Market impact modeling and optimal execution strategies for equity trading

An intraday and multiday study

by

F.J.G. Veldman

to obtain the degree of Master of Science Applied Mathematics
at the Delft University of Technology,
to be defended publicly on August 1, 2024.

Student number:	5170478	
Project duration:	January 8, 2024 – Augustus 1, 2024	
Thesis committee:	Prof. Dr. A. Papapantoleon,	TU Delft, supervisor
	Dr. N. Parolya,	TU Delft, second reader
	Lars ter Braak	Robeco, daily supervisor

TU Delft

Abstract

Electronic trading algorithms are at the centre of every buy-side equity trading desk. These algorithms rely often on market impact models, which are stochastic models for the stock prices that account for the feedback effects of trading. Propagator models are central tools for describing the evolution of market impact during and after a trade. This thesis extends the linear propagator model by proposing a new variant that incorporates time-varying liquidity and general decay kernels. To bridge the gap between theory and practice, we use Robeco's proprietary order data base to calibrate the model and validate its performance. The main findings reveal a two-stage decay pattern of market impact, the absence of a single best admissible decay kernel, and a model performance which is in line with our expectations based on the low-signal to noise ratio of the data. The main application of the model in this research is its use in the optimal execution problem, in both a intraday and multiday setting. In an intraday setting we formulate the problem for a risk aware trader and incorporate short-term alpha signals. The discrete analogs of these problems are solved analytically and we highlight significant cost reduction compared to industry benchmarks. In the multiday framework, we quantify the expected cost of trading two adjacent orders and use this to find optimal multiday execution strategies. In a final simulation study we quantify the expected cost of rebalancing two similar investment accounts on consecutive days with a varying number of overlapping stocks. The simulation study accentuate a significant additional cost for the account trading on the second day, which stresses the importance of multiday cost management in rebalancing investment accounts.

Keywords: Algorithmic trading • Market impact modeling • Optimal control • Financial mathematics • Quantitative analysis

Preface

This research is conducted at Robeco, a global asset manager with over 182 billion AuM (per December 2023), and is specialised in the quantitative and sustainable approach to investing. For the past seven months, between January and July 2024, I have been part of Robeco's Equity Trading Research Team in Rotterdam, where I had the opportunity to write my thesis for the partial fulfilment of the requirements for the degree of Master of Science in Applied Mathematics at Delft University of Technology.

As part of Robeco's Super Quant Program and under the supervision of Lars ter Braak, one of the experienced trading researchers, I had the opportunity to apply my theoretical knowledge about mathematics and quantitative finance in real-word scenarios and contribute to Robeco's innovative environment. I want express my gratefulness to Lars for all the insightful discussions and his professional guidance for the duration of the project.

Furthermore, I would like to thank my colleagues in the research and trading team, who made me feel like I was part of the department as from day one, and who graciously shared their expertise and provided assistance when needed. Moreover, I would like to thank Prof. Dr. A. Papapantoleon for being my responsible supervisor and Dr. N. Parolya for being part of my graduation committee.

Finally, I would like to thank my fellow Super Quants, friends and family, who supported me unconditionally during the past seven months.

F.J.G. Veldman
Rotterdam, July 2024

Contents

1

Introduction

Trading large orders with minimum market impact is one of the major difficulties buy-side trading desks are dealing with every day. Robeco is one of these buy-side firms and with over hundred billion euros AuM in equities, orders arriving at its equity trading desk can be of significant size. Because these orders can comprise up to 30% of the daily traded volume, the resulting impact on the stock price can be notable, which makes effective management of market impact a crucial factor in the profitability of its investment strategies. To reduce market impact, traders split up the order into smaller pieces, so called child orders, that are gradually traded throughout a certain time interval. The question, academics and practitioners are trying to solve for many years, is how to allocate an optimal proportion of the total order to each individual child order such that the market impact is minimized: ”the optimal execution problem”.

