omtales som the `Integrated Postponement model`, foreslås derfor. The `Integrated Postponement model` introduserer utsettelsesopsjonalitet, og dynamiske handelsvolumer basert på markedslikviditet, hvor begge disse aspektene er nye bidrag innenfor feltet. The `Integrated Postponement model` er presentert som en flerstegs stokastisk programmeringsmodel med rullende tidshorisont.

Ytelsen til the `Integrated Postponement model` har blitt tested ved hjelp av en backtest-metodologi som tar i bruk mikrostrukturen til en ordrebok. The `Integrated Postponement model` økte inntektene fra handel av strømderivater med 89 bps sammenlignet med den beste referansestrategien. En kan konkludere fra resultatene at utsettelsesopsjonaliteten øker ytelse ved å kunne være mer selektiv med tanke på handelsavgjørelser. Samtidig ble det vist at dynamiske handelsvolumer økte inntektene med 39 bps.

Contents

Figures

Tables

Chapter 1

Introduction

Power producers often have to acquire or liquidate significant positions of financial instruments. Selling or buying these prominent positions, on the other hand, is easier said than done. Large orders remove liquidity from the market, leading to price impacts due to clearing the best bid or ask orders. Price impact is already an issue in liquid markets but is magnified in illiquid markets due to a lack of trading activity. Power producers, therefore, seek order execution strategies that limit the influence on the security price.

Order placement strategies are created to limit adverse price movement by deciding upon order size, order type and trading frequency. These decisions are part of the order placement problem or the portfolio execution problem and has been researched extensively. Bertsimas and Lo (1998) and Almgren and Chriss (1999) are some of the first to implement mathematical models to optimise the acquisition or liquidation of significant positions. A commonality for most papers researching the order execution problem is the balancing between price impact and price risk. Traders are inclined to reduce their price impact, yet they must consider price risk. Securities that experience severe volatility incentivises early execution. Since illiquid securities are more volatile than liquid securities, they yield a higher risk of large price fluctuation during the day.

Electricity producers face two major decisions with regards to optimising their operations: production planning and financial hedging. The production plan is a long term production schedule, where producers seek to maximise their expected revenue. The second major decision is constructing a hedging plan, where the producer decides how to sell the power. Producers can either sell power on the spot market or the derivatives market. Risk-averse power producers may hedge their production portfolio against price shocks by trading power futures. For instance, Norway experienced negative power prices for the first time in 2020 (E24, 2020). The negative power prices was due to mild weather combined with much rainfall.

With over 90% of all power coming from hydropower production, rainfall greatly affects the price due to overflowing reservoirs in the hydro production facilities (SSB, 2020). Power producers are aware of these systematic risk factors and use meteorological forecasts to reduce these risks. However, the Corona lockdown also caused lower demand for power, pushing the prices downwards. In contrast to the weather, the lockdown situation is an event that was not possible to predict. However, Hedging your position would reduce the impact of the corona lockdown situation on revenue streams since you would already have secured cash flows before the sharp price decline. Power producers can trade derivatives on markets such as NASDAQ Nordic or the European Energy Exchange (EEX) to hedge their production portfolio. Some power producers use mathematical models to determine hedging policies. The hedging policies describe spot and derivatives allocation of production volumes. One such model is presented by Dimoski et al. (2019), where a dynamic hedging model determines the hedging plan based on expected future cash flows. Their model considers stochastic reservoir inflow, market risk and currency risk. However, to the authors' best knowledge, research on liquidity risk in the European power market is limited.

On the 10th of September, 2018, NASDAQ Clearing's default fund experienced a catastrophic loss of £100m, more than two-thirds of the fund, due to a member of NASDAQ failing to meet their margin call. This trading member was Einar Aas, the notorious Norwegian power trader known for his risky trades. Einar Aas noted that "[His] exposure to the market was too big relative to the liquidity in the market." Senior analyst John Brottemsmo at Kinect Energy pointed out that the market prices would eventually stabilise, however, the trading volumes would likely experience a long-term hit due to Aas' huge market position (Paulson & Starn, 2018). Considering this event, there is an increasing interest in research within this field. Low liquidity in the power derivatives market leads to unfavourable conditions for power producers as they struggle to trade power derivatives. Placing large orders in such markets will lead to adverse price movement. Adverse price movement is most likely not an issue for small and midsize producers, but large scale producers must consider this. Therefore, the incentive to use trading algorithms to optimise order execution is substantial as trading decisions can result in significant financial impact.

This thesis will propose an integrated hedging model, combining a dynamic hedging plan and dynamic order placement. The hedging problem is solved sequentially. First, a hedging model decides on a daily trading plan. This solution is then used as input in the order execution model. The integrated hedging model is implemented in a case study for a European power producer. The model has been backtested and compared to 6 benchmarks using a microstructural limit order book simulation.

We have three main contributions to the field of study. The first and most noteworthy contribution is the inclusion of postponement optionality in the order execution model, a novelty within the field. The contemporary literature is restricted to the liquidation of a portfolio within the predefined trading period. We expand on the contemporary models to allow for the postponement of liquidation to subsequent trading periods. The proposed order execution model referred to as the `Postponement model`, does perform better than other trading benchmarks tested in this thesis, signifying that power producers will potentially benefit by using the `Postponement model`.

Our second contribution is the introduction of dynamic trading volume allocation to the order execution problem. By using the `QUASAR Dynamic Hedging model`, we are able to construct a hedging plan that considers market liquidity. Instead of treating the daily trading volumes as exogenous, daily trading volumes are then included as a variable to decide upon in the order execution model.

Lastly, we propose a backtesting methodology that employs the limit order book microstructure for illiquid markets. Using order flow data from the EEX futures market, we have created a backtesting framework where the limit order book is updated chronologically according to the real-time order arrival for the exchange. Additionally, we have included the option of placing your own orders on the trading exchange. The idea of this approach is to get more realistic price impact estimates.

This thesis is structured in the following way: Chapter 2 presents background information relevant to the problem while 3 covers relevant literature to enlighten the reader about the context of this thesis. Following, the `Integrated Postponement model` is presented in chapter 4. Chapter 5 describes the data sets, while chapter 6 highlights the methodologies that have been used in this thesis. Next, chapter 7 presents the results from the backtest, which is followed by a discussion. Lastly, the findings are summarised in chapter 8, along with mentions of potential further work.

Chapter 2

Background

The purpose of this chapter is to provide context for the central topics of this thesis. In section 2.1, the European power market is briefly explained, followed by a more detailed exploration of power futures. Then, the central topics for the limit order book, such as microstructure and liquidity measurements, are introduced in section 2.2.

2.1 The European Power Market

As this thesis studies power producers' decision making regarding hedging, it is appropriate to introduce the power market and products used for hedging. In this section, important characteristics of the European power market are presented in brief. In section 2.1.1, the role of the most important physical power markets is explained. After that, the characteristics of volatile electricity prices are outlined in section 2.1.2, in addition to an introduction to the practice of hedging. Section 2.1.3 concludes with an exploration of power derivatives and their use cases for hedging.

2.1.1 Short-term physical power markets

Three important short-term physical power markets are the day-ahead market, the intraday market and the balancing market. Market participants can buy or sell power for the next 24 hours through a blind auction in the day-ahead market. The participants place buy or sell orders with a specified volume of electricity to deliver at different price levels. In the day-ahead auction, bidding occurs for all the hours of the consecutive day. Thus, a market participant can transact in 24 different power instruments. In addition to conventional limit orders, it is possible to place block orders. A block order involves specifying a volume and price for a set of consecutive hours within the same day. Regular block orders are all-or-nothing orders, which means that either they all have to be entirely accepted or all fully rejected. Block orders can be linked together so that the acceptance of one block order is conditional on the acceptance of another block order. The market-clearing price is found at the aggregated demand and supply curves intersection and applies as the price for all transactions. Power can be continuously bought and sold on the same day as delivery occurs in the intraday market. Thus, the intraday market enables participants to balance their position closer to the physical delivery. The balancing market is the final stage for electricity trading and ensures the closing of any real-time deviations between supply and demand.

2.1.2 Electricity prices

Among all energy prices, electricity prices are particularly volatile in most spot markets. In an empirical study, Haar (2010) found the average annualised volatility of the EEX spot price from 2002 to 2008 to be 520%. The high volatility is due to the non-storability of electricity, the large fluctuations in the demand level with the hour of the day and the day of the year; the inelasticity of the electricity demand, and the required balance between production and consumption (Pineda & Conejo, 2013). Hydropower producers experience considerable uncertainty in respect to both electricity price and water inflow. As these are the two most important components of their revenues, hydropower producers have a solid incentive to actively managing their exposure to these risk factors. Power producers limit their risk exposure through hedging.

Hedging is a risk management practice employed to reduce the risk of an investment or cash flows through entering an offsetting position. Hedging is widely practised by companies that have incentives to avoid financial losses or operate in volatile markets. A hedged position can secure profits or cash flows through an increase of the value in the offsetting position when the original position decreases in value and vice versa.

Considering that the total supply of power is inversely related to the spot price, the inflow risk has a natural hedging effect in markets dominated by hydropower (Fleten, 2000). Regardless, the regional power markets in Scandinavia are integrated with the European power market. Thus, low water supplies in Norwegian reservoirs is balanced by the aggregated supply of the European market. The revenue risk has caused power producers to implement hedging strategies into their financial operation. The field of risk management for power producers has seen an increasing interest from both practitioners and academia over the last two decades (Fleten et al., 2010).

Hedging is normally achieved through investing in derivatives contracts, whose value is dependent on its underlying asset (McDonald, 2013). Examples include power derivatives, whose underlying is the electricity spot price. Power producers can thus reduce the cash flow uncertainty by investing in power derivatives, offsetting the spot price volatility.

2.1.3 Electricity derivatives

The electricity derivatives market provides power producers with an attractive opportunity of hedging their exposure to the various risk factors. Forward contracts with physical delivery are one such derivative and is an agreement to deliver electricity during a future delivery period at a pre-determined price. Selling electricity by forward contracts shield the producer from future price uncertainty, eliminating the risk associated with spot price volatility. A drawback of a forward contract is the obligated delivery. Under the circumstance of production failure or shutdowns, the producer must obtain the missing energy on the day-ahead market. If the spot price is much higher than the contract price in such an event, the company may incur large financial losses. This risk is termed availability risk (Pineda & Conejo, 2013). Power exchanges typically also offer Electricity Price Area Differential (EPAD) contracts, which can hedge the area price difference. The reference is the difference between the system spot price and the price in a specific bidding area. In addition, power options are offered on the exchange. These alternatives are commonly used (Sanda et al., 2013), but the contracts are significantly less liquid than the power futures.

A futures contract is a financial derivative and a legal agreement between two parties of buying or selling the underlying asset at a specified time in the future. Futures contracts specify the quantity of the asset delivered at expiration, in addition to the price, and are commonly traded on electronic exchanges. Futures contracts can be settled either by physical delivery or by cash settlement, according to the contract details. Cash settlement entails that the buyer does not receive the actual asset but rather the difference between the asset spot price at delivery and the current futures price.

The underlying asset for physically settled power futures is the delivery of electricity over a specified period. In contrast to traditional futures contracts, power futures are not settled at one specific point in time. Instead, the delivery occurs throughout the entire settlement period. The continuous delivery differentiates power futures from most other commodity futures. Power futures are typically offered through daily, weekly, monthly, quarterly, and yearly contracts, indicating the length of the settlement period. Futures can be listed as base load or peak load contracts. Baseload contracts cover delivery for all hours of the settlement period, while peak load contracts deliver during peak hours (typically between 8 a.m. and 8 p.m.) on a given day. Thus, a monthly power future contract of size 720 MWh, with baseload delivery, implies delivery of 1 MW throughout the whole month.

Most power futures are cash-settled. Financially settled power futures are the most liquid contracts in the European power markets. They are also the most commonly used hedging derivatives among power producers. The underlying of the financially settled power futures is commonly the day-ahead price for the market area.

Haar (2010) studies the price of EXX power futures. Empirical data indicates that power futures prices are less volatile than the spot price. The average annual volatility was 22%, which is notably less than for the spot price. The study also implies that the volatility of power futures is dependent on the duration of delivery and time to maturity. Futures contracts are priced based on the average expectational spot prices over the delivery period. A price shock for a given day or week will have a lesser effect on the futures price than the spot price. A more extended delivery period will therefore lead to lower volatility. The fact that volatility tends increase when the time to maturity approaches is known as the "Time to Maturity" hypothesis, or the Samuelson hypothesis (Duong & Kalev, 2008).

2.2 The limit order book

Due to sufficient market liquidity, traders usually do not have to concern themselves with adverse price movement for liquid securities such as blue-chip stocks. However, for an illiquid market, large orders can severely affect the market price. The mechanism by which adverse price movement occurs is the clearing of one or more order levels in the limit order book (LOB). While liquid securities also employ the LOB to structure orders, adverse price movement can be seen as more influential for trading of less liquid securities. In particular, there tend to be larger gaps in the price levels, and order depth is usually lower (Lee et al., 1993). As power futures markets can be categorised as illiquid, it is essential to consider the market microstructure when trading power. Therefore, we introduce central aspects of LOBs in the following section.

2.2.1 Limit order book dynamics

The LOB is a list of all orders with their corresponding prices for a financial instrument. LOBs are divided into a bid, and an ask side, representing the orders placed to buy and sell the security, respectively. The LOB is structured into order levels, where each level is described by its aggregated quantity and price. The LOB is sorted by prices such that the best order is shown first. Thus, the buy orders in the LOB are sorted in descending order, and sell orders are sorted in ascending order.

When limit orders are placed on the exchange, they are added to the LOB. If the price of a new limit order is equal to, or better, than the best price on the opposing side of the LOB, the matching mechanism of the exchange will match the order against the current LOB and execute it at the best available price. If there is more than one order at any particular order level, priority is determined by the first-in, first-out (FIFO) method in the case of order execution. A market order will match instantly with the best order(s) of the opposing side of the LOB and does not need a specified price. Thus, market orders do not enter the LOB. Market orders remove liquidity, as the total quantity left in the LOB is reduced by matching against orders from the LOB. Limit orders, however, increase liquidity as they provide liquidity by increasing the total volume in the LOB. For illustration, one instance of a LOB is seen in figure 2.1.

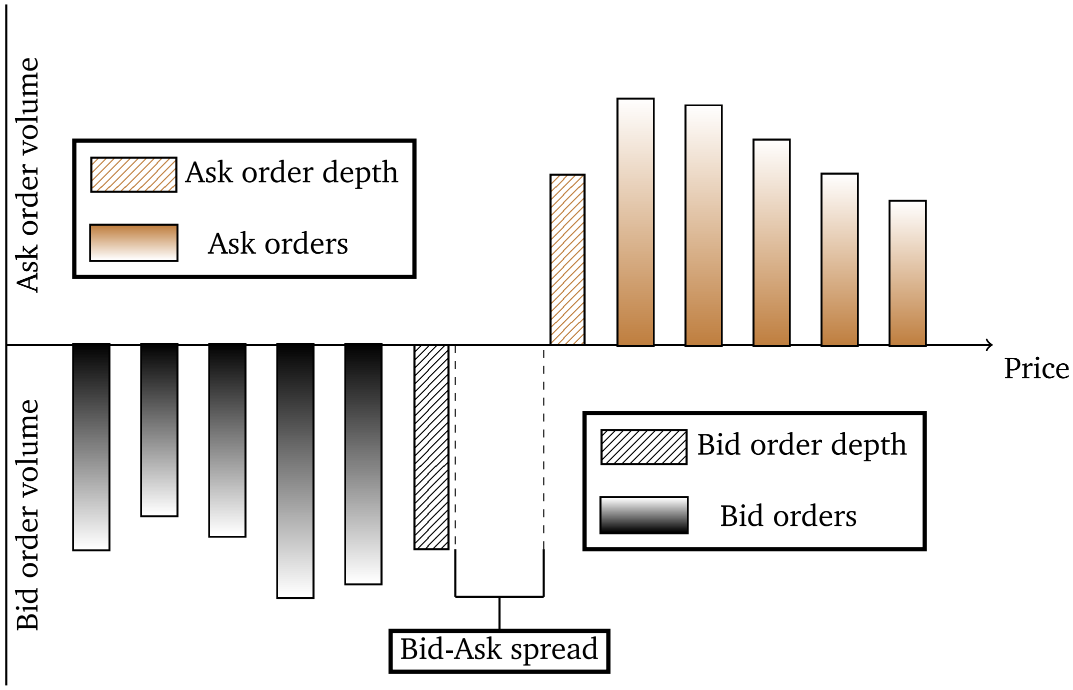

Figure 2.1: An instance of a limit order book. The rectangles represent the aggregated volume of limit orders at different prices in the LOB.

2.2.2 Liquidity measurements of the LOB

Liquidity is defined as the ease of acquiring or liquidating a significant position with minimal price impact (W. Liu, 2006). Several indicators in the microstructure of the LOB reflect the liquidity of a market, such as depth, bid-ask spread, and resiliency.

The volume of the first level in the LOB is called market depth. The depth of the LOB is often called the price impact trigger, as clearing this depth would lead to a new best price. Liquid markets tend to have higher depth than illiquid markets (Frestad, 2012). The bid-ask spread (BAS) is defined as the price difference between the best ask and the best bid in the LOB. Liquid markets tend to have tighter bid-ask spreads (Frestad, 2012). Usually, there is a supply and demand imbalance after a trade of significant size has been executed. This imbalance leads to a temporary price shift away from the equilibrium price level. Kyle (1985) defines the rate at which the price returns to equilibrium as the resiliency rate. If the resiliency rate is high, the trader can place orders frequently and still prevent large temporary price impacts due to the frequent arrival of other orders. Liquid markets tend to have a sufficient number of orders in the LOB and higher resiliency rates.

Other liquidity measures in the LOB include aggregated traded volume, order book imbalance, price volatility, and the LOB shape. It is more difficult to liquidate or acquire a large position when there is a low aggregated trading volume. Amihud and Mendelson (1986) find that the bid-ask spread tends to be negatively correlated with the aggregated trading volume. Order book imbalance measures the difference in volume in the buy and sell-side of the LOB, as seen in equation (2.1). A clear deviation from zero signals that there is market imbalance. Market imbalance can be used as a predictor of future price movement (Lehalle & Mounjid, 2018).

$$Imbalance(t) = \frac{V_{bid}(t) - V_{ask}(t)}{V_{bid}(t) + V_{ask}(t)} \tag{2.1}$$

Price volatility is not observable in the LOB. Rather it is a measurement of the fluctuation of either the mid-price, best bid or best ask. Illiquid markets tend to be more volatile than liquid markets (Cheriyan & Daniel, 2019). The higher volatility is due to the lack of volume at the best bid and ask, resulting in drastic price changes when a market participant places large market orders.

The last liquidity indicator that will be mentioned is the LOB shape. The LOB shape, also referred to as the slope of the LOB, is portrayed by transforming a LOB instance into a plot where the price is expressed as a function of the cumulative volume in the LOB. Examples of two LOB shapes are presented in figure 2.2. The shape indicates the change in the best bid price or the best ask price that will occur if a market order is placed. If the slope of the LOB is relatively flat, as seen in panel (a), the price will experience a minor change, whereas if the slope is steep, as seen in panel (b), the price change will be more significant. Consequentially, the market is less liquid if the shape of the LOB is steep.

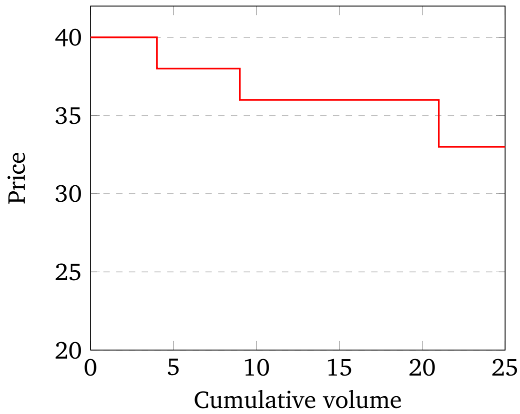

(a) Flat LOB shape

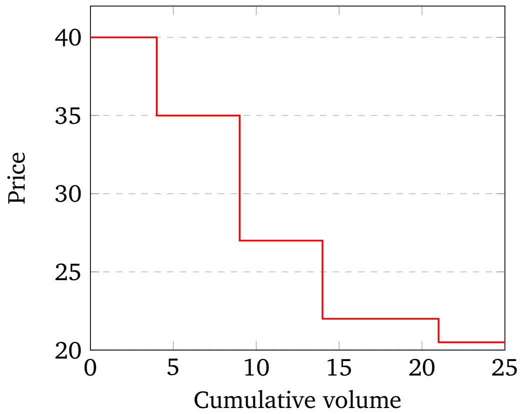

(b) Steep LOB shape

Figure 2.2: Two different instances of LOB shapes. Panel (a) shows a relatively flat LOB shape, while panel (b) shows a steeper LOB shape. LOBs are less liquid if their shape is steep rather than flat.

Chapter 3

Literature review

The purpose of this chapter is to explore trading models that have previously been used to solve the order execution problem and to place our contributions in the context of the related literature. Section 3.1 will introduce the order execution problem and details modelling assumptions found in the literature. Then in section 3.2, we present the motivation for hedging and discuss the theoretical frameworks for hedging.

3.1 The order execution problem

A market participant with a large market share needs to be aware of their impact on the market price when they liquidate or acquire many shares. This problem is known as the order execution problem and has been researched extensively for stock markets. The order execution problem is often formulated as an optimisation problem where the objective is to find the optimal trading frequency. The work of Almgren and Chriss (1999) was one of the first significant contributions to the field of study. Almgren and Chriss (1999) propose a dynamic programming model where the trader must liquidate (acquire) a position of X shares by the end of period T. Dividing the time period T into N intervals, each of length $\tau = T/N$, Almgren and Chriss (1999) define a set of discrete events at times $t_i = i\tau$, where the trader may reduce (add to) their position by selling (buying) x_t shares of the relevant stock. The objective is then to find the optimal trading trajectory, $\{x_0, x_1, x_2, ..., x_{T-2}, x_{T-1}, x_T\}$. Two fundamental elements are needed to formulate an order execution model. First, a process of how the price evolves during the period, referred to as the price dynamics, needs to be specified. Second, the objective of the liquidation (acquisition) process must be defined.

In section 3.1.1 we present previous literature about price dynamics formulations, while section 3.1.2 presents different ways the objective of the liquidation process has been defined. Last, the novel aspect of postponement optionality with reference to previous literature is proposed in section 3.1.3.

3.1.1 Price dynamics

The contemporary literature models price dynamics in many ways, yet they all tend to be composed of exogenous and endogenous price factors.

Exogenous price factors: Exogenous factors of the price dynamics include drift and volatility. Papers such as Bertsimas and Lo (1998), Almgren and Chriss (1999) and Bertrand (2021) assume the exogenous price factor follows a discrete arithmetic random walk process. This assumption allows for an optimal discrete trading trajectory. On the other hand, Almgren et al. (2005), Almgren (2012), Shen (2014), and Shen (2017) assume that the exogenous price factor follows a linear Brownian motion. It is common to assume that the drift is zero due to a short time-interval of the trading period. The problem formulation in these papers is slightly different from those assuming a discrete arithmetic random walk since Brownian motions are time-continuous. Those papers assuming a Brownian motion aim to find the optimal continuous trading frequency function rather than the optimal discrete trading trajectory.

Bertsimas and Lo (1998) also include price driving features as part of their exogenous price process. Price driving features are exogenous variables that are correlated with the price of the asset. The logic behind this inclusion is that price driving features send signals about the market's view of the price in the future. Bertsimas and Lo (1998) use S&P500 returns as a price driver.

A majority of the contemporary literature studies order placement for blue-chip stocks. To obtain closed-form solutions, it is necessary to assume that the volatility and liquidity are constant or deterministic. This is assumed to be the case for blue-chip stocks, but not for less liquid stocks. However, Almgren (2012) does obtain a closed-form solution for a model where the volatility varies, with the assumption that the product of volatility and the permanent price impact is constant. With this in mind, implementing a model with varying volatility or liquidity would benefit the field.

Endogenous price factors: The endogenous factor of the price dynamics is the price impact of endogenous orders[1]. Price impact is often divided into subcategories. Almgren et al. (2005) make the distinction between temporary and permanent price impact, illustrated in figure 3.1. These subcategories are often modelled separately.

Permanent price impact: The price impact element that affects the market price over an extended period of time is referred to as the permanent price impact. The permanent price impact is related to the informational signal that a trader sends when placing orders. Bertsimas and Lo (1998) and Almgren and Chriss (1999) model the permanent price impact as a linear function of the traded volume. Almgren et al. (2005) explore the possibility that the permanent price impact function is a power law of the trading rate x_t/τ. By analysing the trading activity of US stocks over 19 months starting from December 2001, they infer that the hypothesis of linear permanent price impact cannot be rejected. However, the validity of their approach is limited as they do not consider changing market liquidity throughout the day. According to Shen (2017), price impact is always relative to the liquidity of the security. During periods of low market liquidity, large market orders will impact the price to a greater extent than when the market liquidity is high. To include the market liquidity, Shen (2017) models the permanent price impact instead as a linear function of the *participation of volume* (PoV), which is the ratio between the order size and the total market volume in the LOB at time t. One flaw with this approach is that Shen (2017) assumes that the

[1] Endogenous orders: Self-placed orders.

market volume is deterministic and known.

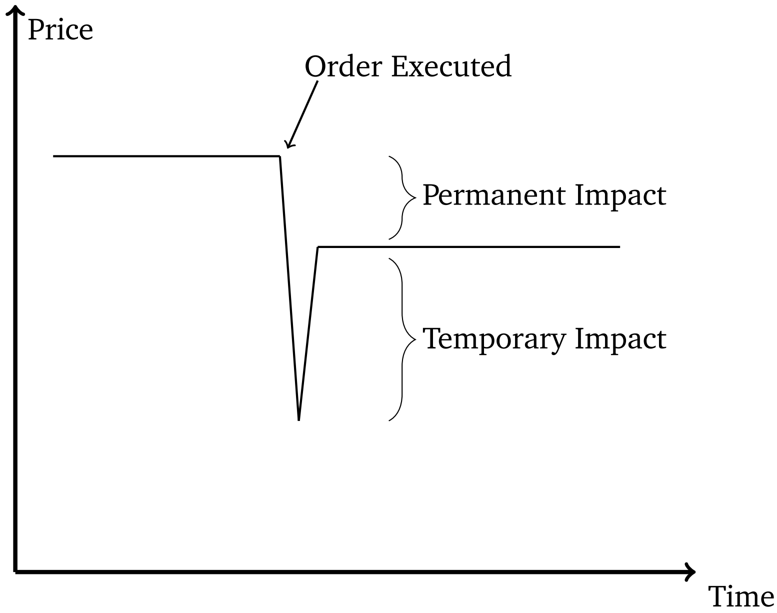

Figure 3.1: Illustration of the permanent and the temporary price impact of a trade (Mete Soner, 2015). Note that the market does not fully recover, which is due to the permanent impact of the executed order.

Temporary price impact: The endogenous price impact element that affects the price due to changes in the market liquidity after a large market order is referred to as the temporary price impact. Shen (2014) and Akersveen and Graabak (2018) distinguish between *transient* and *instantaneous* price impacts. The transient price impact reflects that a large market order has removed a large fraction of the market liquidity. The transient price impact creates a temporary supply and demand imbalance in the market but recovers over time due to resiliency. Given that there is sufficient time between trades, the transient price impact may be neglected, while the effect of instantaneous price impact must still be considered.

Instantaneous price impact results from the LOB depth being smaller than the order volume. Thus the average clearing price is worse than the best bid (ask) before placing a sell (buy) order. Instantaneous price impact is the only price impact factor that directly affects the price of a trade, as opposed to permanent and transient price impacts, which relate to the market price after a trade. The instantaneous price impact is illustrated in figure 3.2.

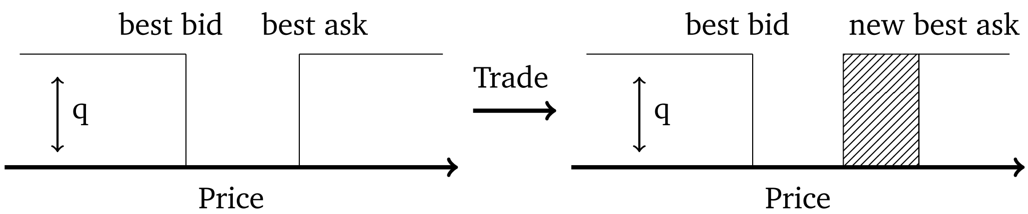

Figure 3.2: Illustration of the instantaneous price impact of a trade. The dashed area is the volume which the incoming order removes from the LOB (Mete Soner, 2015). Note that the average execution price will be the volume-weighted average price of the volumes cleared from the LOB.

Bertrand (2021) models temporary price impact strictly as instantaneous price impact. The instantaneous price impact is approximated by using the slopes of the LOB. Figure 3.3 illustrates how the slopes are modelled using the shape of the LOB. This method for modelling the LOB slope yields a continuous function for the price impact rather than the discrete structure of the price levels in the LOB. Bertrand (2021) assumes that the slopes are deterministic. As opposed to Bertrand (2021), Shen (2014) models the instantaneous price impact as a linear function, with PoV as the explanatory variable, again taking market liquidity into account.

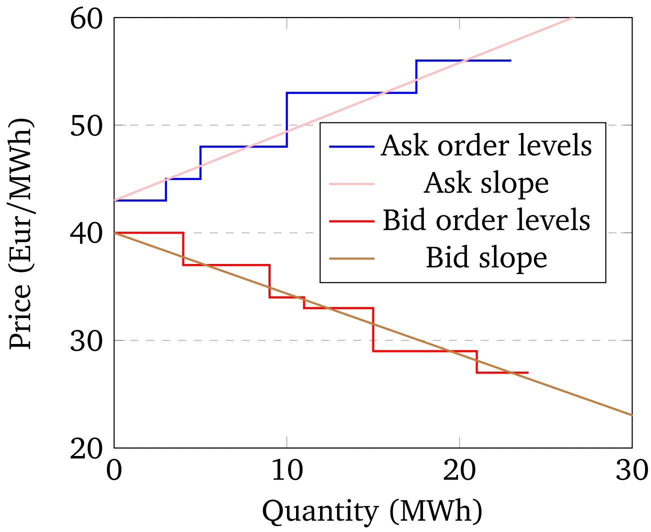

Figure 3.3: Illustration of approximated slopes of the bid and ask side of the LOB.

3.1.2 Defining the model objective

The objectives of the order execution models are what distinguishes their solutions from each other. A trader needs to make the trade-off between two factors, namely price impact and price risk. A risk-averse trader prefers to liquidate their position as quickly as possible to reduce the price risk, yet this could lead to significant price impact. A simplistic model would ignore price risk and maximise the total expected revenue. A heuristic in line with this model is the time-weighted average price (TWAP) trading strategy, which allocates volume evenly throughout the trading period.

Other models, such as the one formulated by Shen (2014), instead consider an objective function that combines revenue and price risk. The risk preferences of the trader are introduced by adding the variance of the revenues multiplied with a Lagrange multiplier, λ, in the objective function. The trading trajectory is then dependent on the trader risk preference, as seen in figure 3.4. By increasing the value of λ, the tolerance of price risk is reduced. Therefore, the trading trajectory will be more convex, resulting in front-loading. Almgren and Chriss (1999) explain that the convexity of the trading trajectory curve will not only be dependent on the risk tolerance but also the size of the price impact. Almgren and Chriss (1999) argue that if the price impact is small, the risk term will be dominant in the objective function, resulting in a more convex trading trajectory. However, if the price impact is large, the trajectory will be close to linear, as the price impact will carry a higher relative weight in the objective function.

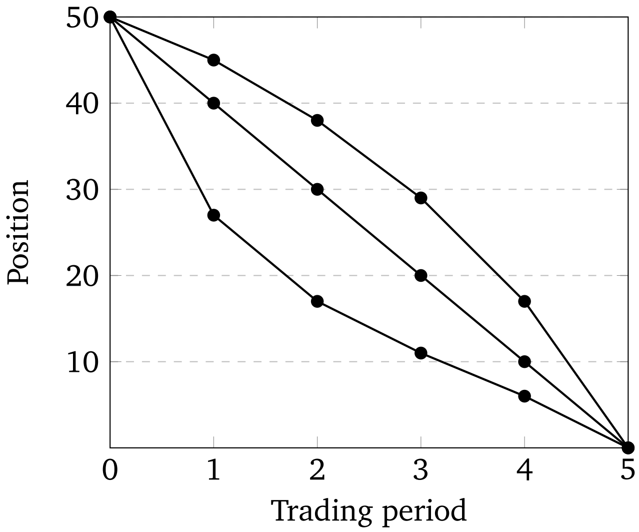

Figure 3.4: Trading trajectories based on a trader's risk preference. The convex trading trajectory belongs to a risk-averse trader, corresponding to a high λ value. The linear trajectory belongs to a risk-neutral trader, corresponding to a λ value of 0. The concave trajectory belongs to a risk-seeking trader. Risk-averse traders will front-load, liquidating their position quicker than risk-neutral investors to minimise their exposure to price risk.

Other risk measures are also applicable to ensure models take risk preference into account. For instance, Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR) are risk measures that only consider tail risks. VaR estimates the expected loss, given a predetermined confidence interval, α. For instance, a VaR_α value of -5.0% means that there is a probability of α that the trader will not experience a loss larger than -5.0%, where $0 \leq \alpha \leq 1$. CVaR builds upon VaR and quantifies the expected loss, given that the loss exceeds the VaR. VaR and CVaR are expressed in equation (3.1) and (3.2), where $F(x)$ is the cumulative distribution of variable x. For instance, x can be returns or revenues. Feng et al. (2012) use CVaR rather than variance as the risk determinant in their order execution model, arguing that using variance penalises both negative and positive deviation. In contrast, CVaR only penalises the negative tail risk.

$$\text{VaR}_\alpha = inf\{x|F(x) \leq \alpha\} \tag{3.1}$$

$$\text{CVaR}_\alpha = \mathbb{E}[x|x \leq \text{VaR}_\alpha] \tag{3.2}$$

In contrast to Almgren and Chriss (1999) and Feng et al. (2012), Hora (2006) models the price risk as a delay cost term. The delay cost is a quadratic term of the daily volume left to trade at time t, I_t, multiplied with a scaling parameter ρ. This approach encourages front-loading, yet Shen (2017) criticises this approach since the delay cost remains constant throughout the day. To improve upon Hora (2006), Shen (2017) models the scaling parameter as a monotonically increasing term. Despite Shen's criticism, one could argue that the proposed approach by Hora (2006) captures the preference for early trading, as the penalty cost for the same volume will be included twice if the trader does not place any volume for subsequent trades.

Ruszczyński and Shapiro (2006) suggest conditional risk mappings for modelling of risk preferences in multistage stochastic optimisation models. Due to the dynamic structure and the sequential composition of such models, there is an argument of using conditional risk measures that represent future risk based on information available at the current stage t (Ruszczyński & Shapiro, 2006). The nested conditional value-at-risk (nested CVaR) is one such risk measure, applied by Löhndorf and Wozabal (2021) for a gas storage valuation model. Following the notation of Löhndorf and Wozabal for the nested CVaR, a set of random variables are defined as $X_1, \ldots, X_T$. Then, stage t random variables are defined as

$$\mathbb{V}_{t,\alpha,\lambda}(X_t) = \lambda CVaR_{t,\alpha}(X_t) + (1-\lambda)\mathbb{E}(X_t|X_{t-1}) \tag{3.3}$$

The nested CVaR can then be defined as

$$CVaR^{NEST}_{1,\alpha,\lambda}(X_1, X_2, X_3, \ldots) = X_1 + \mathbb{V}_{1,\alpha,\lambda}(X_2 + \mathbb{V}_{2,\alpha,\lambda}(X_3 + \cdots)) \tag{3.4}$$

The nested CVaR is thus a convex combination of expected value and CVaR with a recursive combination of other convex combinations of expected value and CVaR (Löhndorf & Wozabal, 2021).

3.1.3 Expanding the model objective

The order execution models discussed so far share the commonality of requiring order execution within the specified time domain. Additionally, they have also been restricted to studying the price impact of market orders. By introducing the option of placing limit orders, the trader experiences the risk of no (or partial) execution. A majority of the research conducted on limit order placement addresses how to model execution risk. Cont and Kukanov (2017) derive a closed-form solution for the dynamic decision between placing market or limit orders. The risk of non-execution of limit orders is modelled as a penalty, increasing linearly with the outstanding volume. Agliardi and Gençay (2017) build on the work of Cont and Kukanov (2017) but allow for choosing the limit order price freely. In their model, order placement frequency is treated in a similar fashion to Almgren and Chriss (1999), where N trades are scheduled with equidistant time between trades. If the volume requirement is not met by the end of the trading period, a terminal market order is sent with the outstanding volume. In their model, the cost of the terminal market order increases quadratically with order size. The aggressiveness of the limit order price relates to the trade-off between the cost of non-execution and revenues. Unlike the model of Cont and Kukanov (2017), the cost of non-execution is captured in the transaction cost of the terminal market order. As will be explained next, a similar approach for quantifying the cost of execution risk will be used to model the cost term for the postponement optionality.

The cost of non-execution builds on the assumption that the required volume needs to be executed for the specified period. However, not all traders operate with such rigid constraints. By this token, one can distinguish between traders who need to trade and those who may choose to trade. Traders who choose to trade are said to be tolerant of execution risk. These traders are faced with the decision of trading at current prices or postponing order execution. By the same reasoning underlying an execution risk cost term, we introduce a cost term in the objective function that quantifies the value of postponing order execution. To the authors' best knowledge, there is a gap in the order execution literature concerning quantifying the value of waiting with order execution.

Therefore, this thesis will contribute to the literature by providing a method to quantify the value of postponing order execution for an electricity producer classified as a "choose-to" trader.

3.2 The hedging problem

Hedging has traditionally been perceived as a means for risk-averse producers to reduce the diversifiable risk of their profits through offsetting in financial instruments (Anderson & Danthine, 1980). Under the assumption of efficient markets, Modigliani and Miller (1958) argue that hedging cannot increase firm value, although it can be successful in reducing risk exposure. However, both empirical evidence and theoretical arguments exist for occurrences of a hedging premium in the literature on hedging. Lin and Chang (2009) find that U.S. airlines hedging their jet fuel costs, on average, are valued higher than similar airlines with no hedging policy. Similarly, Allayannis and Weston (2001) examine the relation between firm value and the use of foreign exchange derivatives. Using 720 large U.S. nonfinancial firms between 1990 and 1995, they find a hedging premium of 4.9% on firm value.

3.2.1 Predictive and selective hedging

Stulz (1996) finds that most companies allow their views on price and market movement to influence their hedge ratios. Incorporating this type of speculation in the hedging practice is referred to as selective hedging. Supporting the concept of selective hedging, Adam and Fernando (2006) differentiate between selective and predictive hedging practices. In contrast to selective hedging, predictive hedging is the method of hedging predicted cash flows from a company's operations, independently of market view. Adam and Fernando (2006) study hedging in the gold mining sector and find significant evidence of selective hedging in their sample. Selective hedging was found to yield, at best, small increased firm value and large cash flow variance.

Sanda et al. (2013) study the hedging policies of 12 Norwegian hydropower companies and finds that selective hedging is widely practised in the sample companies. Furthermore, it was found that the firms obtained substantial profits from their hedging activities. Stulz (1996) argues that while selective hedging, mainly practised by large corporations, contrasts with theoretically prescribed risk management methods, it can increase corporate value if the hedging firm has a comparative informational advantage. This informational advantage, acquired through its ordinary business activities, allows the firm to predict price movements more accurately than other market participants.

Table 3.1: Hedging practice of 12 different power producers, studied by Sanda et al. (2013). Static hedging policies are the most common hedging practices among the Norwegian power producers.

Practice	Number of companies
No written policy	2
Static hedging policies	8
Cash Flow at Risk requirement	2
Total	12

3.2.2 Hedging procedures

Dupuis et al. (2016) separate hedging procedures into two categories: static and dynamic. With a static hedging procedure, hedging assets are bought at once and not rebalanced in the following periods. For dynamic hedging, the hedging portfolio is rebalanced over time as new information becomes accessible. Dynamic hedging procedures can be further separated into two sub-categories: local and global hedging (Dupuis et al., 2016). Local hedging procedures aim to minimise the portfolio risk until the next rebalancing, while global hedging minimises the risk associated with all future cash flows.

Näsäkkälä and Keppo (2005) consider an electricity producer with static hedging strategies that maximise the risk-adjusted expected value of its cash flows. Fleten et al. (2010) propose an optimisation model to find static hedge strategies for a hydropower producer. Dynamic hedging strategies are studied by Dupuis et al. (2016), Fleten et al. (2002), Zanotti et al. (2010), S. D. Liu et al. (2010), Kettunen et al. (2007) and Pineda and Conejo (2013). Pineda and Conejo (2013) propose a multi-stage stochastic model to dynamically obtain the most suitable portfolio of options and forward contracts subject to uncertain power production. Fleten et al. (2002) coordinate the power production of a hydropower producer with the global dynamic hedging of forward and option contracts by use of a stochastic programming model. The study finds that the dynamic hedging approach yields higher expected returns than static hedging.

Chapter 4

Formulating the Integrated Postponement model

In this chapter we present the `Integrated Postponement model`. This is done by decomposing the model into two sub-models: A dynamic hedging model and an order execution model, which will be referred to as the `Postponement model`. The hedging problem is solved by the `QUASAR Dynamic Hedging model`, a stochastic dynamic programming model developed by Quantego. As the authors of this papers have not contributed to this model, we simply summarise how the model is solved and its features concerning risk-preferences and liquidity. For the order execution model, we decompose the `Postponement model` such that each aspect of the model design is explained.

The rest of this chapter is structured in the following way; in section 4.1, the `QUASAR Dynamic Hedging model` is introduced. In particular, the impact of the LOB slopes on the optimal hedging policy is outlined. Section 4.2 presents the `Postponement model` in detail. The underlying assumptions are outlined, and the objective function and constraints are presented in a stepwise manner. The integration of the `QUASAR Dynamic Hedging model` and the `Postponement model` is explained in section 4.3. Finally, alternative trading strategies that will serve as benchmarks are presented in section 4.4.

4.1 The QUASAR Dynamic Hedging model

This paper uses a hedging model developed by Quantego, which from this point on will be referred to as the `QUASAR Dynamic Hedging model`. This model is an alteration of the model proposed by Dimoski et al. (2019). Dimoski et al. (2019) presents a global dynamic hedging model for a Norwegian hydropower producer participating in the Nordic electricity market. The authors use a sequential approach, first running a dynamic production planning model to obtain optimal production policies. By applying these production decisions as endogenous variables in a dynamic hedging model, they obtain trading decisions for derivatives contracts. The hedging model uses power futures contracts and currency futures contracts to hedge price and exchange risk, respectively. Risk-preferences are modelled by the nested conditional value-at-risk (nested CVaR).