Electronic trading algorithms are at the centre of every equity trading desk. However, it is in many cases not the trader who decides how to split the order into child orders, but the trading algorithm selected by the trader. The mathematical foundation for many of these trading algorithms is a market impact model, which is a stochastic model for the stock price that takes into account the feedback effects of trading. In the first generation of market impact models, the stock price is only affected in two ways: a temporary component which only affects the current trade and a permanent component which represent the lasting change in the price and affect all future trades. These models are based on the seminal papers by Bertsimas and Lo (1998), Almgren and Chriss (1999) and Almgren and Chriss (2001). However, more recent research in market microstructure (see e.g., Bouchaud et al. (2003), Obizhaeva and Wang (2013), Bouchaud et al. (2018)) has shown that temporary impact can further be decomposed into instantaneous impact and transient impact, in the sense that, each trade causes an immediate impact and a subsequent decay which “propagates” across time. This leads to the second generation of market impact models also known as ”propagator models”.

This new wave of market impact models, quantitatively describe the decay of market impact after every trade and how future trades are affected by this. One of the pioneering models in this category is the discrete-time, log-impact model, introduced by Bouchaud et al. (2003). This model is based on empirical observations of market microstructure and further developed by Bouchaud et al. (2009). Conversely, Obizhaeva and Wang (2013) proposed a model that derives its dynamics from a simplified mathematical description of the limit order book. This model is extended by Alfonsi et al. (2008) and Alfonsi et al. (2010) to incorporate nonlinear price impact on a global level. Gatheral (2010) generalizes the before mentioned models with general decay and nonlinear instantaneous impact on a local level such as in the original Bouchaud model.

In this thesis we propose a new variant of the propagator model originally introduced by Gatheral (2010). This new variant is a linear model that combines the concept of time-varying liquidity parameters (see, Cont et al. (2014) and Fruth et al. (2014)) and general forms for the market impact decay into one model. The underlying economic idea behind intraday time-varying liquidity is that there is generally more liquidity available in the end of the trading day then there is at the middle of the day. Therefore, trading 100 shares just before the close makes less market impact than during lunch time when volumes are the lowest. The time-varying liquidity parameter we use is the intraday volume curve such that the resulting model is a linear propagator model on the traders participation rate. To ensure the viability of the model, we derive some conditions in discrete and continuous time such that the model is free of price manipulation in the sense of dynamic arbitrage (see, Huberman and Stanzl (2004) and Gatheral (2010)).

A unique aspect of this thesis, is that we have access to Robeco's proprietary intraday order database. Up until today, academics have faced a significant challenge in empirically evaluating market impact models due to the lack of access to proprietary order data. Consequently, much of the existing literature is focused on analyzing market impact models using public trading data (see, Webster (2023) and references therein). The availability of Robeco's data allows for a thorough analysis and validation of the new introduced linear propagator model specifically applied to its US orders, thereby bridging the gap between theoretical models and their practical implementation. Notably, Capital Fund Management (CFM), a France quantitative asset managers, has provided some insights into this area with the papers Hey, Bouchaud, et al. (2023) and Hey, Mastromatteo, et al. (2023), but comprehensive access to proprietary data remains rare. We further provide a detailed methodology to calibrate the model to intraday order data, where we leverage the method proposed by Neuman et al. (2023) but adapt it to our own model specifications.

Building on the calibrated linear propagator model, we turn to the optimal execution problem for a single order. We formulate the associated optimal control problem for a deterministic trading strategy as a mean-variance optimization and extend the objective function to include a short-term alpha signal, modeled as an Ornstein–Uhlenbeck process. For the discrete-time versions of these optimal execution problems, we derive analytical solutions using the Lagrange multiplier method. These theoretical results are illustrated through realistic examples, in which we execute an order and compare the performance of different strategies in terms of expected cost and expected impact. These examples demonstrate the practical applicability of the model in realistic trading scenarios.

The final contribution of this thesis extends the intraday optimal execution setup to a multiday framework, a relatively new research direction with limited existing results (see, Harvey et al. (2022) and Bordigoni et al. (2022)). This extension is primarily motivated by the autocorrelation observed in Robeco's metaorders. Robeco manages multiple accounts for each investment strategy, necessitating the rebalancing of accounts within the same strategy on consecutive days. As a result, the same stock may be traded for different accounts over successive days. To quantify it implications, we derive the expected cost of trading two adjacent orders under the linear propagator model. This not only allows us to calculate optimal multiday execution strategies but more importantly enables us to quantify the hidden costs of rebalancing similar accounts on two consecutive days. These theoretical results are demonstrated through a realistic simulation study.