In contrast to Dimoski et al. (2019), this thesis will not apply a production planning model, instead, the `QUASAR Dynamic Hedging model` obtains exogenous production volumes as input and solves for the optimal daily hedging targets. Currency risk will also not be considered. As opposed to Dimoski et al. (2019), the `QUASAR Dynamic Hedging model` model considers the liquidity of the LOB. The liquidity aspect is manifested in the `QUASAR Dynamic Hedging model` as a price impact penalty with the use of the LOB slope. In this thesis, the producer trades financial futures contracts with monthly and quarterly delivery periods. Power futures contracts were deemed most suitable because they are the most liquid derivatives in the European power market (EXX, 2021b). Liquidity is crucial as this thesis takes the perspective of a power producer with a large market share. Furthermore, futures contracts are the most common hedging derivatives used by Norwegian power producers (Sanda et al., 2013). The model has daily granularity and allows for both selling and buying of futures contracts. It is, however, restricted to only taking short positions in power futures. As the producer has long positions in their physical production and aims to reduce their risk exposure to the spot price, short positions are sufficient to meet this objective.

4.1.1 Solution method

Given the large number of variables and stages in the hedging model, an efficient solution algorithm is required for the model to be computationally tractable. The main issue with such a high-dimensional problem is that the decision space can become too large to find the optimal decisions for all stages within a reasonable amount of time. Therefore, we require a method that resolves this issue by obtaining decision policies that are approximately optimal. The ADDP algorithm (Löhndorf et al., 2013) serves this purpose and is used to solve the `QUASAR Dynamic Hedging model` efficiently. ADDP integrates stochastic dual dynamic programming (SDDP) (Pereira & Pinto, 1991) with methods from approximate dynamic programming (ADP). SDDP involves formulating the problem as a dynamic program and then applying Bender's decomposition to recursively construct the value function at each stage around a set of sample decisions (Pereira & Pinto, 1991). SDDP can handle problems with a large number of stages as long as the optimisation problem at each stage is convex and the stochastic process is stage-wise independent. ADP algorithms simulate the state transition process of a Markov Decision Problem (MDP) and use the sampled information to approximate the high-dimensional value function by a function of much lower complexity (Powell, 2011). As with SDDP, ADDP iteratively solves the decision problem using forward simulation to obtain possible optimal solutions and backwards recursion to construct the approximate future cost function (Löhndorf et al., 2013; Pereira & Pinto, 1991). In contrast to SDDP, however, ADDP assumes that random variables follow a Markov process.

ADDP requires discretising the evolution of the state variables into a scenario lattice. Future scenarios of spot and futures prices are generated using movements in a forward curve. A forward curve estimates the future spot price for delivery at specific points in time, based on all contracts available in the market. As the time of delivery approaches, the forward price for delivery on that specific day will tend towards the spot. As done by Dimoski et al. (2019), the model uses the HJM framework (Heath et al., 1992) to generate future scenarios of the underlying spot price.

4.1.2 LOB slope and risk measure

The slope of the LOB is a measure of the average price elasticity across all price levels with the corresponding volumes. A price impact penalty term that incorporates the slope of the LOB regulates the trading volume of contracts with respect to their slopes. Consequently, the hedging model regulates the trading volumes of contracts with a steep slope. Such contracts are traded in smaller volumes in periods where the LOB slope is steep, compared to periods where the LOB slope is lower. In the case of power futures contracts, the LOB slope is generally higher far from delivery when the contracts are less liquid. Therefore, the inclusion of a price impact term will have a backloading effect, meaning that futures contracts will be traded with higher volumes closer to their delivery. Using the LOB slope creates a trade-off between entering a significant hedging position early, potentially with a hefty price impact penalty, versus waiting, thus bearing a higher risk associated with the unhedged position.

4.1.3 Hedging model objective

For the `QUASAR Dynamic Hedging model`, the objective function is defined as a linear combination of the stage d cash flows and a dynamic risk measure term, reflecting the risk-averse preferences of a power producer. Stage d cash flows are composed of cash flows from power futures trading and spot sale from physical power production. The hedging model is solved at every stage d, corresponding to each trading day, to maximise the corresponding objective function. The optimisation yields the stage d optimal decision policy π_d, from which the first-stage decisions, π_d^0, are obtained. π_d^0 represents the daily trading target of every tradable futures contract on the current stage d, and is used in the `Postponement model`.

4.1.4 Modelling the constraints of the QUASAR Dynamic Hedging model

The `QUASAR Dynamic Hedging model` includes variables and balance constraints for tracking financial short positions in power futures and committed cash flows, which reflects the actual payoff structure of a producer.

Let u_{d,M_i}, u_{d,Q_j} and u_{d,Y_k} [MWh] denote the aggregate short position at stage d in futures contracts with delivery in month i, quarter j and year k, respectively. In addition, let w_{d,M_i}, w_{d,Q_j}, w_{d,Y_k} [MWh] denote new short positions in a futures contract entered into at stage d, for month i, quarter j and year k, respectively. We let $\mathcal{D}^{M_i}$, $\mathcal{D}^{Q_j}$ and $\mathcal{D}^{Y_k}$ denote the sets of trading days d for futures contracts with delivery in month i, quarter j and year k, respectively.

While there are more types of constraints included in the hedging model, this thesis will only present the position balance constraints, as formulated in equation (4.1). These constraints have an important implication for the modelling of the postponement option in the `Postponement model`. For a more complete formulation of the constraints in the hedging model, the reader is referred to Dimoski et al. (2019).

$$
\begin{aligned}
u_{d,M_i} &= u_{d-1,M_i} + w_{d,M_i}, \quad d \in \mathcal{D}^{M_i} \\
u_{d,Q_j} &= u_{d-1,Q_j} + w_{d,Q_j}, \quad d \in \mathcal{D}^{Q_j} \\
u_{d,Y_k} &= u_{d-1,Y_k} + w_{d,Y_k}, \quad d \in \mathcal{D}^{Y_k}
\end{aligned}
\tag{4.1}
$$

The total position in a futures contract after trading day d is equal to the position in the contract before the trading day, plus the amount traded during day d. From the constraints in equation 4.1, a shadow price η_d can be obtained for every tradable contract d. This shadow price represents the change in the objective function by relaxing the constraint by one small increment. Thus, these shadow prices represent a quantifiable value of waiting to enter a hedged position in a futures contract.

4.2 The Postponement model

The order execution problem is restricted to trade execution. Trading target volumes are obtained as the first stage decisions of the hedging policy π_d^0 for trading day d, as described in section 4.1. A MDP model is proposed with inspiration from previous literature where the price impact is considered. Additionally, postponement optionality is introduced. Note that we present the model without including a contract index. This is done to present the model with a reasonable level of complexity. The structure of the model is similar to the model in Almgren and Chriss (1999), where the policy to be determined is the trading trajectory $\mathbb{X} \in \{x_0, x_1, x_2,..., x_{T-2}, x_{T-1}, x_T\}$. Elements in $\mathbb{X}$ are the volumes to be traded at stage t. Order placement is treated as a discrete process, and the trader is restricted to placing orders at these points during the daily trading period. Expressing the order execution problem as a dynamic model necessitates a set of governing assumptions, in addition to defining a price dynamics model. These aspects are considered in the following sections.

4.2.1 Trading assumptions

Several assumptions are made in order to formulate the problem as a multistage stochastic programming model. The assumptions are stated and justified below.

Assumption 1 - Neglecting transaction costs: Trading exchanges charge transaction fees for trading activity. This could either be a flat fee, or a variable fee per MWh traded. As of June of 2021, the transaction costs on EEX are limited to a variable fee of €0.0075/MWh (EXX, 2021a). Since all strategies trade equal volumes for all contracts, the trading costs will not be considered.

Assumption 2 - Only market order placement: Limit order placement could improve trading performance, yet runs the risk of non-execution. Combining limit and market orders in the `Postponement model` would involve creating execution scenarios, a non-execution penalty term, as well as the price impact. Due to the high computational complexity of creating execution scenarios, this thesis is limited to market order placement.

Assumption 3 - Constant volatility: Prior to constructing the `Postponement model`, a price volatility study was conducted for the EEX data. This study is presented in Appendix A. Volatility has been tested as a function of time to maturity and an AR(1) process. The results indicate that neither of these models are fit as predictors. Therefore, the volatility is set to be constant, as in Almgren and Chriss (1999). Haar (2010) finds the annualised volatility of futures contracts to be 22%. The findings of Haar (2010) correspond well to the volatilities of the futures contracts in the data set, seen in Appendix B. Therefore it is assumed that volatility is constant with a value of 22% annualised.

Assumption 4 - Best bid as price: Since the `Postponement model` only places market sell orders, we use the best bid price to model the price dynamics.

Assumption 5 - Continuous trading volumes: Mixed integer problems are more computationally expensive compared to continuous problems. It is therefore assumed that the trading volumes are continuous. This reduces the run time of the backtest substantially. One flaw with this assumption is that the minimal tick size[1] for orders on the EEX exchange is 1 MW.

Assumption 6 - Only instantaneous price impact: It is assumed that the time interval between order placement is sufficient for the transient price impact to have vanished before the subsequent trade. Thus, transient price impact is not considered. Ignoring permanent price impact is justified by the modelling restrictions. The herd effect is mostly caused by traders acting on emotions rather than their market view (Bikhchandani & Sharma, 2001). This behaviour is what causes permanent price impact but will not be present in the historical data for endogenous orders.

4.2.2 Price dynamics

The trading trajectory will be dependent on the price dynamics process of the security. A price process is necessary to describe the state of the price variable at future stages for the `Postponement model` to make trading decisions. For the `Postponement model`, the features included in the price model can be categorised as exogenous or endogenous.

Exogenous price features: The exogenous feature of the price is modelled as an arithmetic random walk. In preparation for constructing the `Postponement model`, a study on price drivers was conducted to check whether microstructural features are correlated to the price and if they could be used for predicting future prices. The feature study is presented in Appendix C. Had features turned out to be good price predictors, they would have been included in the price model. Since this was not the case, the arithmetic random walk in equation (4.2) was used as the exogenous price process.

$$p_{t+1} = p_t + \varepsilon_t \quad \text{for } t = 0, 1, \dots, T-1. \tag{4.2}$$

$$\varepsilon_t \sim N(0, \sigma^2) \tag{4.3}$$

Endogenous price features: The endogenous price feature is the price impact caused by endogenous order placement. In this paper, price impact is limited to temporary price impact. As in Bertrand (2021), only instantaneous price impact is accounted for. The impact is modelled by using the slope of the LOB, similarly to in Bertrand (2021). The LOB slope is assumed to be constant throughout the day. The marginal price with respect to x_t is given by equation (4.4).

$$\tilde{p}_t(x) = p_t - \omega \cdot x_t \tag{4.4}$$

Here, the LOB slope is given by ω. As shown by Bertrand (2021), the total revenue from trading x_t units at stage t is:

[1] Tick size: the minimum incremental amount at which one can trade a security.

$$R_t(x_t) = p_t \cdot x_t - \frac{\omega}{2} \cdot {x_t}^2 \tag{4.5}$$

where equation (4.5) is equation (4.4) integrated with respect to x_t from 0 to $\mathbf{x_t}$.

4.2.3 Trader risk preference

The `Postponement model` is a multistage stochastic programming model, thus a dynamic risk measure reflecting the sequential decision making will be used. The risk preferences of the producer is expressed by the nested CVaR, similarly to Shapiro et al. (2013) and Löhndorf and Wozabal (2021). For a sequence of revenues $R_0(x_0)$, $R_1(x_1)$, ... , $R_T(x_T)$, corresponding to each stage t of the multistage problem, we can define the stage t random variables

$$\mathbb{V}_{t,\alpha,\lambda}[R_t(x_t)] = \lambda CVaR_{t,\alpha}(R_t(x_t)) + (1-\lambda)\mathbb{E}[R_t(x_t)|R_{t-1}(x_{t-1})] \tag{4.6}$$

Then, the nested CVaR of all future revenue streams, $R_t(x_t)$, is expressed by equation (4.7).

$$CVaR^{NEST}_{0,\alpha,\lambda}\Big[R_0(x_0),...,R_T(x_T)\Big] = R_0(x_0) + \mathbb{V}_{0,\alpha,\lambda}\Big[R_1(x_1) + \cdots + \mathbb{V}_{T,\alpha,\lambda}[R_T(x_T)]\Big] \tag{4.7}$$

The weighting λ and the significance level α of the nested CVaR terms in the objective function, are calibrated to suit the risk preferences of the Norwegian power producer, who tends to be risk-averse (Dimoski et al., 2019).

4.2.4 Postponement optionality

In contrast to the existing literature, we propose a model that allows for the postponement of trading volume to the subsequent trading period. This optionality reflects the "choose-to-trade" characteristic of a power producer. We introduce the state variable, Y_t, which denotes the volume left to trade for the given day, at trading stage t. Consequently, the `Postponement model` follows the transition function in equation (4.8).

$$Y_{t+1} = Y_t - x_t \tag{4.8}$$

Y_t will then be the volume that has not been traded through the day, which will correspond to the amount of volume that we decide to postpone to the next trading period.

Further on, let η_d denote the shadow price of the position balance constraint (4.1) in the `QUASAR Dynamic Hedging model`. The shadow price expresses the additional value of increasing the balance by one unit. We use the shadow price to express the postponement optionality value. The cost term of the postponement optionality, $C(Y_t)$,is defined in equation (4.9). Here, ζ is a scaling factor. Going forward, $\zeta \cdot \eta_d$ will be referred to as the price-equivalent shadow price, p_{SP}.

$$C(Y_t) = \zeta \cdot \eta_d \cdot Y_T \tag{4.9}$$

4.2.5 Postponement model formulation

The `Postponement model` is presented below. Combining the risk preferences of the producer, expressed by the nested CVaR, with the postponement optionality, the objective function is expressed as equation (4.10). The recursive formulation of the objective function ensures that the optimal decisions x_t are non-anticipative, meaning that decisions are made only based on information available until the current stage t. Constraint (4.11) represents the transition function for the postponement volume, and constraint (4.12) the exogenous price process. The `Postponement model` is formulated as a multistage stochastic programming model. The model is implemented over a descending horizon to solve the optimisation problem at every trading point throughout the trading period. To the authors' knowledge, this is a novel approach in the order execution problem field.

$$Max \quad Z = CVaR^{NEST}_{0,\alpha,\lambda}\Big[R_0(x_0), R_1(x_1), ..., R_T(x_T)\Big] + \zeta \cdot \eta_d \cdot Y_T \tag{4.10}$$

s.t

$$Y_{t+1} = Y_t - x_t \quad \text{for } t = 0, 1, \ldots, T-1. \tag{4.11}$$

$$p_{t+1} = p_t + \varepsilon_t \quad \text{for } t = 0, 1, \ldots, T-1. \tag{4.12}$$

$$\varepsilon_t \sim N(0, \sigma^2) \tag{4.13}$$

$$Y_0 = \pi_d^0 \tag{4.14}$$

$$0 \leq x_t \leq \pi_d^0 \quad \text{for } t = 0, 1, \ldots, T. \tag{4.15}$$

$$0 \leq Y_t \leq \pi_d^0 \quad \text{for } t = 0, 1, \ldots, T. \tag{4.16}$$

4.3 Combining the hedging and order placement models

The `QUASAR Dynamic Hedging model` and the `Postponement model` have been presented in the two prior sections. In this section, we describe how the models are implemented together. This includes a description of the informational flow of data used as input factors in the models.

The `QUASAR Dynamic Hedging model` is initially solved for the remaining trading period. The set of first-stage hedging decisions for tradeable futures contracts, π_d^0, are retrieved from the `QUASAR Dynamic Hedging model`, in addition to the set of associated shadow prices η_{cd} for all traded contracts c. This information is transferred to the `Postponement model` before the daily trading period starts, as seen in figure 4.1. The trading decisions and shadow prices are used as parameters in the `Postponement model`. At each trading stage t, corresponding to a new hour in the daily trading period, the `Postponement model` solves the optimisation problem, and obtains the set of first-stage trading decisions, x_{tc}, for the daily order execution. At the end of the daily trading period, the aggregated trading volumes are transferred back to the `QUASAR Dynamic Hedging model`. New trading targets are then obtained as the first-stage decisions from the optimal decision policy π^0_{d+1}, by solving the `QUASAR Dynamic Hedging model` for the next stage $d+1$.

The `Integrated Postponement model` will dynamically adjust the daily trading targets in response to the outcome of trade execution for the previous day. If the `Postponement model` deems complete or partial postponement to be the optimal decision on trading day d, the shadow prices for trading on trading day $d+1$ should readjust to reflect the current state of the balance constraint (4.1). This dynamic approach to order placement is motivated by the natural tendency of market conditions to change. The fact that a decision policy is currently considered ideal does not automatically entail that it is ideal at a future stage. For instance, the price may change dramatically through the day, which the `QUASAR Dynamic Hedging model` does not take into account. By including postponement optionality, trading decisions are reactive to the market conditions.

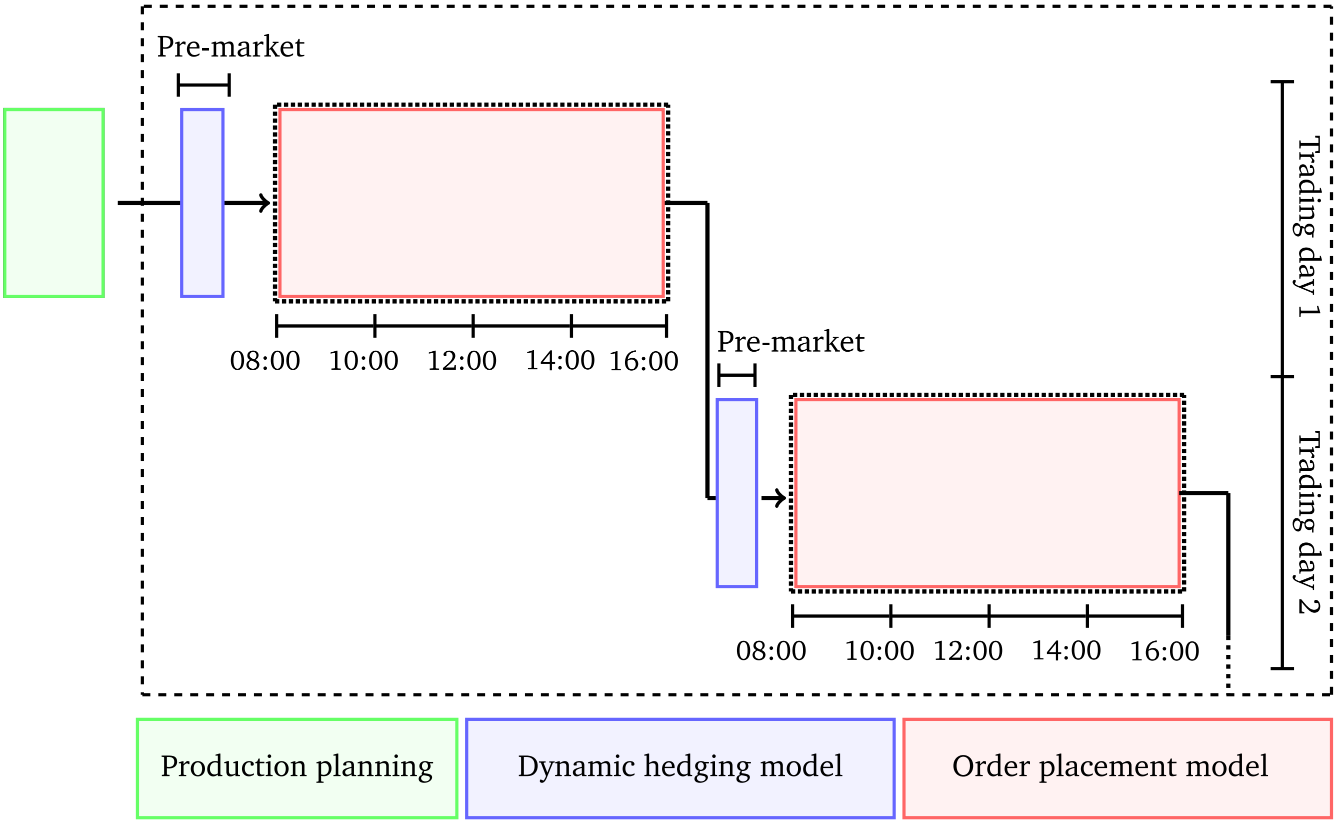

Figure 4.1: Integrated hedging and order execution process. The dynamic hedging plan is rebalanced every morning, pre-market, to decide the daily trading targets. Based on the trading results from the `Postponement model`, the dynamic hedging plan is rebalanced again the next morning.

4.4 Trading benchmarks

The `Integrated Postponement model` should be compared to other trading algorithms to verify whether it creates value for the trader in light of their risk preferences. The benchmark strategies have been selected such that the impact of each novel element in the `Integrated Postponement model` ought to be clear. The alternative algorithms are either well known heuristical trading strategies or existing models from the contemporary literature. The purpose of comparing performance is to see which elements of the `Integrated Postponement model` offer the most value to the trader.

We divide the alternative trading strategies into two subgroups. One subgroup of trading strategies is those which do not use the `QUASAR Dynamic Hedging model`. The process for determining the size of the daily volumes for these strategies is described in Appendix D. These strategies are referred to as static. The other subgroup consists of the strategies that utilise the `QUASAR Dynamic Hedging model`. We refer to these as dynamic hedging model (DHM) strategies. By comparing the subgroups, we can infer whether the `QUASAR Dynamic Hedging model` creates value for the trader. The strategies are listed below and summarised in table 4.1.

IOBE strategy: A single block order. The most basic trading strategy. The daily target volume, π_d^0, is placed as a single order in the opening hour. The daily volume target is statically determined.

IOBE DHM strategy: Similar to IOBE, but uses the `QUASAR Dynamic Hedging model` to determine the hedging targets.

TWAP strategy: A common trading strategy where the daily trading target is divided into orders of equal size, with hourly order placement. $x_t = \pi_d^0/|T|$. The daily volume target is statically determined.

TWAP DHM strategy: Similar to TWAP, but uses the `QUASAR Dynamic Hedging model` to determine the hedging targets.

Bertrand strategy: Identical to the `Postponement model`, but without the postponement option. Risk tolerance is modelled with a CVaR term in the objective function. This will be referred to as the Bertrand strategy as it is the strategy formulated in Bertrand (2021). The daily volume target is statically determined.

Bertrand DHM strategy: Similar to the Bertrand strategy, but uses the `QUASAR Dynamic Hedging model` to determine the hedging targets.

Table 4.1: A summary of the characteristics of the trading strategies. Note that strategies with order volume marked as NA place volumes dynamically, thus the volume is not determined before the current trading stage.

Strategy	Order volume	Dynamic hedging	Risk term	Postponement option
IOBE	π_d^0	✗	✗	✗
IOBE DHM	π_d^0	✓	✗	✗
TWAP	$\frac{\pi_d^0}{\lvert T\rvert}$	✗	✗	✗
TWAP DHM	$\frac{\pi_d^0}{\lvert T\rvert}$	✓	✗	✗
Bertrand	NA	✗	✓	✗
Bertrand DHM	NA	✓	✓	✗
`Postponement model`	NA	✓	✓	✓

Chapter 5

Data

To test the performance of the `Integrated Postponement model` we conduct a backtest. This chapter details the data sources that have been collected for this purpose. First, section 5.1 provides detail on the production plan and LOB data. Then, in section 5.2 we validate these data sources by commenting on data cleaning, data configuration and data quality.

5.1 Data sources

To use the `QUASAR Dynamic Hedging model`, a power production plan is required to determine an appropriate hedging plan. Furthermore, LOB data is needed to conduct the backtest, which is used to evaluate the trading performance of the `Integrated Postponement model`. Next, we describe the characteristics of these data.

5.1.1 The production plan

The production plan has been collected from one of Quantego's clients. Noise is added to distort sensitive information. The production profile is presented in figure 5.1. There is clear seasonality in the data. Notably, the production is lower during the winter even though this period has the highest electricity demand. This is a consequence of lower inflow. In addition to seasonal variation, the production deviates significantly in the short term as well.

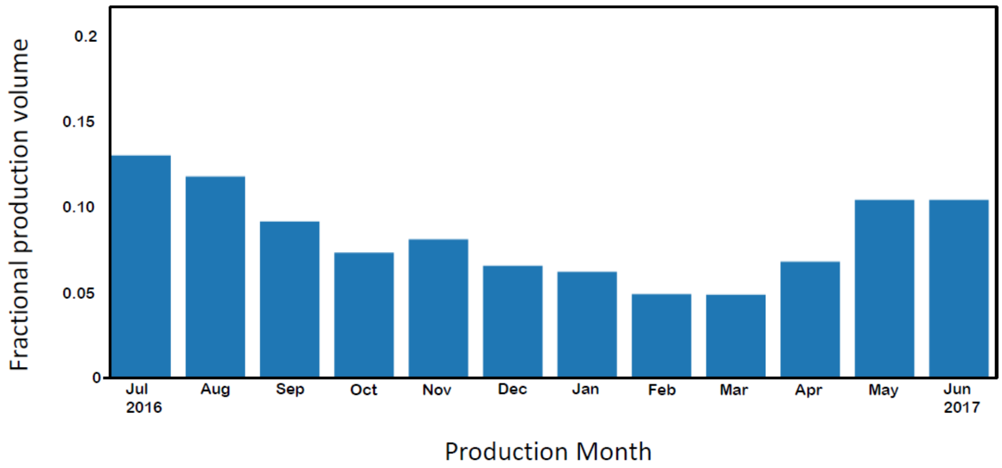

Figure 5.1: The reference production plan used by the `QUASAR Dynamic Hedging model`. The volumes are aggregated for each month and presented as fractions of the total power production over the period to distort any sensitive information. The volumes are notably lower during the winter due to low inflow to the reservoirs.

5.1.2 The LOB data

The LOB data has been collected from EEX, specifically for the German power futures market. The data set includes all monthly, quarterly and yearly baseload contracts for the period 2014 through 2017. The EEX data is structured as a set of LOB instances, presented in table 5.1. Every new data point is a snapshot of the LOB after an order has been placed, executed or removed.

Table 5.1: The data structure of the data points collected from the EEX exchange. Each data point consists of a contract name, timestamp and the best five bid and ask order levels.

Contract name	
Timestamp	
Best bid prices	$p_1^b, p_2^b, p_3^b, p_4^b, p_5^b$
Best bid volumes	$q_1^b, q_2^b, q_3^b, q_4^b, q_5^b$
Best ask prices	$p_1^a, p_2^a, p_3^a, p_4^a, p_5^a$
Best ask volumes	$q_1^a, q_2^a, q_3^a, q_4^a, q_5^a$

Data from 2015 is used to conduct a price driver study in an out-of-sample regressional analysis, with reference to section 4.2.2. The price driver study explained further in Appendix C. Data from 2014 is used for testing and validation of the price driver study. The data from June 2016 to June 2017 is used to conduct the trading backtest. Descriptive statistics for the data are presented in table 5.2. Noticeably, there are more orders for quarterly contracts than monthly contracts. Contracts with delivery during the winter also experience higher trading activity than contracts with delivery during the summer.

Table 5.2: Descriptive statistics for monthly and quarterly futures contracts tradeable between June 2016 and June 2017.

	JUL-2016	AUG-2016	SEP-2016	OCT-2016	NOV-2016
Number of orders	53 662	64 243	73 506	65 413	104 115
Mid Price(€)					
median	26.450	27.125	29.190	29.575	32.325
1st quantile	25.980	26.575	28.775	28.725	30.905
3rd quantile	27.015	27.865	29.685	30.100	33.550
Daily Volatility(%)					
Median	1.55	1.53	1.10	1.24	1.88
1st quantile	0.85	1.07	0.79	0.68	0.65
3rd quantile	2.55	2.15	1.58	1.95	1.95
Trading Price(€)					
Median	26.525	27.281	29.392	29.601	35.472
1st quantile	26.033	26.733	28.901	28.957	32.331
3rd quantile	27.201	28.051	29.873	30.491	37.777
Spread(%)					
Mean	0.23	0.35	0.33	0.40	0.61
1st quantile	0.18	0.18	0.17	0.17	0.17
3rd quantile	0.39	0.54	0.52	0.83	1.13
Bid depth					
Median	3.9	3.8	3.7	5.0	4.5
1st quantile	2.5	2.5	2.5	2.5	2.5
2nd quantile	5.0	5.0	6.5	8.9	8.6
Ask depth					
Median	2.5	2.5	2.5	2.5	4.0
1st quantile	2.5	2.5	2.5	2.5	2.5
2nd quantile	5.0	5.0	5.0	5.0	7.5

	DEC-2016	JAN-2017	FEB-2017	MAR-2017	APR-2017
Number of orders	138 808	92 749	155 517	166 373	107 105
Mid Price(€)					
median	30.065	38.690	38.125	34.500	31.210
1st quantile	29.460	33.940	36.135	32.490	29.950
3rd quantile	38.465	41.715	41.275	36.605	32.725
Daily Volatility(%)					
Median	1.21	1.84	2.21	1.67	1.22
1st quantile	0.28	0.48	0.88	0.91	0.62
3rd quantile	2.15	3.42	3.41	2.86	2.38
Trading Price(€)					
Median	39.081	38.952	40.484	36.255	31.403
1st quantile	36.051	36.403	37.777	34.075	29.605
3rd quantile	43.211	41.956	42.691	37.721	32.805
Spread(%)					
Mean	0.23	0.39	0.44	0.43	0.70
1st quantile	0.50	0.17	0.19	0.17	0.18
3rd quantile	1.01	0.85	1.00	1.01	1.76
Bid depth					
Median	4.1	5.0	2.5	4.2	5.0
1st quantile	2.5	2.5	5.0	2.33	2.5
2nd quantile	8.7	11.2	10.4	8.1	9.8
Ask depth					
Median	2.5	2.5	5.0	3.3	5.0
1st quantile	2.5	2.5	2.5	2.5	2.5
2nd quantile	5.0	12.5	12.5	5.0	7.5

	MAY-2017	Q3-2016	Q4-2016	Q1-2017
Number of orders	103 084	345 756	398 860	401 658
Mid Price(€)				
median	30.150	27.450	30.235	30.830
1st quantile	28.950	27.005	29.150	29.470
3rd quantile	31.170	28.055	30.820	36.075
Daily Volatility(%)				
Median	1.22	1.58	1.05	1.54
1st quantile	0.59	0.79	0.76	1.05
3rd quantile	2.01	2.69	1.64	2.74
Trading Price(€)				
Median	31.15	27.56	30.40	34.96
1st quantile	29.491	27.154	29.444	30.801
3rd quantile	31.821	28,211	31.000	38.601
Spread(%)				
Mean	0.56	0.18	0.17	0.32
1st quantile	0.14	0.17	0.13	0.16
3rd quantile	0.36	0.85	0.32	0.56
Bid depth				
Median	3.5	2.5	2.5	2.5
1st quantile	2.1	1.5	1.5	2.2
2nd quantile	7.5	4.0	4.5	4.35
Ask depth				
Median	3.9	2.5	2.5	2.5
1st quantile	2.5	1.5	1.5	2.0
2nd quantile	8.0	3.5	2.5	3.0

The price development for selected contracts is presented in figure 5.2. As seen, the price levels appear positively correlated. Contracts such as FEB-17 and Q1-17 share the characteristic of having delivery of power in February as part of, or as fully, the contract's underlying asset, justifying the contracts showing price correlation. Prices for the futures contracts in the data set demonstrate price appreciation as maturity approaches. The appreciation should be seen in the context of the price development before the trading period. Reduced fuel prices and high wind power production in the first quarter of 2016 put downward pressure on the spot price (AleaSoft, 2017). Subsequently, the spot price experienced mean reversion to the upside during the period of the backtest, which was also reflected in the power futures prices.

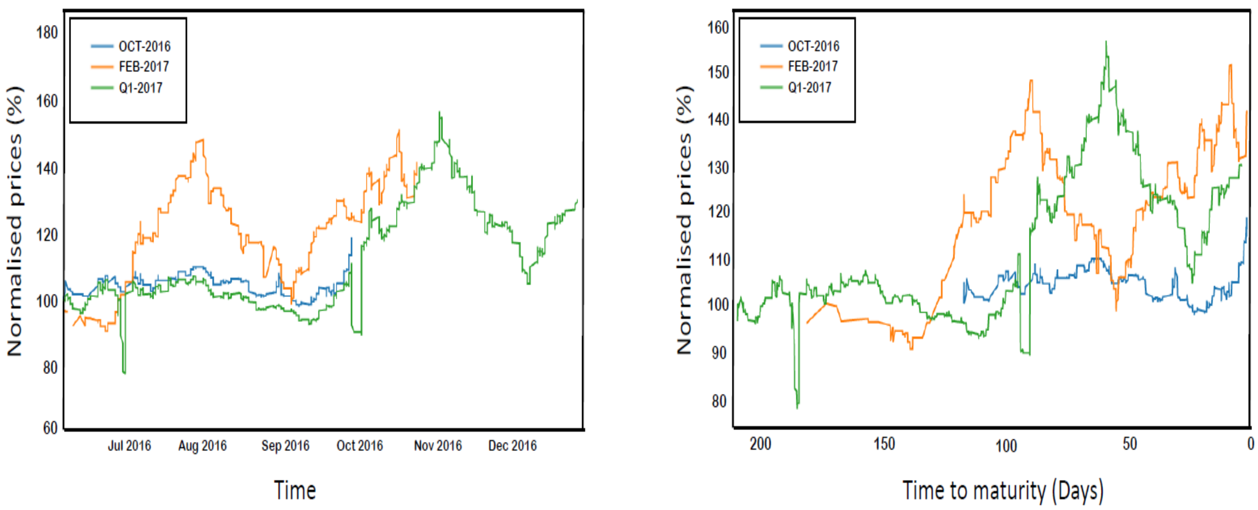

Figure 5.2: The normalised time series of price for three futures contracts as a function of time and time to maturity. There is positive correlation between the prices of each contract. Prices for the futures contracts in the data set demonstrate appreciation as maturity approaches.

5.2 Data validation

In this section, we comment on the central aspects of data validation. Specifically, how the data has been configured to conduct the backtest, the data cleaning process, and the data quality.

5.2.1 Data configuration

The data structure of the EEX LOB instances is not compatible with the microstructural backtesting framework used in this thesis, which requires single limit orders. To address this issue, the EEX LOB instances have been decomposed into a set of individual orders. For a detailed description of the transformational limit order process, the reader is referred to Appendix E.

5.2.2 Data cleaning

The raw LOB data has been put through a cleaning process before the LOB is deconstructed into individual orders. In particular, a large portion of the data points registered on the exchange before 07:30, and after 15:30 were found to be incomplete or empty. The incompleteness would correspond to a missing value for the characteristics that describe an order level, such as price or volume. It is speculated that these inaccuracies are related to the starting and stopping of order recording on the exchange. Therefore, data points with timestamps before 07:30 and after 15:30 were removed.

The order generation process described in Appendix E requires that the data points follow the LOB convention, where order levels are sorted by price level in ascending order for the ask side and descending order for the bid side. For some of the data points in the EEX data set, this convention is not upheld. These data points are then removed from the dataset. The reason for the occurrence of these faulty data points is unknown. However, discarding them is appropriate, seeing that these data points do not correspond to a legitimate state of the LOB.

5.2.3 Data quality

A few comments are made concerning the quality of the LOB data. First, the production plan is only given for a period of one year. In reality, power producers make hedging decisions with a longer planning horizon than one year. Another unfortunate consequence of the short horizon for the production plan is that yearly futures contracts are not traded, based on the hedging decisions of the `QUASAR Dynamic Hedging model`. A yearly contract implies the delivery of power for an entire year. However, since there are no full years of production decisions in the production plan, only six months for 2016 and six months for 2017, the `QUASAR Dynamic Hedging model` does not include yearly contracts in its hedging decisions. Yearly contracts are seen to be the most liquid contract type and would allow for a more favourable allocation of hedging volumes across contracts by evaluating contract liquidity.

Second, the raw EEX data only contains information about the five best bid and ask order levels. The remaining order levels on both sides of the LOB are not provided in the data. While the five best order levels may provide sufficient volume in the LOB for an incoming order to not completely clear the whole LOB, the fact that the other order levels are missing means that the actual state of the LOB is not reflected in the data. One could also argue that the absence of order levels for backtesting in an illiquid market is a bigger issue than for a liquid market.

Third, LOB data for the period of 2018 through 2020 appears to be incomplete. While a more recent period for the backtest than 2016 and 2017 would have been preferable, the incomplete LOB data means that conducting the backtest for 2016-2017 is the most appropriate choice.

Fourth, for the 2016 and 2017 period, electricity prices experienced price appreciation, demonstrating a structural trend in the price data. Ideally, the LOB data would have reflected multiple market trends, such that the inferences from the backtest are applicable regardless of the market context.

Chapter 6

Methodology

This chapter outlines the process of calibrating model parameters and introduces the backtesting methodology. First, section 6.1 describes the methods for calibration of the auxiliary model parameters. Then section 6.2 introduces the notion of backtesting and its use cases. Assumptions and rules governing the backtest are also described. Section 6.3 details the modelling of the LOB microstructure. Section 6.4 concludes the chapter with a description of how strategy performance is evaluated.

6.1 Calibrating auxiliary model parameters

The solution of the `Postponement model` is dependent on the values of the parameters that are used. This section introduces the methodology for calibrating two such parameters, namely the LOB slope and the scaling parameter for the postponement option.

6.1.1 Modelling the LOB bid slope

The LOB bid slope is modelled by the same method as Bertrand (2021). Each order level in the LOB is described by the pair of price and cumulative volume, Vol_i and price p_i. Each data point is then denoted as (Vol_i, p_i). The cumulative volume at level i is then the aggregated volume of all limit orders, q, with a price higher than p_i, plus half of the volume at price p_i.

$$Vol_i = \frac{q_i}{2} + \sum_{j \epsilon (p_j > p_i)} q_j \tag{6.1}$$

The pairs of price and cumulative volume are then used to conduct linear regression, described by equation (6.2). $\hat{\beta}_1$ is then the bid slope. The bid slope is re-calibrated before every trading stage. Every contract has its own bid slope. Approximating the shape of the LOB as linear, as opposed to the discrete nature of a LOB, is necessary due to high computational complexity. Using non-continuous parameters will result in a discrete solution space. To solve such problems, the solver needs to test all possible solutions or use algorithms such as branch and bound, increasing computational complexity.

$$p_i^{bid} = \hat{\beta}_0 + \hat{\beta}_1 \cdot Vol_i \tag{6.2}$$

6.1.2 Calibrating the shadow price parameter ζ

The postponement optionality term in the `Postponement model` contains of a scaling parameter ζ. This parameter is calibrated to improve the `Postponement model` trading performance.

In order to calibrate the postponement parameter ζ, an out-of-sample test for 2015 is conducted where several values of ζ are tested. By testing different values of ζ, we notice that monthly and quarterly contracts behave differently. Therefore we decide to separate between ζ for monthly and quarterly contacts. We define ζ as a function of delivery duration (DD) and a calibration variable γ, as expressed in equation (6.3).

$$\zeta = \frac{-1}{\mathrm{DD} \cdot \gamma} \tag{6.3}$$

By using a value of γ of 1.37 for monthly contracts, and 1.57 for quarterly contracts, revenues are maximised over the short period of the out-of-sample test.

6.2 Backtesting

Backtesting is a well-established methodology in finance, with several applications. These include determining the accuracy of VaR models and simulating the performance of trading strategies using historical data. Questions have been raised regarding the validity of backtesting results. Harvey and Liu (2015) argue that using historical data limits the generalisation of the results and inferences. Despite this potential drawback, backtesting is prevalent among brokers and investors. Matras (2011) makes the point that backtesting is useful for evaluating trading strategy performance. In this thesis, backtesting is used to compare the performance of the `Postponement model` with benchmark strategies. This section will describe the backtesting framework that has been used in this thesis, which includes assumptions, incrementation, order placement and model implementation.

6.2.1 Backtest assumptions

To conduct a case study with a realistic application, a set of underlying assumptions should be in place to reflect the limitations of the model and methodology. Inferences made out ought to be seen in the light of these assumptions. The most significant assumptions are outlined and discussed below.

Assumption 1 - Order placement with complete information: Orders carry information about a trader's market view. Consequently, one trader's order can affect the trading behaviour of other traders. In a traditional market setting, traders place orders with full knowledge of the state of the LOB. A distinction is made between exogenous and endogenous orders. In the context of a backtest, exogenous orders are those orders retrieved from the data set, while endogenous orders are those placed in the LOB by the backtest user. Endogenous orders are placed with full knowledge regarding the state of the LOB. The same characteristic does not hold for historical exogenous orders. For backtesting, exogenous orders are assumed to be placed in the LOB with knowledge of the LOB state, including all endogenous orders.

Assumption 2 - Event-driven simulation: The LOB clock is incremented discretely by the sequence of exogenous orders instead of continuous clock incrementation. This sequential starting and stopping allows a trader to interpret the LOB microstructure and calculate market statistics between simulation increments. As a result, the trader receives an advantage from acting before other market participants that does not reflect reality. With hourly trading, there is sufficient time to make inferences about the state of the market, such that the advantage is assumed to be insignificant for trading performance.

6.2.2 Trading frequency

The backtest is conducted on EEX data, where orders arrive between 07:30 to 15:30 on trading days. Trading will therefore occur during these hours in the backtest. For all strategies except the IOBE strategies, orders are placed once every hour, i.e., between 7:30-8 and 8-9 until 15:30. There will therefore be a maximum of 9 order placements each day for every contract.

6.2.3 Backtesting sequence

The `QUASAR Dynamic Hedging model` requires the accumulated short position of every tradeable contract to solve the optimisation problem for the following trading period. Therefore, it is necessary to simulate order placement for all contracts collectively on each trading day. As the `QUASAR Dynamic Hedging model` rebalances the hedging plan daily, the accumulated short positions in the power futures contracts are updated after each trading day. After the backtest has iterated through all orders for the trading day, the accumulated short positions are updated and used as input in the `QUASAR Dynamic Hedging model` to rebalance the hedging plan for the subsequent trading day. This process is described in algorithm 3 in Appendix E.

6.2.4 Execution price

If a market order clears the LOB depth, the order is filled by multiple order levels in the LOB. The execution price must then be calculated. The execution price is calculated as the volume-weighted average price of the limit order prices that the market order clears.