The outline of this thesis is as follows. Chapter 2 delves into the fundamentals of algorithmic trading for asset managers, starting with an exploration of electronic markets and limit order books. It covers empirical analysis of trading data, including data cleaning, order flow analysis and continues to preliminary concepts of market impact with an empirical study on the concavity and transient nature of market impact. Chapter 3 introduces the propagator model for market impact, beginning with Bouchaud's model, followed by the Obizhaeva and Wang model, before we present the new variant of linear propagator model. We moreover discuss the viability of the model by deriving the set of admissible kernels, and outline in detail the calibration process. Numerical results on calibration, decay kernel estimates, and performance evaluation are also provided.

Furthermore, Chapter 4 focuses on optimal intraday execution strategies, formulating the optimal execution problem for the Obizhaeva and Wang model and linear propagator model and discussing the resulting optimal strategies. It explores minimizing execution costs, mean-variance optimal strategies, and strategies incorporating short-term alpha signals. Chapter 5 extends the discussion to optimal multiday portfolio rebalancing, examining the expected cost of trading adjacent metaorders, formulating optimal multiday execution strategies, and addressing the complexities of rebalancing portfolios over multiple days. In the conclusion we summarize the key findings and contributions of the thesis, while in the discussion section we reflect on the implications of the research and potential future directions.

2

Preliminaries on algorithmic trading

We start this chapter by introducing general concepts such as electronic equity markets and the limit order book before we look at some empirical properties of Robeco's proprietary order database. We define filters to enhance the data quality and look for statistical patterns in the order flow. Furthermore, we introduce preliminary concepts on market impact and look at some stylized facts in the data.

2.1. Electronic equity markets and limit order books

Electronic equity market functions through sophisticated systems where orders get matched and executed without human intervention. The shift towards automation is driven by the introduction of electronic communication networks, algorithmic trading and high-frequency trading. These advancements facilitate the rapid execution of orders and provide enhanced liquidity, which results in more dynamic and responsive markets. However, these advancements have also led to a shift from traditional centralized exchanges, such as the New York Stock Exchange and NASDAQ, to alternative trading venues, such as dark pools and crossing networks. These venues are introduced to facilitate the execution of large amounts of shares anonymously, with the aim on reducing market impact and protecting trading intentions.

For buy-side firms such as asset managers, the shift to electronic trading has brought many advantages. Firstly, it allows them to execute orders more efficiently and discreetly by using dark pools and crossing networks. Secondly, they are able to deploy algorithmic trading strategies, which allow for a more systematic and cost-efficient approach to order execution. Common algorithmic trading strategies used by asset managers are Volume-Weighted Average Price (VWAP), and Implementation Shortfall (IS) strategies. A VWAP strategy aims to execute orders in line with the historical/predicted volume curve. IS strategies, also known as arrival price strategies, aim to minimize the difference between the decision price (price at the time of order placement) and the final volume weighted execution price. The IS strategy trades more aggressive compared to VWAP to balance the risk of price movements and market impact. Besides these industry standard strategies, asset managers might also employ more sophisticated algorithms that dynamically adjust order sizes based on short-term predictions of the stock price known as alpha signals or exploit different trading venues to capture favorable prices and volumes.

Orders in electronic equity markets can broadly be categorized into two types: market orders and limit orders. On the one hand, market orders are executed immediately at the current best bid/ask price, ensuring immediate execution but with the risk of price slippage by crossing the spread. On the other hand, limit orders specify a price at which the trader wants to buy or sell, which provides more control over the execution price but without the guarantee of immediate execution. These orders are routed to one of the many trading venues where buy and sell orders are matched through their matching engines and order books. The order book is a real-time, continuously updated list of buy and sell orders, and displays the quantities and prices at which market participants would like to trade. The matching engine prioritizes orders based on price and timing, where they often use a first in first out method.