6.2.5 Endogenous order placement

Endogenous orders are generated and added to the LOB. Due to the incremental sequence of the simulation, endogenous orders are placed between exogenous orders entering the LOB. The frequency and size of these orders will depend on the order execution strategy. The IOBE strategies place one order with the total daily trading target π_d^0 after the first exogenous order of the day arrives in the LOB. The TWAP strategies place orders of equal volumes after the first exogenous order arrival every hour. Suppose there is no exogenous order arrival during a trading hour. In that case, the trading volume for the next hour is adjusted such that the outstanding volume that was scheduled for the prior trading hours is included in the current trade. I.e., if there has been one hour without trading, the order volume for the subsequent order is doubled. If there have been two hours without exogenous order arrival, the order volume is tripled and so on. For the dynamic programming strategies, orders are placed with the same frequency as the TWAP strategy, but instead, the order volumes are determined by solving their respective optimisation problems.

6.2.6 Implementing the Integrated Postponement model

The `Postponement model` is solved by using the stochastic optimisation software QUASAR. Since the solver requires relatively complete recourse, artificial boundaries and variables are included in the model to provide solving stability. The `Postponement model` which includes artificial boundaries and variables is presented in Appendix F.

As of model implementation, the backtest algorithm recognises when there is a new hour. The `Postponement model` is then initialised. It uses the best bid at that time as the input value for price p_t, as well as the previous stage value for volume left to trade, I_{t-1}. The `Postponement model` then returns the first-stage trading decision x_t as well as the value left to trade at the next stage, I_t. The first-stage trading decisions are then submitted to the LOB as endogenous orders.

6.3 Modelling the LOB microstructure

In conventional backtests, the execution price of an endogenous order is set to the market price at the point in time where the order was placed, regardless of the order volume. One weakness with this method is that the LOB microstructure is not considered. As a result, conventional backtests do not consider temporary or permanent price impact. The novelty of the backtest conducted in this thesis relates to the implementation of a LOB microstructure model. In a conventional backtest, the market state is described solely by the best bid and ask prices. Instead, this thesis employs the collection of all limit orders in the LOB to describe the market state. The following section describes the modelling of the LOB microstructure used as the market's state variable.

Cont et al. (2010) was the first to model the multiple price levels of the LOB as a multiclass queueing system. Zheng (2016) models the LOB using a pseudo-continuous[1] price grid for fluid, dynamic limit order arrival. In contrast to Zheng (2016)'s LOB microstructure model, we discretise the price grid to reflect all price levels in the LOB. Additionally, limit order arrival is treated as a discrete arrival process.

Price levels: For the ask side of the LOB we consider prices $p_i^a(t)$ indexed by $i \in I$, at time t, where prices are in ascending order such that $p_1^a(t) < p_2^a(t) < \ldots < p_N^a(t)$. Index i corresponds to the ith price level of the multilevel ask order queue, where I is denoted as the set all ask order

[1]The continuity is restricted due to minimal tick size on the exchange.

levels. Let $P^a(t)$ denote the set of all ask order prices at time t, such that

$$p_i^a(t) \in P^a(t) \quad \text{for all} \quad i \in I.$$

Prices $p_j^b(t)$ indexed by $j \in J$, for the bid side follows the same argument but where $p_1^b(t) > p_2^b(t) > \ldots > p_M^b(t)$. Index j corresponds to the jth price level of the multilevel bid order queue, where J is denoted as the set all bid order levels. Like for the ask order queue, let $P^b(t)$ denote the set of all bid order prices at time t, such that

$$p_j^b(t) \in P^b(t) \quad \text{for all} \quad j \in J.$$

Because bid and ask orders are matched and removed from the LOB if $p_1^b(t) > p_1^a(t)$, it follows that $p_1^b(t) < p_1^a(t)$ for every state of the LOB.

Order queues: For index $i \in I$ at time t, there is a quantity of MW available for purchase, denoted by $q_i^a(t)$ at price $p_i^a(t)$. By the same token, for index $j \in J$, there is a quantity of MW available for sale, denoted by $q_i^b(t)$ at price $p_i^b(t)$. $q_1^a(t)$ is defined as the ask order depth. Similarly, $q_1^b(t)$ is defined as the bid order depth. The state of the LOB at time t, is described by the multilevel bid and ask order queues, $Q^a(t)$ and $Q^b(t)$, formulated in equation (6.4) and (6.5) respectively.

$$Q^a(t) \triangleq \{(q_1^a(t), p_1^a(t)), \ldots, (q_{N^a}^a(t), p_{N^a}^a(t))\} \tag{6.4}$$

$$Q^b(t) \triangleq \{(q_1^b(t), p_1^b(t)), \ldots, (q_{N^b}^b(t), p_{N^b}^b(t))\} \tag{6.5}$$

Here, N^a and N^b denotes the number of order levels in the multilevel bid and ask order queues, respectively.

Limit order arrival: Upon arrival, limit orders are placed in the LOB. If the order price already exists in the price grid, the order volume is added to the existing order queue. Otherwise, a new order queue is created for the order price. Let R^l and S^l be the sets of all limit ask and bid orders, respectively. Now let L_k^a be the kth limit ask order, and L_o^b be the oth limit bid order where $L_k^a \in R^l$ and $L_o^b \in S^l$. Limit orders are defined by their quantity, price and time of placement. Thus we have

$$L_k^a \triangleq \{\hat{q}_k^a, \hat{p}_k^a, \hat{t}_k\} \quad \text{for all} \quad L_k^a \in R^l. \tag{6.6}$$

$$L_o^b \triangleq \{\hat{q}_o^b, \hat{p}_o^b, \hat{t}_o\} \quad \text{for all} \quad L_o^b \in S^l. \tag{6.7}$$

Considering the arrival of order L_k^a to the multilevel ask order queue, at the time of its placement, $Q^a(\hat{t}_k)$, if $\hat{p}_k^a \in P^a(\hat{t}_k)$, $\hat{q}_k^a$ is added to the volume of the existing order queue with the corresponding price level $\hat{p}_k^a$. Otherwise, a new order queue $q_i^a(t)$ is created, where the index i is given by the location of $\hat{p}_k^a$ in the price grid:

$$p_{i-1}^a(\hat{t}_k) < \hat{p}_k^a < p_i^a(\hat{t}_k)$$

The index of all order queues in $Q^a(t)$ beyond i, will then be incremented by one.

Figure 6.1 illustrates a snapshot of a LOB before and after the arrival of an ask limit order. Because the price of the new limit order, $\hat{p}_k^a$, is better than the previous best ask price, p_1^a, the new limit order is now the best ask order. The indices describing the position of existing price levels in the LOB are now shifted to reflect the new order entry. The old p_1^a is now p_2^a, and so on.

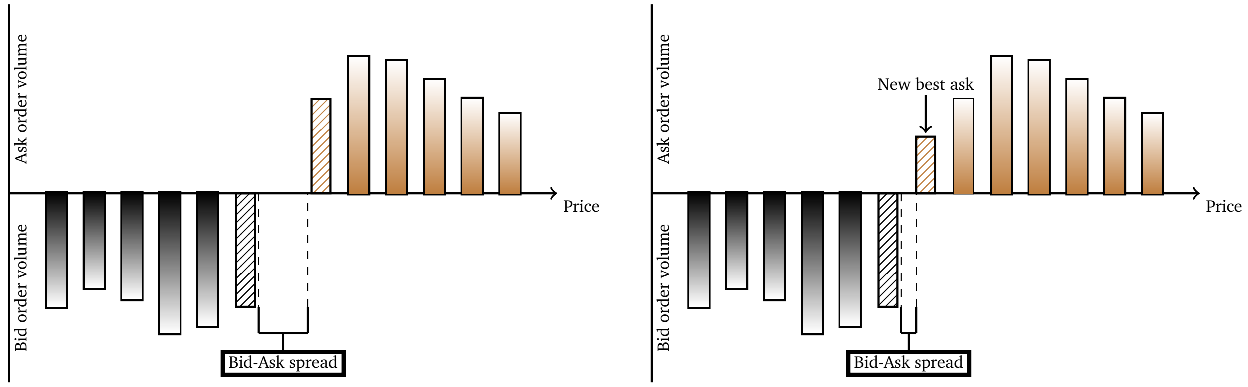

Figure 6.1: A new best ask order arrives in the LOB. The bid-ask spread and the ask depth are reduced.

Market order arrival: Market orders do not enter the LOB upon arrival. Instead, they fill existing limit orders. Let R^m and S^m denote the sets of all market ask and bid orders, respectively. Now let M_v^a be the vth market ask order, and M_u^b be the uth market bid order, where $M_v^a \in R^m$ and $M_u^b \in S^m$. Market orders are denoted only by their quantity and time of placement. Thus, we have

$$M_v^a \triangleq \{\tilde{q}_v^a, \tilde{t}_v\} \quad \text{for all} \quad M_v^a \in R^m. \tag{6.8}$$

$$M_u^b \triangleq \{\tilde{q}_u^b, \tilde{t}_u\} \quad \text{for all} \quad M_u^b \in S^m. \tag{6.9}$$

Consider the arrival of M_v^a at its time of placement $\tilde{t}_v$. Let N^j denote the set of volumes for the j first bid order queues, $\{q_1^b, \dots, q_j^b\}$. Now let N^{j-1} be the set of volumes for the $j-1$ first bid order queues, $\{q_1^b, \dots, q_{j-1}^b\}$. If j is given as the integer for which the following inequality holds:

$$\sum_{j \in N^j} q_j^b(t) \geq \tilde{q}_v^a \geq \sum_{j \in N^{j-1}} q_j^b(t)$$

then all limit bid orders in the first $j-1$ order queues are fully filled and removed from the LOB. If the following equality holds,

$$\sum_{j \in N^j} q_i^b(t) - \tilde{q}_j^a = 0$$

then the jth bid order queue is fully filled as well. Otherwise the volume for the jth bid order queue is adjusted to

$$\sum_{j \in N^j} q_j^b(t) - \tilde{q}_v^a.$$

Limit order cancellation: Limit orders can be cancelled and removed. Upon cancellation, the existing order is removed from the order queue corresponding to its price level. If the price or quantity of an observable limit order is modified, the change is treated as a cancellation and then re-entry of the modified order.

6.4 Quantifying strategy performance

To the best of the authors' knowledge, there is no established framework for liquidation strategy comparison as most contemporary literature solve the problem analytically. Sharpe ratio is commonly used to evaluate trading strategy performance in conventional backtest. However, given that the Sharpe ratio is computed from returns, it is not an appropriate metric for a liquidation strategy. Instead, this section introduces other metrics that may be used to compare trading strategies. Trading strategy performance is evaluated using these four metrics; average revenue per MWh, standard deviation of revenue per MWh, relative performance and price impact.

6.4.1 Revenue per MWh

Revenue per MWh is a suitable measurement of the trading performance of an order execution strategy for a power producer, as it reflects the total cash flows from the power derivatives trading. Since power futures contracts have different delivery durations, the revenues accrued from a contract is proportional to its duration.

Let trades be denoted as $i_c \in I_c$ where i_c is a single trade for contract c and I_c is the set of all trades for contract $c \in C$. Let p_{i_c} be the execution price of trade i_c for contract c, w_{i_c} the trading volume of trade i_c, and d_c is the delivery duration of contract c.

Let Ω denote a set of contracts. Revenue per MWh, $\bar{p}_\Omega$, is then calculated using equation (6.10). Notice that the product of the trading volume, w_{i_c} and contract delivery duration, d_c is used as the weights for trade i_c. Applying the weighting method is done because not all trades are of equal volume. Additionally, their accrued revenues are different due to varying delivery duration.

$$\bar{p}_\Omega = \frac{\sum\limits_{c\in\Omega} \sum\limits_{i_c \in I_c} d_c \cdot w_{i_c} \cdot p_{i_c}}{\sum\limits_{c\in\Omega} W_{1c}} \tag{6.10}$$

where

$$W_{1c} = \sum_{i_c \in I_c} d_c \cdot w_{i_c} \tag{6.11}$$

6.4.2 Standard deviation of revenues per MWh

The standard deviation of revenues per MWh measures the variation in execution prices of trades i_c. A risk-averse trader seeks to decrease the standard deviation of revenue per MWh to reduce cash flow uncertainty. The standard deviation for contract c, s_c, is calculated using reliability weighted sample variance, with volumes w_{i_c} as weights. The standard deviation for a single contract is calculated by using equation (6.12). Here, $\bar{p}_c$ is the revenue per MWh for contract c calculated using equation (6.10).

$$s_c^2 = \frac{\sum_{i_c \in I_c} d_c \cdot w_{i_c}(p_{i_c} - \bar{p}_c)^2}{W_{1c} - \frac{W_{2c}}{W_{1c}}} \tag{6.12}$$

where

$$W_{2c} = \sum_{i_c \in I_c} w_{i_c}^2 \tag{6.13}$$

The standard deviations are then converted from hourly to daily for the strategies with hourly order placement. Standard deviations are presented as percentages of the respective average revenues per MWh.

The standard deviation for a set of contracts, Ω, is also calculated using reliability weighted sample variance, with the total number of MWh traded for contract c, W_{1c}, as its weights. Contract standard deviation, s_c relative to the average revenue per MWh for contract c, $\bar{p}_c$, is used in the calculation of standard deviation for Ω. The standard deviation for a set of contracts is formulated in equation (6.14). This approach is used to calculate the standard deviations of the following sets of contracts: monthly contracts, quarterly contracts and all traded contracts.

$$s_\Omega^2 = \frac{\sum_{c \in \Omega} W_{1c} \cdot (\frac{s_c}{\bar{p}_c})^2}{Z_1 - \frac{Z_2}{Z_1}} \tag{6.14}$$

where

$$Z_1 = \sum_{c \in \Omega} W_{1c} \tag{6.15}$$

and

$$Z_2 = \sum_{c \in \Omega} W_{1c}^2 \tag{6.16}$$

6.4.3 Relative performance

Another metric is introduced, which will be referred to as the relative performance (RP). Relative performance measures how a strategy performs relative to the IOBE strategy by comparing revenue per MWh. Equation (6.17) illustrates how relative performance is calculated.

$$RP_{Strategy} = \frac{Rev_{Strategy} - Rev_{IOBE}}{Rev_{IOBE}} \tag{6.17}$$

6.4.4 Price impact

An important aspect of this thesis is the influence of price impact on trading performance. Price impact is calculated as the relative difference between the execution price of trade i_c, p_{i_c} and the best bid prior to order placement, $p^b_{1c}(t)$, for contract c. The product of the trading volume w_{i_c} and the delivery duration, d_{i_c} for all trades i_c are used as weights. These weights are used to consider delivery duration. The average price impact for a set of contracts Ω is calculated using equation (4.4). As for the standard deviation of revenue per MWh, the average price impact for the sets of monthly, quarterly and all contracts are presented and discussed in the next chapter.

$$PI_\Omega = \frac{\sum_{c\in\Omega}\sum_{i_c\in I_c} d_c \cdot (1 - \frac{p_{i_c}}{p^b_{1cm}}) \cdot w_{c_i}}{\sum_{c\in\Omega} d_c \cdot w^{total}_c} \tag{6.18}$$

where

$$w^{total}_c = \sum_{i_c \epsilon I_c} w_{i_c} \tag{6.19}$$

Chapter 7

Results

The `Integrated Postponement model` from section 4.2 and the alternative trading strategies from section 4.4 have been backtested using the microstructural LOB framework introduced in section 6.3. This chapter gives emphasis to the `Postponement model` and its decision policies. The motivation behind this area of focus is the novel postponement optionality and the model's sensitivity to changing market liquidity on an hourly basis. Additionally, IOBE and TWAP are static strategies with no reactivity to changing market liquidity on an hourly basis and will serve as trading performance benchmarks.

An Intel Xeon 2.1 GHz processor with 256 GB RAM was used to run the backtest. The runtimes of the backtest for each strategy are presented in table 7.1.

Table 7.1: Runtimes for the trading strategies.

Strategy	Runtime [minutes]
IOBE	148
IOBE DHM	489
TWAP	163
TWAP DHM	501
Bertrand	836
Bertrand DHM	1259
`Postponement model`	1401

Runtime is seen to increase with model complexity. Using the `QUASAR Dynamic Hedging model` increases runtime by 367 minutes on average. Bertrand DHM and the `Postponement model` naturally have the longest runtimes, as these incorporate two rolling horizon stochastic optimization problems on each trading day.

This chapter is divided into four sections. First, we present the trading performance of the strategies in section 7.1. In section 7.2 we present the decision policies specific to the `Postponement model`, namely trading trajectories, and postponement decisions, in the context of changing market liquidity. Next, we interpret and discuss trading performance results in section 7.3. Section 7.4 concludes the chapter with interpretation and discussion of the `Postponement model` decision policies.

7.1 Trading strategy performance

Trading strategy performance is presented next with reference to revenue per MWh and price impact.

7.1.1 Revenue per MWh

Revenue per MWh (€/MWh), standard deviations and relative performances (RP) have been computed using the methods described in section 6.4. The results are presented in table 7.2. Results have been aggregated across contracts of the specific delivery duration, i.e. monthly and quarterly contracts. For a comprehensive review of the results on a per contract basis, the reader is referred to Appendix G.

Table 7.2: Revenue per MWh, standard deviation and relative performances for the trading strategies. The Postponement model has the highest revenue per MWh, followed by the Bertrand DHM strategy. Bertrand DHM has the lowest standard deviation of revenue per MWh.

Strategy	Monthly		Quarterly		Total		
	€/MWh	St.dev	€/MWh	St.dev	€/MWh	St.dev	RP (bps)
IOBE	31.560	1.37%	31.320	1,38%	31.404	1,38%	-
IOBE DHM	31.986	1.30%	31.187	1.12%	31.516	1.13%	35.7
TWAP	31.664	1.50%	31.714	0.62%	31.697	1.02%	93.3
TWAP DHM	32.101	1.44%	31.654	1.31%	31.838	1,25%	138.2
Bertrand	31.775	0.48%	31.731	0.48%	31.739	0.48%	106.7
Bertrand DHM	32.220	0.57%	31.612	0.48%	31.854	0.49%	143.3
`Postponement model`	32.350	0.37%	31.994	0.82%	32.134	0.68%	232.5

Revenue per MWh is seen to increase by allocating trading volume across trades on a daily basis. Strategies employing the `QUASAR Dynamic Hedging model` outperform their equivalent strategy without the `QUASAR Dynamic Hedging model` by 39 bps on average. The Bertrand strategies experience the lowest daily standard deviation of 0.48-0.49%, followed by the `Integrated Postponement model` with 0.68%.

7.1.2 Daily price impact

The volume-weighted price impact for each strategy has been calculated using the methods described in section 6.4. The price impact results are presented in table 7.3.

Table 7.3: Average relative price impact, measured in basis points. The single-order strategies experience considerably higher price impact than the strategies placing multiple orders through the day. The strategies placing multiple daily orders experience similar price impacts, where the Bertrand model experiences slightly lower price impact than the TWAP strategy.

Strategy	Price impact (bps)
IOBE	87.7
IOBE DHM	101.1
TWAP	5.05
TWAP DHM	7.46
Bertrand	4.37
Bertrand DHM	7.44
`Postponement model`	7.09

Price impact is reduced significantly by moving from daily to hourly trading. Among the strategies that place orders every hour, the strategies without the `QUASAR Dynamic Hedging model` experience the lowest price impact. For the strategies using the `QUASAR Dynamic Hedging model`, the `Integrated Postponement model` has the lowest price impact.

7.2 Postponement model decision policies

In this section, the trading decisions obtained for the `Integrated Postponement model` are presented, specifically the trading trajectory and postponement decisions. These are highlighted due to their sensitivity to changing market liquidity.

7.2.1 Trading trajectory throughout the day

Figure 7.1 presents the trading trajectories throughout the day for the `Integrated Postponement model`. Each curve represents one particular realization of a daily trajectory. The colour of each curve is given by its time to maturity. Green curves correspond to trajectories close to maturity, and blue curves correspond to trajectories far from maturity.

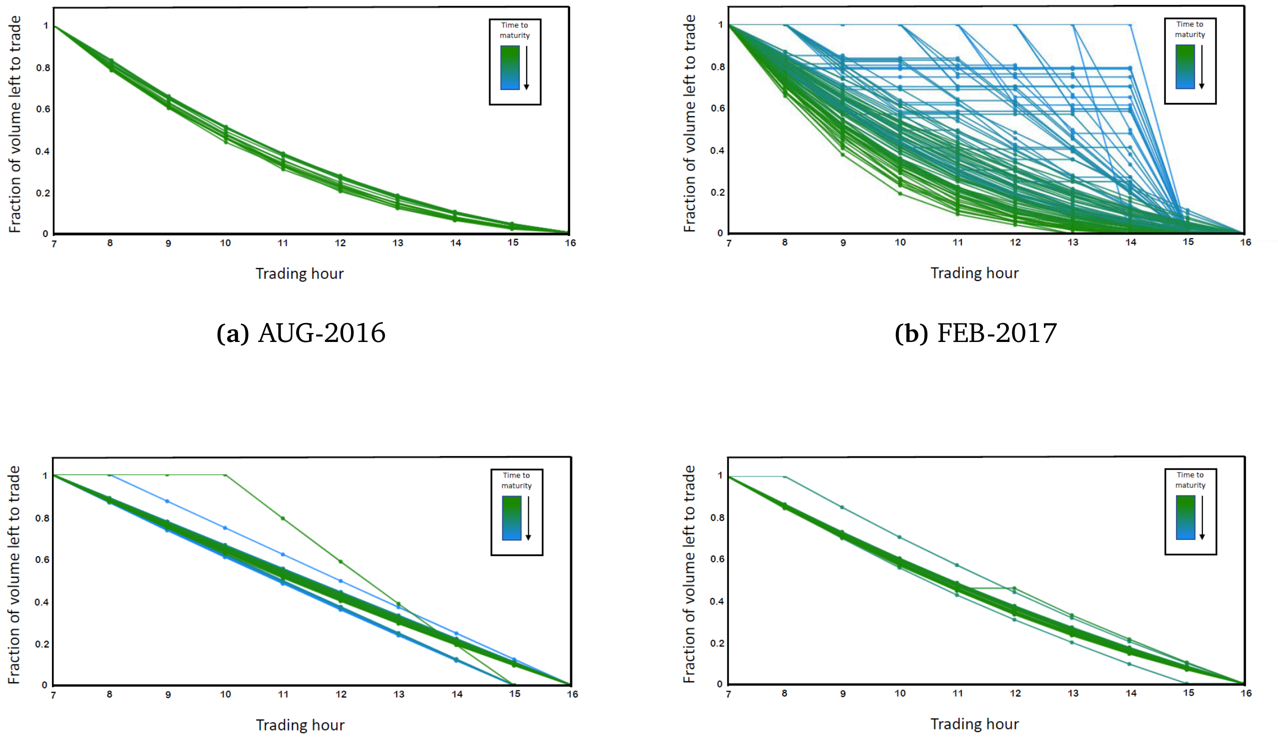

(a) AUG-2016

(b) FEB-2017

(c) Q1-2017

(d) Q4-2016

Figure 7.1: Trading trajectories for selected contracts. The top panels show monthly contracts, while the bottom panels show quarterly contracts. The contracts in the right hand side panels are more liquid than the contracts in the left hand side panels. The liquid contracts experience a greater degree of convexity as maturity approaches than the less liquid contracts.

As seen from figure 7.1, the trading trajectories exhibit a greater degree of convexity as maturity approaches. This distinction is more clear for monthly rather than quarterly contracts. Additionally, one can see that order placement occasionally starts at later trading stages on days further from maturity. Since order arrival occurs discretely, the limited order arrival on illiquid trading days prevents the `Integrated Postponement model` from placing orders until the first exogenous order arrives.

7.2.2 Postponement decisions

Figure 7.2 shows the fraction of daily postponed volume in conjunction with the ratio of best bid price, p_1^b, to price-equivalent shadow price (p_1^b/p_{SP}), towards contract maturity. Here the postponement fraction is plotted in red, with its vertical axis on the left side of the figure. The ratio of the best bid price to price-equivalent shadow price is shown as the black plot, with its values denoted on the right-hand vertical axis.

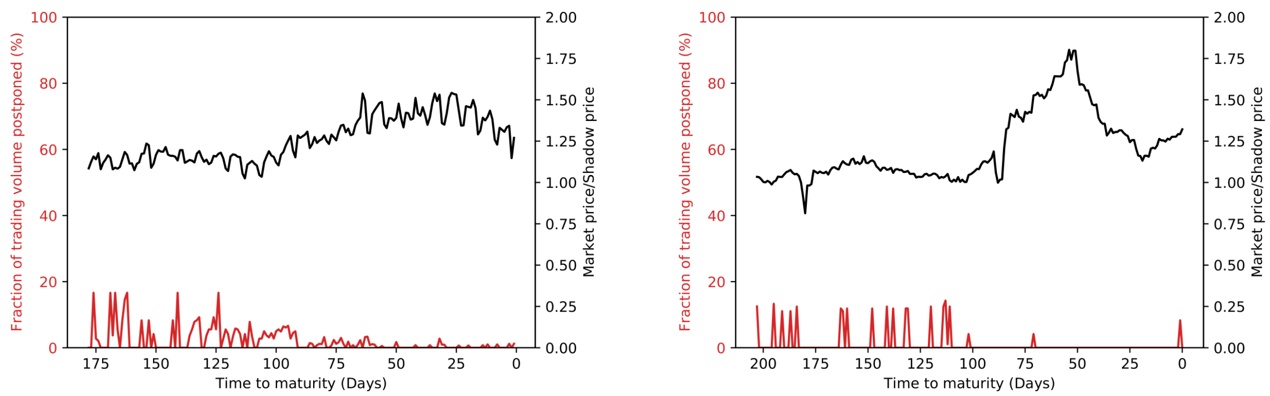

(a) Monthly contracts

(b) Quarterly contracts

Figure 7.2: Postponement decisions of the `Postponement model` as a function of time to maturity. The red plot illustrates the fraction of daily volume postponed, while the black plot illustrates the ratio between the best bid and the price-equivalent shadow price. The fraction of postponed volume decreases towards maturity. This trend is seen in light of the increasing p_1^b/p_{SP} value towards maturity.

The ratio p_1^b/p_{SP} is seen to increase as contract maturity approaches, decreasing the propensity for the `Integrated model` to postpone. As seen in the figure, the frequency of days where postponement occurs decreases towards maturity. Additionally, the fraction of postponed trading volume decreases as contracts approach maturity.

7.3 Discussion of trading strategy performance

To assess the performance of the different trading strategies, some key metrics are calculated and compared across strategies. In section 7.3.1, the obtained revenue per MWh and the experienced price impact are discussed. The effect of different model features on the standard deviation of revenues are interpreted and discussed in section 7.3.2. Section 7.3.3 discusses the differences in performance between strategies with and without the `QUASAR Dynamic Hedging model` to assess its value.

7.3.1 Revenue per MWh and price impact

Significant value can be accrued by moving from daily to hourly trading. By placing market orders with lower volumes, price impact is reduced by as much as 96 bps as seen in table 7.3. These results are consistent with the findings of Predoiu et al. (2011), highlighting the relationship between order size and price impact.

In contrast, the Bertrand model experiences less price impact than the TWAP strategy. Based on Almgren and Chriss (1999), we would expect the opposite as the model accepts additional price impact to reduce price risk. However, since the difference is slight, with a difference of 0.68 basis points between the strategies that do not use dynamic hedging and 0.02 basis points between the strategies that use dynamic hedging, the difference in price impact may not be statistically significant.

Additionally, the lower revenue per MWh could also result from the trading trajectories and price movements during the daily trading period. If the price level is consistently higher during the hours the Bertrand model allocates larger volumes than TWAP, it can obtain a higher average revenue per MWh. Upon inspection of the price development on trading days, prices tend to depreciate through the trading period, which explains the Bertrand model outperformance. Considering that the Bertrand model possibly outperforms TWAP due to an exogenous factor that is not part of the model formulation, one should not conclude that the Bertrand model is a better order execution model for the risk-neutral trader.

When the Bertrand model is extended to allow for postponement, i.e. the `Integrated Postponement model`, average revenue per MWh is increased by 80 bps. By introducing the price-equivalent shadow price as a threshold for when prices are considered favourable, the `Integrated Postponement model` can be more selective in its trading decisions. Given that the postponement optionality comes at no expense, one would expect better performance than a model that has not been giving this costless optionality. Considering that the price-equivalent shadow price is an objective metric of the value to wait, the trading strategy could be said to provide a favourable trade-off between risk and reward.

7.3.2 Revenue per MWh standard deviation

Moving from daily to hourly trading not only increases revenue per MWh, but also reduces the standard deviation of the revenues per MWh. By introducing trader risk aversion in the Bertrand model, the variance is reduced by nearly half, compared to the risk-neutral TWAP strategy. When the trading trajectory is convex, which corresponds to risk-averse preferences, a larger fraction of the total trading volume will have been executed earlier in the day than will be the case for a risk-neutral trading trajectory. Since a more significant proportion of the trading volume has already been executed during a short period, a smaller proportion of the trading volume is exposed to future price volatility. Thus, the total standard deviation is decreased.

The `Integrated Postponement model` has a higher standard deviation than the equivalent model without the postponement option, i.e. the Bertrand model. By introducing daily postponement optionality, the trading trajectory reflects a risk preference somewhere between the risk neutrality of TWAP DHM and the risk aversion of the Bertrand DHM strategy, which gives it a standard deviation value between those strategies.

7.3.3 Value of QUASAR Dynamic Hedging model

The `QUASAR Dynamic Hedging model` improves revenue per MWh by 39 bps when comparing strategies with and without the `QUASAR Dynamic Hedging model`. Interestingly, strategies with the `QUASAR Dynamic Hedging model` tend to outperform for monthly contracts but perform worse for quarterly contracts. One possible explanation for this tendency could be that the `QUASAR Dynamic Hedging model` is able to evaluate the trade-off between trading monthly and quarterly contracts. By sacrificing the performance of one contract type, the model can outperform in the aggregate.

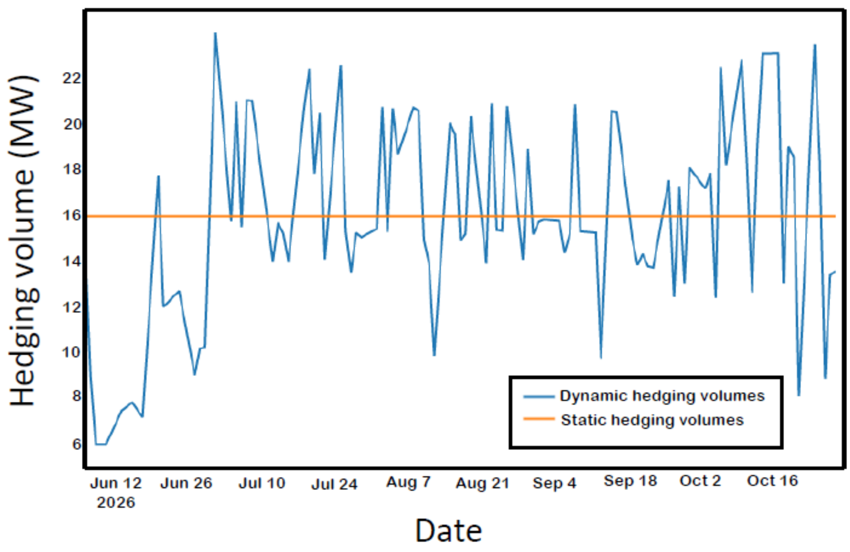

Figure 7.3: The hedging plan for the NOV-2016 contract for the `Integrated Postponement model`, and the static volume strategies. Note that the trading volumes for static volume strategies are evenly distributed across the contract lifecycle. The DHM strategies have lower trading volumes far from maturity, indicating a back-loaded trading volume allocation. Since the price of the NOV-2016 contract appreciates towards maturity, the `QUASAR Dynamic Hedging model` achieves higher revenues than the static model.

A different explanation for the outperformance relates to the trading volume allocation seen in the context of the price data trend. The `QUASAR Dynamic Hedging model` allocates trading volumes based on market liquidity, preferring to allocate larger volumes in periods of high liquidity. In the case of electricity futures which are more liquid close to maturity, this tendency leads to backloading throughout the contract lifecycle, seen in figure 7.3. Figure 7.4 illustrates the price development of monthly contracts where the `QUASAR Dynamic Hedging model` outperforms, and falls short, respectively. A common factor for the contracts showing outperformance is the trend of substantial price appreciation as maturity approaches. Contracts where the model falls short experience less price appreciation. For these contracts, the increased price impact experienced due to higher volume allocation results in lower revenue per MWh even for the contracts that slightly appreciate. Thus, it appears that the outperformance can be partly attributed to the favourable allocation of trading volume, seen in the light of price trends.

7.4 Discussion of Postponement model decision policies

The `Integrated Postponement model` makes trading decisions based on the market conditions at the time of order placement. By taking the liquidity of the LOB into account, better-informed decisions on the allocation of trading volumes can be reached. This dynamic allocation of volume is reflected in the trading trajectories. Furthermore, the postponement optionality offers unique flexibility of postponing trading when the market conditions are deemed unfavourable. In this section, the `Integrated Postponement model` decision policies will be discussed in light of changing market conditions (liquidity and time to maturity).

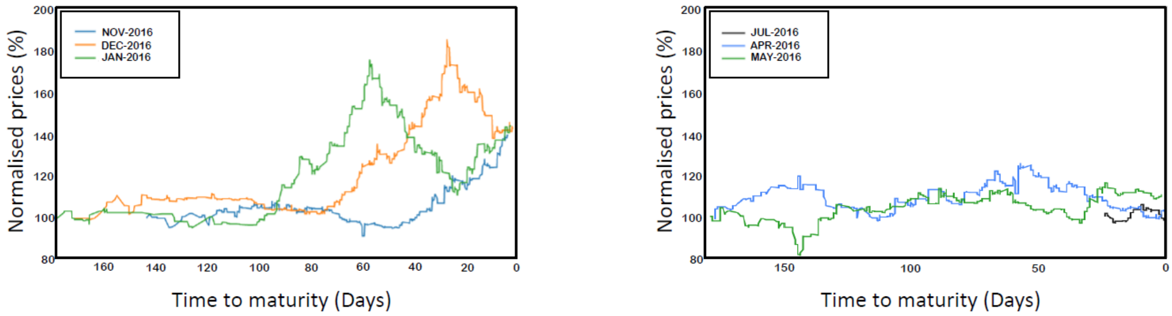

(a) DHM winner contracts

(b) DHM loser contracts

Figure 7.4: Price time series of contracts where DHM performs better (a) and worse (b) than a static hedging approach. The prices of the contracts that show outperformance increase substantially more than for the contracts that show underperformance. The price trend could partly explain why backloading gives higher average revenues.

7.4.1 Daily trading trajectories

The `Integrated Postponement model` attempts to make a trade-off between the preference for large order volumes in the short term and the price impact associated with these trades. The trading trajectory is then a result of the market conditions that influence the steepness of this trade-off, namely the slope of the LOB and the market volatility. As the regression analysis did not uncover a statistically significant volatility process, the volatility used in the `Integrated Postponement model` was assumed to be constant at 22% annualized. Thus, it will be the slope of the LOB that influences the convexity of the trading trajectory.

As discussed in section 7.2.1, the convexity of the trading trajectories increases as contract maturity approaches. By inspection of the development of the LOB slopes, available in figure 7.5, one can see that the absolute value of the slope decreases towards maturity. This evolution of the slope is a consequence of increased contract liquidity towards its maturity. The price impact term carries less weight in the objective function when the slope decreases. In that case, the risk term dominates the objective function and leads to convex trading trajectories close to maturity. The LOB slope is larger far from maturity, so the price impact term dominates the decision policy, which gives a more linear trading trajectory. These results are consistent with the findings of Almgren and Chriss (1999).

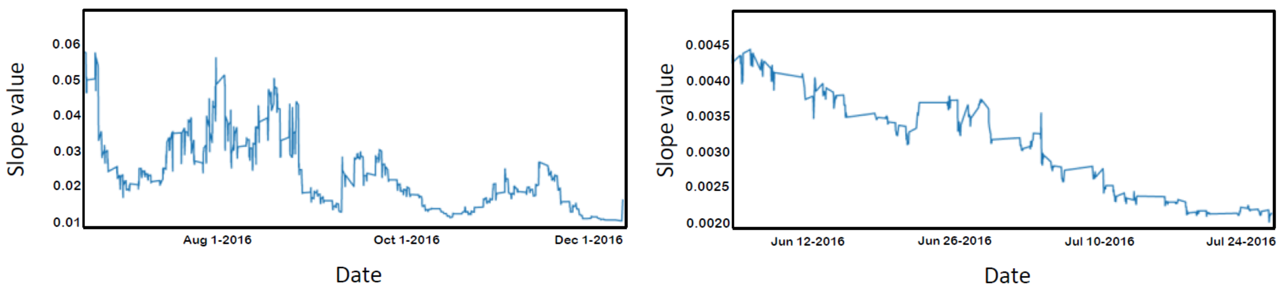

(a) Q1-2017

(b) AUG-2016

Figure 7.5: The absolute value of the slope of the LOB towards time to maturity. The slopes decrease as maturity approaches, indicating that the contracts become more liquid.

The quarterly contracts have close to linear trading trajectories, despite the market liquidity effect towards maturity. The lower degree of convexity can be explained by the higher trading volumes allocated to quarterly contracts. Since the price impact is modelled to increase quadratically with respect to order volume, the price impact term carries a higher weight in the decision policy relative to the risk term. Therefore, the trading volume is spread more evenly to reduce price impact.

7.4.2 Postponement decisions

Postponement decisions are made by comparing the best bid price, p_1^b, to the price-equivalent shadow price, where the ratio of bid price to price-equivalent shadow price, p_1^b/p_{SP}, can be used to interpret the postponement decisions of the `Integrated Postponement model`. The postponement option is speculated to improve revenue per MWh by two distinct mechanisms. First, when the price-equivalent shadow price is higher than the market price, the `Integrated Postponement model` will outright postpone the remaining volume. This type of outright postponement can be seen as the model being selective, only accepting prices above the threshold represented by the price-equivalent shadow price. This mechanism ought to prevent order execution at drastically unfavourable prices. By this token, trading only occurs in periods with favourable prices.

The second mechanism by which the postponement option improves revenue per MWh is limiting price impact. If p_1^b/p_{SP} is above unity, the `Integrated Postponement model` will favour order placement, but not without considering order volume. The price impact is modelled such that the marginal revenue decreases linearly with respect to order volume. As more volume is added to an order, at one point, the marginal revenue will fall below the price-equivalent shadow price. At this point the `Integrated Postponement model` will decide to postpone the remaining trading volume. This second mechanism can be seen as working to limit the price impact of a trade.

Since an electricity producer can choose to be a flexible trader far from maturity, but can not afford this flexibility close to delivery, one should aim to implement a trading model where postponement diminishes towards a contracts delivery date, which is the case for the `Integrated Postponement model`. A decreasing postponement fraction should reflect the transition from a "choose-to" to a "need-to" trader. If the price-equivalent shadow price decreases relative to the best bid price, preference is given to trading rather than a postponement. In the data set used for the backtest, the time series of prices have generally shown price appreciation towards maturity. Additionally, shadow prices are seen to decrease towards maturity. Both of these trends lead to less postponement as maturity approaches.

Chapter 8

Conclusion

This thesis has explored the order execution problem in the context of an illiquid financial market, specifically the electricity futures market. The goal of the order execution problem is to develop a set of trading decision policies to maximise trading performance, measured by trading revenues, the standard deviation of trading revenues and price impact. The objective of the exploration has been to develop an order execution model that introduces the option of trading postponement, integrated with a dynamic hedging model with daily granularity.

From this thesis, there are three main contributions to the study of order execution. First, we have introduced postponement optionality to the order execution literature, a novelty in the field. The optionality is proposed as a result of the distinction made between two types of traders. We make the distinction between traders who need to trade and traders who choose to trade. Additionally, we have proposed a novel and intuitive method for determining the value of waiting with order execution, built on treating the shadow price of the `QUASAR Dynamic Hedging model` as the opportunity cost of trading.

The second contribution is the integration of a multistage stochastic hedging model with a multistage stochastic order execution model. Trading volumes have conventionally been treated as fixed, and therefore not as a variable to decide upon. We expand the decision space to include trading volume allocation, decided daily. Daily trading targets are decided by using the `QUASAR Dynamic Hedging model`, which considers prior trading and market liquidity in its decision making.

The third contribution of this thesis has been the construction and application of a backtesting framework that incorporates the microstructure of limit order books. Backtesting frameworks that incorporate discrete order arrival have previously been employed to study order execution of blue-chip stocks. However, to the authors' best knowledge, the methodology has not previously been used to explore the aspect of price impact in an illiquid market context. We argue that a microstructural backtesting framework that considers discrete order arrival is appropriate for evaluating trading performance in an illiquid market. Additionally, it is worth mentioning that using microstructural order books to simulate the actual trading of futures contracts in the market as part of hedging is a novelty in the research field of hedging models.

The results from the backtest are summarised next. The `Integrated Postponement model` improves average revenue per MWh by 234 bps, compared to placing a single daily block order, i.e. the IOBE strategy. By introducing the postponement option, revenue per MWh was improved by 88 bps, compared to the equivalent model without the postponement option. Including risk-aversion in the model objective function reduces the variance of trading revenues. Its inclusion does not appear to reduce trading revenues or increase price impact. This result is inconsistent with the findings of Almgren and Chriss (1999), commenting on the trade-off between price impact and risk-aversion.

Next, we comment on a few possible flaws and discuss the application of our findings. The backtest has been conducted for the period of June 2016 through June 2017. While the computational demands of the backtest put a restriction on the length of the simulation period, it would have been preferable to conduct the backtest on a larger sample of market data had there not been restrictions on the available memory on the server. Ideally, a larger sample should have been used to see strategy performance across different parts of the market cycle. A larger sample would also allow for the trading of yearly contracts.

Furthermore, the price volatility for all futures contracts was assumed to be constant and equal across all the contracts in this thesis. This assumption may not be appropriate for several reasons. First, the various futures contracts are different both in respect to liquidity and their delivery period. The quarterly contracts have a longer delivery period than the monthly contracts, and the liquidity of futures contracts varies based on the time of delivery during the year. Thus, all contracts are fundamentally different. Second, assuming the volatility to be constant in respect to time may also be flawed. The trading volumes on a contract tend to increase towards maturity. The increased liquidity towards maturity, in addition to a maturity effect on volatility, yield strong arguments for incorporating a time-dependent price volatility process.

The application of our findings deserves a few comments. The `Integrated Postponement model` is best suited for the liquidation of power portfolios for a power producer. The `QUASAR Dynamic Hedging model` is able to make a better allocation of power production to the spot and futures market while also considering the relative liquidity of futures contracts in its policy. The `QUASAR Dynamic Hedging model` also retrieves the shadow prices that determine the value of waiting with order execution, which reflects the "choose-to" trader characteristic of a power producer. The `Integrated Postponement model` is also suited for the liquidation of other commodities. However, if the commodity does not share the characteristic of power where contracts can share the underlying asset, the `QUASAR Dynamic Hedging model` will not make trading decisions based on the relative liquidity of contracts. The `Postponement model` carries application beyond the commodities market and could be used for liquidation of stocks. However, its use relies on quantifying a value for the postponement optionality. The outperformance of the postponement strategy should be seen in the context of the trader's market share. A smaller market participant will to a lesser extent, experience price impact and therefore not draw as much value from using the `Integrated Postponement model`.

Two areas of future research are proposed here. First, we propose the inclusion of limit order placement in the order execution model. By allowing the placement of limit orders, the model will reflect the decision space of a trader to a greater extent. Additionally, it could improve trading performance if the trade-off between risk and reward is modelled appropriately. Second, one should explore other methods for assigning a value of postponing order placement. The price-equivalent shadow price has been evaluated as the value of waiting, and it has an intuitive interpretation as the opportunity cost of trading. However, due to the novelty of the postponement idea, one should not rule out that other methods for assessing the postponement reference price yield better model performance. One such extension could be to explore different methods of calibration for the scaling parameter ζ.

References

Adam, T. R. & Fernando, C. S. (2006). Hedging, speculation, and shareholder value. *Journal of Financial Economics*, *81*(2), 283–309. https://doi.org/https://doi.org/10.1016/j.jfineco.2005.03.014

Agliardi, R. & Gençay, R. (2017). Optimal trading strategies with limit orders. *International Journal of Theoretical and Applied Finance*, *20*(1). https://doi.org/https://doi.org/10.1142/S0219024917500054

Akersveen, S. & Graabak, J. (2018). *A multistage stochastic optimization model for trading on the german intraday power market* (Master's thesis). Norwegian University of Science and Technology.

AleaSoft. (2017). *Assessment of prices of the electricity european markets at the end of 2016*. https://aleasoft.com/assessment-prices-electricity-european-markets-end-2016/

Allayannis, G. & Weston, J. (2001). The use of foreign currency derivatives and firm market value. *Review of Financial Studies*, *14*, 243–276. https://doi.org/https://doi.org/10.1093/rfs/14.1.243

Almgren, R. (2012). Optimal trading with stochastic liquidity and volatility. *SIAM J Financial Math*, *3*(1), 163–181. https://doi.org/https://doi.org/10.1137/090763470

Almgren, R. & Chriss, N. (1999). Optimal execution of portfolio transactions. *Journal of Risk*, *3*(2), 5–39.

Almgren, R., Thum, C., Hauptmann, E. & Li, H. (2005). Direct estimation of equity market impact. *Risk*, *18*(7), 58–64.

Amihud, Y. & Mendelson, H. (1986). Asset pricing and the bid-ask spread. *Journal of Financial Economics*, *17*(2), 223–249. https://doi.org/https://doi.org/10.1016/0304-405X(86)90065-6

Anderson, R. W. & Danthine, J.-P. (1980). Hedging and joint production: Theory and illustrations. *The Journal of Finance*, *35*(2), 487–498. https://doi.org/https://doi.org/10.1111/j.1540-6261.1980.tb02180.x

Bertrand, G. (2021). *Enhancing the value of flexible generation capacity through continuous intraday markets and scarcity pricing* (Doctoral dissertation). Universite Catholique de Louvain.

Bertsimas, D. & Lo, A. W. (1998). Optimal control of execution costs. *Journal of Financial Markets*, *1*, 1–50. https://doi.org/10.1109/5992.805135

Bikhchandani, S. & Sharma, S. (2001). Herd behavior in financial markets. *IMF Staff Papers*, *47*, 1–1.

Cheriyan, N. K. & Daniel, L. (2019). Relationship between liquidity, volatility and trading activity: An intraday analysis of indian stock market. *International journal of economics and financial issues*, *9*(1), 17–22. https://doi.org/10.32479/ijefi.7268

Cont, R. & Kukanov, A. (2017). Optimal order placement in limit order markets. *Quantitative finance*, *17*(1), 21–39. https://doi.org/0.1080/14697688.2016.1190030

Cont, R., Stoikov, S. & Talreja, R. (2010). A stochastic model for order book dynamics. *Operations Research*, *58*(3), 549–563. https://doi.org/https://pubsonline.informs.org/doi/abs/10.1287/opre.1090.0780

Dimoski, J., Fleten, S.-E., Löhndorf, N. & Nersten, S. (2019). *Dynamic hedging for the real option management of hydropower production with exchange rate risks* (Working Paper). Norwegian University of Science and Technology.

Duong, H. N. & Kalev, P. S. (2008). The samuelson hypothesis in futures markets: An analysis using intraday data. *Journal of banking finance*, *32*(4), 489–500. https://doi.org/https://doi.org/10.1016/j.jbankfin.2007.06.011

Dupuis, D., Gauthier, G. & Godin, F. (2016). Short-term hedging for an electricity retailer. *The Energy Journal*, *37*(2), 116–137. https://doi.org/10.5547/01956574.37.2.ddup

Dyrkolbotn, S., Hove, A. & Jørgensen, C. F. (2020). *Order placement in iliquid markets - a case study of nasdaq nordi* (tech. rep.).

E24. (2020). Nå tilbyr flere strømselskaper å betale deg for å bruke strøm. https://e24.no/olje-og-energi/i/dO1v7O/naa-tilbyr-flere-stroemselskaper-aa-betale-deg-for-aa-bruke-stroem/ (Accessed: 2021-03-10).

EXX. (2021a). *List of services and prices of eex ag*. EEX. https://www.eex.com/fileadmin/EEX/Downloads/Trading/Price_Lists/20210504_Price_List_EEX_AG_0088a_E_FINAL.pdf (Accessed: 2021-03-19).

EXX. (2021b). Power futures. https://www.eex.com/en/markets/power-derivatives-market/power-futures (Accessed: 2021-06-08).

Feng, Y., Rubio, F. & Palomar, D. P. (2012). Optimal order execution for algorithmic trading: A cvar approach. *2012 IEEE 13th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, 480–484. https://doi.org/10.1109/SPAWC.2012.6292954

Fleten, S.-E. (2000). *Eportfolio management emphasizing electricity market applications. a stochastic programming approach* (Doctoral dissertation). Norwegian University of Science and Technology.

Fleten, S.-E., Bråthen, E. & Nissen-Meyer, S.-E. (2010). Evaluation of static hedging strategies for hydropower producers in the nordic market. *J. Energy Markets*, *3*. https://doi.org/10.21314/JEM.2010.050

Fleten, S.-E., Wallace, S. & Ziemba, W. (2002). Hedging electricity portfolios via stochastic programming. *Decision making under uncertainty* (pp. 71–93). Springer.

Frestad, D. (2012). Liquidity and dirty hedging in the nordic electricity market. *Energy Economics*, *34*(5), 1341–1355. https://doi.org/https://doi.org/10.1016/j.eneco.2012.06.017

Haar, R. (2010). *On modelling the electricity futures curve* (Bachelor's Thesis).

Harvey, C. R. & Liu, Y. (2015). Backtesting. *The Journal of Portfolio Management*, *42*(1), 13–28. https://doi.org/https://jpm.pm-research.com/content/42/1/13/tab-pdf-disaabled

Heath, D., Jarrow, R. & Morton, A. (1992). Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica*, *60*(1), 77–105. https://doi.org/https://doi.org/10.1111/j.1540-6261.1980.tb02180.x

Hora, M. (2006). *Tactical liquidity trading and intraday volume* (Working Paper).

Kettunen, J., Salo, A. & Bunn, D. W. (2007). Dynamic risk management of electric- ity contracts with contingent portfolio programming. http://www.sal.hut.fi/Publications/m-index.html

Kyle, A. (1985). Continuous auctions and insider trading. *Econometrica*, *53*(6), 1315–1335. https://doi.org/https://doi.org/10.2307/1913210

Lee, C. M., Mucklow, B. & Ready, M. J. (1993). Sprads, depths, and the impact of earnings information: An intraday analysis. *The Review of Financial Studies*, *6*(2), 345–374. https://doi.org/https://www.researchgate.net/profile/Charles-Lee-28/publication/5217314_Spreads_Depths_and_the_Impact_of_Earnings_Information_An_Intraday_Analysis/links/00b495231f6915a111000000/Spreads-Depths-and-the-Impact-of-Earnings-Information-An-Intraday-Analysis.pdf

Lehalle, C.-A. & Mounjid, O. (2018). Limit order strategic placement with adverse selection risk and the role of latency. *Market Microstructure and Liquidity*, *3*(1).

Lin, R. & Chang, Y. (2009). *Does hedging add value? evidence from the global airline industry* (Working Paper).

Liu, S. D., Jian, J. B. & Wang, Y. Y. (2010). Optimal dynamic hedging of electricity futures based on copula-garch models. *2010 IEEE International Conference on Industrial Engineering and Engineering Management*, 2498–2502. https://doi.org/10.1109/IEEM.2010.5674323

Liu, W. (2006). A liquidity-augmented capital asset pricing model. *Journal of Financial Economics*, *82*(3), 631–671. https://doi.org/https://doi.org/10.1016/j.jfineco.2005.10.001

Löhndorf, N., Wozabal, D. & Minner, S. (2013). Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming. *Operations Research*, *61*(4), 810–823. https://doi.org/https://doi.org/10.1287/opre.2013.1182

Löhndorf, N. & Wozabal, D. (2021). Gas storage valuation in incomplete markets. *European Journal of Operational Research*, *288*(1), 318–330. https://doi.org/https://doi.org/10.1016/j.ejor.2020.05.044

Matras, K. (2011). *Finding 1 stocks, screening backtesting and time-proven strategies*.

McDonald, R. L. (2013). *Derivatives markets*.

Mete Soner, H. (2015). Trading with market impact. https://people.math.ethz.ch/~jteichma/its_talks/soner-its.pdf (Accessed: 2021-03-10).

Modigliani, F. & Miller, M. (1958). The cost of capital, corporation finance and the theory of investment. *American Economic Review*, *48*(3), 261–297.

Näsäkkälä, E. & Keppo, J. (2005). Electricity load pattern hedging with static forward strategies. *Managerial Finance*, *31*, 116–137. https://doi.org/10.1108/03074350510769721

Paulson, L. & Starn, J. (2018). Aas lost millions on nordic power, but others had similar bets. https://www.bloombergquint.com/business/aas-lost-millions-on-nordic-power-but-others-had-similar-bets (Accessed: 2021-3-1).

Pereira, M. & Pinto, L. (1991). Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming*, *52*, 359–375. https://doi.org/https://doi.org/10.1007/BF01582895

Pineda, s. & Conejo, A. (2013). Using electricity options to hedge against financial risks of power producers. *Journal of Modern Power Systems and Clean Energy*, *1*(2), 101–109.

Predoiu, S., Shaikhet, G. & Shreve, S. (2011). Optimal execution in a general one-sided limit-order book. *SIAM J. Finan. Math*, *2*(1), 183–212. https://doi.org/https://epubs.siam.org/doi/abs/10.1137/10078534X

Ruszczyński, A. & Shapiro, A. (2006). Conditional risk mappings. *Mathematics of Operations Research*, *31*(3), 544–561. http://www.jstor.org/stable/25151744

Sanda, G., Tandberg Olsen, E. & Fleten, S.-E. (2013). Selective hedging in hydro-based electricity companies. *Energy Economics*, *40*, 326–338. https://doi.org/https://doi.org/10.1016/0304-405X(86)90065-6

Shapiro, A., Tekaya, W., Costa, J. & Soares, M. (2013). Risk neutral and risk averse stochastic dual dynamic programming method. *European Journal of Operational Research*, *224*, 375–391. https://doi.org/10.1016/j.ejor.2012.08.022

Shen, J. (2014). A pre-trade algorithmic trading model under given volume measures and generic price dynamics. *2015*(1), 64–98. https://doi.org/https://doi.org/10.1093/amrx/abu007

Shen, J. (2017). Hybrid is-vwap dynamic algorithmic trading via lqr. *Social Sci. Res. Network (SSRN) Preprint*, 1–23.

SSB. (2020). Elektrisitet. https://www.ssb.no/energi-og-industri/statistikker/elektrisitet/aar/ (Accessed: 2021-03-23).

Stulz, R. M. (1996). Rethinking risk management. *Journal of Applied Corporate Finance*, *9*(3), 8–25. https://doi.org/https://doi.org/10.1111/j.1745-6622.1996.tb00295.x

Zanotti, G., Gabbi, G. & Geranio, M. (2010). Hedging with futures: Efficacy of garch correlation models to european electricity markets. *Journal of International Financial Markets, Institutions and Money*, *20*(2), 135–148. https://doi.org/https://doi.org/10.1016/j.intfin.2009.12.001

Zheng, H. (2016). *Microstructure analysis of dynamic markets: Limit order books and dynamic matching markets* (Doctoral dissertation). Columbia University.

Appendix A

Daily volatility predictor study

Concerning the "Time to Maturity" hypothesis, we construct a volatility predictor, where volatility is a function of time to maturity (TTM). Additionally, we test whether an AR(1) model is applicable to model the volatility. These models are presented in equation (A.1) and (A.2) respectively.

$$\sigma_t = \theta_0 + \theta_1 \cdot t_{TTM} + \epsilon_t^{ttm} \quad \text{where} \quad \epsilon_t^{ttm} \sim N(0, \sigma_{\sigma_t}^2) \tag{A.1}$$

$$\sigma_t^{AR} = \theta_0^{AR} + \theta_1^{AR} \cdot \sigma_{(t-1)}^{AR} + \epsilon_t^{AR} \quad \text{where} \quad \epsilon_t^{AR} \sim N(0, \sigma_{\sigma_t^{AR}}^2) \quad \text{and} \quad 0 \leq \theta^{AR} \leq 1 \tag{A.2}$$

To estimate the model parameter values, we use linear regression on EEX order data from 2015. Monthly and quarterly contracts are regressed separately. The results from the regressional analysis are presented in table A.1. We notice that TTM is statistically significant as a predictor for volatility since it has a P-value less than 0.05 for both monthly and quarterly contracts. The results also indicate that delivery duration affects the volatility, as the absolute value of the TTM coefficient is smaller for quarterly contracts than for monthly contracts.

An out-sample test is conducted to verify model validity, using EEX order data from 2014. These results are presented in table A.2. The AR(1) model performs better for quarterly contracts than for monthly contracts, which is expected due to the contracts' longer delivery duration, resulting in more stable volatility. For monthly contracts, the TTM model performs better than the AR(1) model. This is expected as the delivery duration is shorter for monthly contracts. It is worth noting that the adjusted R-squared values are low. Three of the four R-squared values are negative, indicating that neither of these models are strong predictors. Both models perform similarly in terms of residual standard errors, as seen in A.2. From these results, it can be concluded that neither of the models are good predictors to model future volatility.

Table A.1: In-sample regressional results for volatility parameter estimation.

Results from TTM model					
Contract type	Intersect	Coef.	P-value	Adjusted R^2	Residual standard error
Quarterly	0.1452	-0.0002	0.000***	0.0124	0.2371
Monthly	0.1876	-0.0009	0.000***	0.0394	0.0534

Results from AR(1) model					
Contract type	Intersect	Coef.	P-value	Adjusted R^2	Residual standard error
Quarterly	0.0762	0.2646	0.000***	0.0692	0.2302
Monthly	0.0694	0.3855	0.000***	0.1480	0.2043

*** Significant at P<0.01
** Significant at P<0.05
* Significant at P<0.1

Table A.2: Out-of-sample regressional results for volatility parameter estimation.

Out-of-sample results from the TTM model		
Contract type	Adjusted R^2	Residual standard error
Quarterly	-0.1459	0.11100
Monthly	0.0505	0.14560

Out-of-sample results from AR(1) model		
Contract type	Adjusted R^2	Residual standard error
Quarterly	-0.0697	0.11095
Monthly	-0.4330	0.10709

Appendix B

Out-of-sample price volatility

The average daily volatilities for the contracts from the out-of-sample period have been computed and then annualised. The volatilites are seen in table B.1. The average annualised volatility across these contracts is 22.38%, which supports the empirical findings of Haar (2010). The volatilities are averages across all trading days for each contract. While an average value gives an indication of the volatility level, it can not be used to uncover any underlying trends, such as how the volatility develops towards contract maturity.

Table B.1: Price volatility for monthly and quarterly futures contracts tradeable between June 2015 and June 2016.

	JUL-2015	AUG-2015	SEP-2015	OCT-2015	NOV-2015
Daily volatility(%)					
Median	1.46	1.47	1.24	1.21	1.50
Annual volatility(%)					
Median	23.08	23.24	19.61	19.13	23.72

	DEC-2015	JAN-2016	FEB-2016	MAR-2016	APR-2016
Daily volatility(%)					
Median	1.34	1.64	1.55	1.52	1.37
Annual volatility(%)					
Median	21.19	25.93	24.51	24.03	21.66

	MAY-2016	Q3-2015	Q4-2015	Q1-2016
Daily volatility(%)				
Median	1.29	1.46	1.20	1.57
Annual volatility(%)				
Median	20.40	23.08	18.97	24.82

Appendix C

Price driver study

Part of this thesis has been the attempt to make a price prediction model. By examining the correlation between prices and other variables, we study whether microstructural features can be used to predict future price movements. The features that have been tested are presented in section C.1. If any features had turned out to be statistically significant, with a high adjusted R-squared, they would have been included in the exogenous price dynamic process. Section C.2 will explain the methodology for measuring price correlation of the features, and section C.3 presents the results of the study.

C.1 Tested features

We study whether the following microstructural features are correlated with the price movement of EEX power futures:

The bid-ask spread: As mentioned in section 2.2.2, the bid-ask spread is used as a measure of liquidity. One could argue that when the market is more liquid, one can obtain better prices. Thus the bid-ask spread is considered a potential price driver.

Volatility: Market volatility could indicate market liquidity, which may provide a favourable price environment for a liquidity taker.

Order book imbalance: Order book imbalance reflects overselling or overbuying in the market. An imbalanced LOB may reflect new information, leading to a surplus of traders wishing to buy or sell the asset. Order book imbalance is tested for the best bid and ask and the three best bids and ask levels in the LOB.

Best bid and ask depth: A deep bid depth could indicate overbuying, which could be exploitable for sellers. We test the depth of the best, three best and five best bids and asks.

Time of day: As mentioned in section 2.2.2, hourly traded volume is an indicator of liquidity. It is assumed that the hourly traded volumes will be similar for all days. This assumption appears reasonable as larger trading volumes tend to be transacted right after opening or just before closing. This is speculated to be due to newly available information between trading days. Traders may choose to close their position every day before closing to eliminate the overnight risk. This behaviour could also explain the higher trading activity in the morning and just before closing.

Micro price: The midprice, the best bid and ask average, is often referred to as the market price. However, it does not consider LOB volumes. Therefore, one could propose to test a metric that takes both prices and volumes in the LOB into account. The micro price is a volume-weighted average price of the best bid and ask, as seen in equation (C.1). The micro price relative to the market price could indicate the future price movement.

$$P_{Micro} = \frac{P_a \cdot V_b}{V_a + V_b} + \frac{P_b \cdot V_a}{V_a + V_b} \tag{C.1}$$

Volume of most recent trade: If the trading volume of the most recent trade is substantially large, one would expect market recovery going forward under the assumption of market resiliency. Recovery would imply that the price will change during the next hour.

C.2 Methodology for indicating price correlation

The price correlation of the features in C.1 are determined by conducting linear regression. The EEX data from 2014 through 2015 has been used to conduct the regressional analysis. Variables such as best bid, spread, volatility, order depth and order book imbalance are collected for every trading hour for all contracts. The time of day feature is classified as a set of dummy variables. There is a binary variable for every trading hour throughout the trading day.

The data is split into two sets of equal duration, one for fitting and one for evaluating the model performance as a price predictor. We follow a wrapper method for feature selection, specifically using forward step-wise selection. The features with adjusted R-squared values higher than 0.6 are then analysed further. The logarithmic difference between the best bid price of two consecutive hours is classified as the dependent variable while the features are classified as independent variables, as shown in equation (C.2). The regressions with mid-price and micro price as independent variables are conducted with the next hour's best bid price as the dependent variable, unlike the other features where the logarithmic change in bid price is classified as the dependent variable. This adjustment is made since the mid-price and micro prices' effect on the new best bid price should be seen relative to the old best bid price, unlike the other features where the correlation is independent of the price level. Quarterly and monthly contracts are regressed separately. This is done since they have different delivery duration, thus are expected to behave differently.

$$ln(p_{t+1}^{bid}) - ln(p_t^{bid}) = \alpha_0^{\text{feature}} + \alpha_1^{\text{feature}} \cdot x_t^{\text{feature}} \tag{C.2}$$

C.3 Results from linear regression

The results from the price driver study are presented in tables C.1 and C.2, one for monthly contracts and the other for quarterly contracts. We notice that many features show signs of being statistically significant but with low adjusted R-squared values. The feature with the highest adjusted R-squared is the mid-price. An adjusted R-squared value of 0.8920 does seem promising.

However, by comparing the adjusted R-squared value to an adjusted R-squared value of 0.98-0.99, which was achieved with a random walk process as the dependent variable, the mid-price could not be said to improve the exogenous price process quality. For the other features, the bid-ask spread has the highest adjusted R-squared value. The adjusted R-squared value of the bid-ask spread for the set of monthly contracts is higher than for the set of quarterly contracts, which could be explained by the lower trading activity of monthly contracts. Therefore, these contracts are less liquid, resulting in the bid-ask spread predicting the price movements more accurately than for the quarterly contracts.

Based on the low R-squared values, the price driver study was stopped at the first step, as there seems to be a limited correlation between each feature and the price. A random walk process was therefore used as the exogenous price process for the `Postponement model`.

The low adjusted R-squared values could be explained by the EEX market being one of the most liquid power markets. A liquid market is closer to an efficient market, where the random walk and GBM price processes are more appropriate. One avenue of future work could then be to conduct a similar study on a less liquid market, like the NASDAQ Nordic market. On the other hand, microstructural features are primarily used in high-frequency trading and not on an hour-to-hour trading frequency. Using microstructural abnormalities to exploit the market state primarily works over short time intervals. Therefore, one could think that there is limited value in microstructural feature analysis for someone trading hour-to-hour.

Table C.1: Results from parameter estimation for the price driver study for monthly contracts. A total of 11081 observations of each feature were collected. All features have low R-squared values except mid price, due to its high correlation to the best bid.

Feature	Intercept	Coef.	P-value	Adjusted R^2	Residual error
Bid-ask spread	-0.0048	0.7851	0.0000***	0.2176	0.0640
Volatility	0.0004	-0.0026	0.55987	-0.0001	0.0720
Ask depth	0.0009	-0.0001	0.49500	0.0000	0.0727
Bid depth	0.0027	-0.0005	0.00319***	0.0007	0.0728
Volume of three best bids	0.0025	-0.0001	0.07270*	0.0002	0.0738
Volume of three best asks	0.0014	-0.0001	0.39563	0.0000	0.0736
Volume of five best asks	0.0017	0.000	0.35391	0.0000	0.0743
Volume of five best bids	0.0024	-0.0001	0.12752	0.0001	0.0745
LOB imbalance	0.0004	-0.0052	0.00471***	0.0007	0.0732
LOB imbalance for best 3 levels	0.0006	-0.0054	0.02532**	0.0004	0.0734
Time of day	-	-	0.43144	0.000	0.0720
Micro price	34.8363	0.0674	0.00000***	0.0460	2.4133
Mid price	2.4096	0.9237	0.00000***	0.8920	0.9791
Volume of most recent trade	0.0124	-0.0004	0.00445***	0.0007	0.0739

- since there are a total of 9 binary variables, we do not present the coefficients due to lack of space.
*** Significant at p<0.01
** Significant at p<0.05
* Significant at p<0.1

Table C.2: Results from parameter estimation for the price driving study for quarterly contracts. A total of 5530 observations of each feature were collected.

Feature	Intercept	Coef.	P-value	Adjusted R^2	Residual error
Bid-ask spread	-0.0012	0.4114	0.0000***	0.0601	0.0077
Volatility	0.0004	-0.0077	0.00000***	0.0241	0.0078
Ask depth	0.000	-0.0000	0.88588	-0.0002	0.0079
Bid depth	-0.0005	0.0001	0.00074***	0.0019	0.0079
Volume of three best bids	-0.0006	0.0000	0.00729***	0.0011	0.0079
Volume of hree best asks	-0.0001	0.0000	0.83880	-0.0002	0.0079
Volume offive best asks	-0.0002	0.000	0.56374	-0.0001	0.0079
Volume of five best bids	-0.0006	0.000	0.01872***	0.0008	0.0079
LOB imbalance	-0.0001	0.0005	0.02808**	0.0007	0.0079
LOB imbalance for best 3 levels	-0.0001	0.0006	0.10992	0.0003	0.0079
Time of day	-	-	0.92266	-0.0009	0.0079
Micro price	30.5544	0.1100	0.00000***	0.0566	2.9153
Mid price	0.2162	0.9925	0.000***	0.9887	0.2621
Volume of most recent trade	-0.0003	0.0000	0.87742	-0.0002	0.0079

- since there are a total of 9 binary variables, we do not present the coefficients due to lack of space.

*** Significant at $p<0.01$

** Significant at $p<0.05$

* Significant at $p<0.1$

Appendix D

Miscellaneous processes for the QUASAR Dynamic Hedging Model

This thesis does not detail the implementation of the Quasar Dynamic Hedging model. However, a few aspects of model implementation are worth mentioning. The Quasar hedging model is used to create a hedging plan that dictates trading decisions. The QUASAR Dynamic Hedging model requires a production plan to make hedging decisions. The production plan ought to reflect a power producer with a large market share to fully experience a lack of market liquidity. Consequentially, we have scaled a normalised production plan to fit a large power producer. The process of scaling the production volumes is explained in section D.1. Additionally, the QUASAR Dynamic Hedging model requires a scenario lattice of price forward curves to solve the optimisation problem by using ADDP. The process of creating these lattices will be described in section D.2. Lastly, three of the benchmark models do not use the QUASAR Dynamic Hedging model. Therefore we present the static hedging plan for these strategies in section D.3.

D.1 Scaling the production plan

The hedging decisions will be dependent on the production volumes. Higher production volumes will lead to higher trading volumes of futures contracts. It is necessary to scale the production plan to balance the effects of price impact and price risk. If the production volumes are too low, the power producer will not experience any price impact, and the trade-off between price risk and price impact is irrelevant. If the production volumes are too high, the power producer will drain too much of the market liquidity, resulting in severe adverse price movements. Therefore, it is essential to find a suitable scaling factor to thoroughly analyse the trade-off between price impact and price risk. The scaling factor is determined by running an out-of-sample simulation where different trading volumes were tested to balance price risk and price impact.

D.2 Generating scenario lattices for the QUASAR Dynamic Hedging model

Trading decisions are obtained by solving the `QUASAR Dynamic Hedging model` with daily granularity. As described in section 4.1.1, the ADDP algorithm relies on using forward simulations to approximate the high-dimensional objective function. The continuous stochastic price process has to be discretised to cohere with the solution algorithm. Scenario lattices have been generated for this purpose. The hedging problem is solved for the time interval of a specific start date, T_{SD}, and end date, T_{ED}. Given that T_{SD} is incremented forwards while T_{ED} stays constant, a new scenario lattice is required for each daily optimisation. The number of stages in each lattice equals the number of days in the time interval, $T_{ED} - T_{SD}$. Constructing the lattice such that nodes recombine gives it a constant number of nodes for every stage. In our case, the number of nodes per stage is 10.

The state variable for node i in stage t, N_{ti} corresponds to the realisation of a price forward curve, described by the set of forward prices from stage t to stage T_{ED}, $\{F^i_{t,T_{ED}}, F^i_{t+1,T_{ED}}, \ldots, F^i_{T_{ED}-1,T_{ED}}\}$. For every node in stage t, we construct a probability distribution for state transitions to the nodes in the subsequent stage $t+1$. Löhndorf and Wozabal (2021) have developed the technique of backwards estimation to estimate these probabilities. Backwards estimation ensures that the expected conditional successor state on the lattice does not deviate from the expected conditional successor state of the continuous process. If the expected conditional successor states did not coincide, arbitrage opportunities would be exploitable.

D.3 Hedging volumes for static volume strategies

Three of the benchmark strategies do not use the `QUASAR Dynamic Hedging model`. Consequentially, hedging volumes are decided in another way. Ideally, all the trading strategies should trade equal volumes of all contracts. To ensure this occurs, we construct the hedging plan of the static strategies on the hedging plan retrieved from the dynamic volume-equivalent strategies.

The benchmark strategies that use the `QUASAR Dynamic Hedging model` are backtested first. The traded volumes of each contract are then aggregated to find the total volume traded of each contract. The aggregated contract volumes are spread evenly across all trading days, as seen in figure D.1. This is referred to as the static hedging plan. It could have been referred to as a TWAP hedging strategy as volumes are distributed like a TWAP trading strategy. For instance, the TWAP DHM strategy is backtested first before the static TWAP strategy. The trading volumes are then aggregated and distributed evenly across all trading days. Evenly distributing the trading volumes across the trading days is a natural heuristic for trading volume allocation. Thus it is used as a benchmark hedging strategy to analyse the effect of the `QUASAR Dynamic Hedging model`.

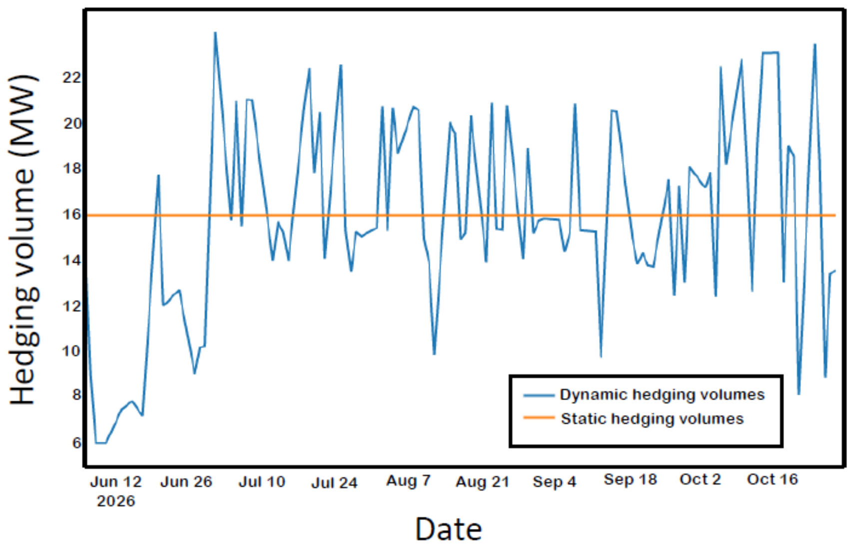

Figure D.1: The static and the dynamic hedging plan for the NOV-2016 futures contract. Note that the trading volumes for static volume strategies are evenly distributed across the contract lifecycle.

Appendix E

Algorithms

Preparing for the backtest consists of conducting several processes, starting with the reverse engineering of the order data to programming the trading algorithms. This has been implemented using the programming languages Jupyter Scala and QUASAR has been used to solve the optimisation problems. Pseudo code for the most essential scripts are provided in this appendix.

E.1 Reverse engineering of order data

The EEX LOB dataset has been reconstructed to comply with the LOB simulation algorithm employed by Dyrkolbotn et al. (2020). Algorithm 1 was used for this purpose. The following section describes this algorithm and its underlying assumptions.

The EEX LOB data details the state of the LOB at an instance, without providing context on how the LOB state was reached. By state, it is meant the combination of prices and volumes for all orders in the LOB at one instance. The LOB state has been reached by the placement of a particular set of orders with respective prices, volumes and timestamps. By observing consecutive LOB instances, one can deduce the characteristics of the order changes that led to the new LOB instance. If the price or volume of the order at a LOB level is different from the price and volume of the order at the same LOB level at the next LOB instance, at least one order change has occurred. By observing the all LOB levels on the bid and ask sides of the LOB, we can deduce the characteristics of the order that was placed, changed or cancelled.

Types of order changes in the LOB

- Change in order price
- Change in order volume
- Cancellation of order
- New order placed in LOB

Assumptions for reverse engineering of order data

- If there is no order matching, it is assumed that maximum one order change for each side of the LOB.
- In the case of an order match, multiple order changes to one side of the LOB can be caused by only one order arrival on the opposite side of the LOB.
- The EEX LOB dataset is incomplete, since it only contains the order data for the first five levels of the LOB. One resulting issue from this is that if an order moves from order level 5 (L5) to order level (L6), it dissappears from the LOB. Therefore it is assumed that if an order is no longer observable in the dataset, the order is deleted, and gets resubmitted if it moves back from L6 to L5.
- For some situations, it is not possible to know with certainty whether an L1 order was matched against an opposing market order of equal volume, or if the order was cancelled. In this case, it is assumed that a market order of the same volume was placed, resulting in a match, as this is the more likely outcome.

```
Data: EEX LOB instances
Result: set of all single orders
Initialization;
while untreated LOB instances exist do
    Check that LOB instance does not contain errors
    if Changes to one or both L1 order exist then
        if Change is a result of an order match then
            Determine volume and price of submitted order
            Create new order that caused the match
        else
            Delete old L1 order
            Create new L1 order
        end
    else
        while not observed order change of ask orders in LOB do
            if new order at current level is better than old order then
                Create new order at current order level
                Determine the type of order change
                if New order at current level has been changed from existing order then
                    Find order that was changed and delete it
                end
            else
                Delete old order at current level
                Determine the type of order change
                if New order at current level has been changed from existing order then
                    Find order that was changed and delete it
                end
            end
        end
        while not observed order change of bid orders in LOB do
            if new order at current level is better than old order then
                Create new order at current order level
                Determine the type of order change
                if New order at current level has been changed from existing order then
                    Find order that was changed and delete it
                end
            else
                Delete old order at current level
                Determine the type of order change
                if New order at current level has been changed from existing order then
                    Find order that was changed and delete it
                end
            end
        end
    end
    Increment to the next LOB instance
end
```

Algorithm 1: Reverse engineering of order data

E.2 Limit order book matching algorithm

An essential part of this thesis is the use of limit order books. Consequentially, an order matching algorithm is needed. Algorithm 2 describes the LOB matching process employed in the backtest.

```
Data: set of buy orders and sell orders, sorted by price in LOB queues
Result: set of all transactions for LOB data
initialization;
while unprocessed orders exist do
    if current order is a market order then
        Initialise tradedVolume to 0
        while tradedVolume is less than market order volume and clearable orders exist do
            if remainingTradingVolume >= OppositeSideL1Order.volume then
                tradedVolume+=OppositeSideL1Order.volume
                delete OppositeSideL1Order
            else
                tradedVolume+=(MarketOrderVolume-RemainingTradingVolume)
                reduce order volume for OppositeSideFirstOrder
            end
        end
        if tradedVolume is less than market order volume then
            send remaining order volume back to LOB as a limit order
        end
        create transaction
        Continue to next order
    else if current order is a limit order then
        if orders with price eligible of match exist then
            while tradedVolume is less than market order volume and clearable orders exist
             do
                if remainingTradingVolume >= OppositeSideFirstOrder.volume then
                    tradedVolume+=OppositeSideFirstOrder.volume
                    delete OppositeSideFirstOrder
                else
                    tradedVolume+=(MarketOrderVolume-RemainingTradingVolume)
                    reduce order volume for OppositeSideFirstOrder
                end
            end
            if tradedVolume is less than market order volume then
                send remaining order volume back to order book as a limit order
            end
            create transaction
            continue to next order
        else
            Continue to next order
        end
    else
        Current order is a cancel order, find corresponding limit order and delete it
        Continue to next order
    end
end
```

Algorithm 2: LOB matching algorithm.

E.3 Backtest Algorithms

The different strategies are backtested using the LOB from E.2. In this section, algorithm 3 explains how the backtest is structured, while algorithm 4 explains the trading sequence for the `Integrated Postponement model`.

Data: LOB data, scenario lattices, production data, liquidity data, first stage hedging decisions
Result: set of all transactions for LOB data
Initialization;
while *in the trading period* **do**
  Run `QUASAR Dynamic Hedging model`
  Store first stage hedging decisions π_d^0
  for *every contract to trade* **do**
    Initialize LOB to current trading date
    while *still on current trading date* **do**
      Place orders according to order strategy
      Increment to next order in LOB
    end
    Store total traded volume on trading date for contract
  end
  Store total traded volumes to be used by `QUASAR Dynamic Hedging model` for the next day
  Increment trading date
end

Algorithm 3: LOB trading simulation sequence.

Data: LOB Data, shadow prices (η_d) from the `QUASAR Dynamic Hedging model`, first stage hedging decisions
(π_d^0) **Result:** Trading policy for order placement problem
Initialization;
while *still on current trading day* **do**
  if *New trading hour* **then**
    Retrieve shadow price from the `QUASAR Dynamic Hedging model`
    Retrieve best bid and slope from the LOB
    Retrieve volume left to trade for the day
    Find number of hours left of trading day
    Run the `Postponement model` with required inputs
    Place order and update volume left to trade based on the first stage solution
  end
  Increment to next order in the LOB
end

Algorithm 4: The sequence of order placement for the `Postponement model`. This is the input for "Place orders according to strategy" in algorithm 3 for the proposed model.

Appendix F

Complete Postponement model formulation

The complete formulation of the `Postponement model` is presented below. An artificial state variable Y_t^{art} is included to ensure relatively complete recourse. The implication of relatively complete recourse is that the model is solvable in all single-stage sub-problems. Y_t^{art} has been added to the objective function with negative big M coefficient to introduce a large penalty to non-feasible solutions at each sub-stage.

$$Max \quad Z = CVaR_{0,\alpha,\lambda}^{NEST}[R_0(x_0), R_1(x_1), ..., R_T(x_T)] + \sum_{t \epsilon T} M \cdot Y_t^{Art} + \zeta \cdot \eta_d \cdot Y_T \tag{F.1}$$

s.t

$$Y_{t+1} = Y_t + Y_t^{Art} - x_t \quad \text{for } t = 0, 1, \ldots, T-1. \tag{F.2}$$

$$p_{t+1} = p_t + \varepsilon_t \quad \text{for } t = 0, 1, \ldots, T-1. \tag{F.3}$$

$$\varepsilon_t \sim N(0, \sigma^2) \tag{F.4}$$

$$Y_0 = \pi_d^0 \tag{F.5}$$

$$0 \leq x_t \leq \pi_d^0 \quad \text{for } t = 0, 1, \ldots, T. \tag{F.6}$$

$$0 \leq Y_t \leq \pi_d^0 \quad \text{for } t = 0, 1, \ldots, T. \tag{F.7}$$

$$Y_t^{art} \geq 0 \quad \text{for } t = 0, 1, \ldots, T. \tag{F.8}$$

Appendix G

Trading strategy performance for all contracts

An overview of the trading strategy performance for the sets of monthly, quarterly and all traded contracts are presented in chapter 7. The trading strategy performance for all individual contracts is presented in tables G.1 and G.2.

Table G.1: Revenue per MWh and standard deviation for all quarterly contracts

	Quarterly contracts					
Strategy	Q3-2016		Q4-2016		Q1-2017	
	€/MWh	St.dev	€/MWh	St.dev	€/MWh	St.dev
IOBE	27.226	0.22 %	29.831	1.08%	32.139	1.49%
IOBE DHM	27.230	0.01%	29.811	0.39%	31.989	1.31%
TWAP	27.494	1.41%	30.018	1.05%	32.602	0.01%
TWAP DHM	27.405	0.33%	29.991	0.96%	32.568	1.21%
Bertrand	27.492	1.35%	30.011	0.72%	32.639	0.04 %
Bertrand DHM	27.402	0.27%	29.971	0.72%	32.515	1.04%
`Postponement model`	27.408	0.30%	29.960	1.48%	33.168	0.95%

Table G.2: Revenue per MWh and standard deviation for all monthly contracts

	Monthly contracts							
Strategy	JUL-2016		AUG-2016		SEP-2016		OCT-2016	
	€/MWh	St.dev	€/MWh	St.dev	€/MWh	St.dev	€/MWh	St.dev
IOBE	26.346	0.33%	26.959	0.66%	29.080	0.31%	29.281	0.99%
IOBE DHM	26.285	0.26%	27.040	0.61%	29.090	0.34%	29.231	1.09%
TWAP	26.403	0.63%	27.084	0.56%	29.182	0.14%	29.387	1.5%
TWAP DHM	26.342	0.57%	27.147	0.60%	29.210	0.16%	29.366	1.64%
Bertrand	26.405	0.37%	27.076	0.24%	29.183	0.09%	29.391	0.55%
Bertrand DHM	26.345	0.38%	27.140	0.63%	29.206	0.12 %	29.376	0.97%
`Postponement model`	26.350	0.36%	27.140	0.64%	29.207	0.11%	29.376	0.01%

	Monthly contracts							
Strategy	NOV-2016		DEC-2016		JAN-2017		FEB-2017	
	€/MWh	St.dev	€/MWh	St.dev	€/MWh	St.dev	€/MWh	St.dev
IOBE	32.392	3.19%	32.746	1.57%	35.734	1.59%	36.891	0.99%
IOBE DHM	32.458	2.95%	33.077	1.41%	35.888	1.48%	36.901	0.92%
TWAP	32.605	3.24%	32.908	1.59%	35.790	1.85%	36.916	1.78%
TWAP DHM	32.699	2.98%	33.245	1.44%	35.989	1.71%	37.036	1.66%
Bertrand	32.723	1.26%	33.087	0.15%	36.603	0.16%	37.406	0.26%
Bertrand DHM	32.820	0.87%	33.471	0.53%	36.751	0.39%	37.308	0.83%
`Postponement model`	32.838	0.01%	33.473	0.41%	36.962	0.39%	37.528	0.76%

	Monthly contracts					
Strategy	MAR-2017		APR-2017		MAY-2017	
	€/MWh	St.dev	€/MWh	St.dev	€/MWh	St.dev
IOBE	33.153	0.68%	31.025	0.54%	29.273	0.42%
IOBE DHM	33.039	0.57%	30.869	0.39%	29.086	0.53%
TWAP	33.216	0.28%	31.081	0.49%	29.363	0.48%
TWAP DHM	33.112	0.25%	30.936	0.36%	29.150	0.63%
Bertrand	33.559	0.07%	31.091	0.34%	29.374	0.23%
Bertrand DHM	33.450	0.08%	30.978	0.01%	29.158	0.47%
`Postponement model`	33.676	0.06%	31.049	0.23%	29.535	0.44%

NTNU

Norwegian University of
Science and Technology