Quotes in the Limit Order Book (LOB) are essentially limit orders containing the price, direction (buy/sell), and quantity. Figure 2.1 illustrates a LOB, with bid (buy) quotes on the left and ask (sell) quotes on the right. The highest bid and lowest ask prices are known as the best bid and best ask prices, with the difference called the bid-ask spread and their average is referred to as the mid-price. A trade in a LOB exchange occurs when a trader submits a market order, which contains the size and direction. The market order is than matched with existing limit orders in the LOB.

When a buy market order is placed, multiple outcomes are possible based on the order size of the market order compared to the volume present on the best ask price:

- If the order size is less than volume quoted on the best ask, the order matches with limit orders at the best ask price, reducing the volume at this level without changing prices.
- If the order size is equal to the volume quoted on the best ask, the order fully matches all limit orders at the best ask price, increasing the best ask price afterward.
- If the order size exceeds the volume quoted on the best ask, the order fully matches all limit orders at the best ask price, and the remaining order size trades at the next best ask price, resulting in an average price higher than the initial best ask.

This means that a market order can not only raise the price for future trades but also impact the average price of the current trade by targeting limit orders deeper in the book.

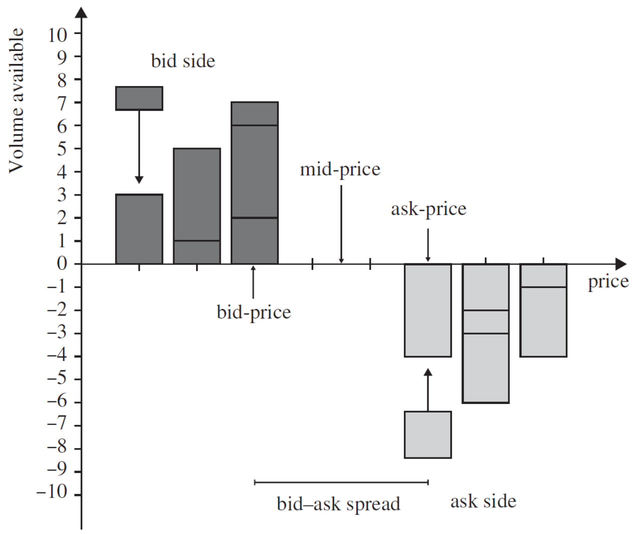

Figure 2.1: Illustration of the limit order book. Source: Bouchaud et al. (2018), Figure 3.1.

2.2. Empirical analysis of trading data

The empirical analysis of trading data is a critical component in the study of market impact and the development of market impact models. Therefore, in this section we have a closer look at some properties of Robeco's trading data. We define some filters to improve the quality of the data and we investigate some stylized facts of order flow.

2.2.1. Data cleaning

The dataset consist of all US equity metaorders executed by Robeco between June, 2020, and September, 2023. With a metaorder we mean an order of one stock that is executed within one day. In particular, for every metaorder we know the aggregated signed fill quantities on 5 minute frequency. To streamline the analysis, we refer to the aggregated signed fill quantity on a 5 minute frequency as a child order and every 5 minute interval we call an intraday time bin. Furthermore, the dataset is structured in such a way that we only have order data and market data on an intraday level without any link to subsequent days.

The dataset includes Robeco's US orders, and only involve stocks that are part of the MSCI US index. The MSCI US index includes approximately 600 stocks and is designed to measure the performance of the large and mid-cap segments of the US market. As of January 2024, this index represents around 85% of the market capitalization in the US (see MSCI Inc. (2024)).

To define filters for improving the data quality, we summarize the average trading behaviour in Figure 2.2 by making several diagnostics plots. In the upper left figure we have a histogram of the length of different metaorders, where we define the length by the difference in time bins between the first and the last trade. In the upper right figure we see a histogram of the number of child orders per metaorder. In the lower left corner we have a plot which displays the amount of trades in each intraday time bin. Lastly, in the lower right corner we see a boxplot of the of the intraday 5 min volatility in bps.

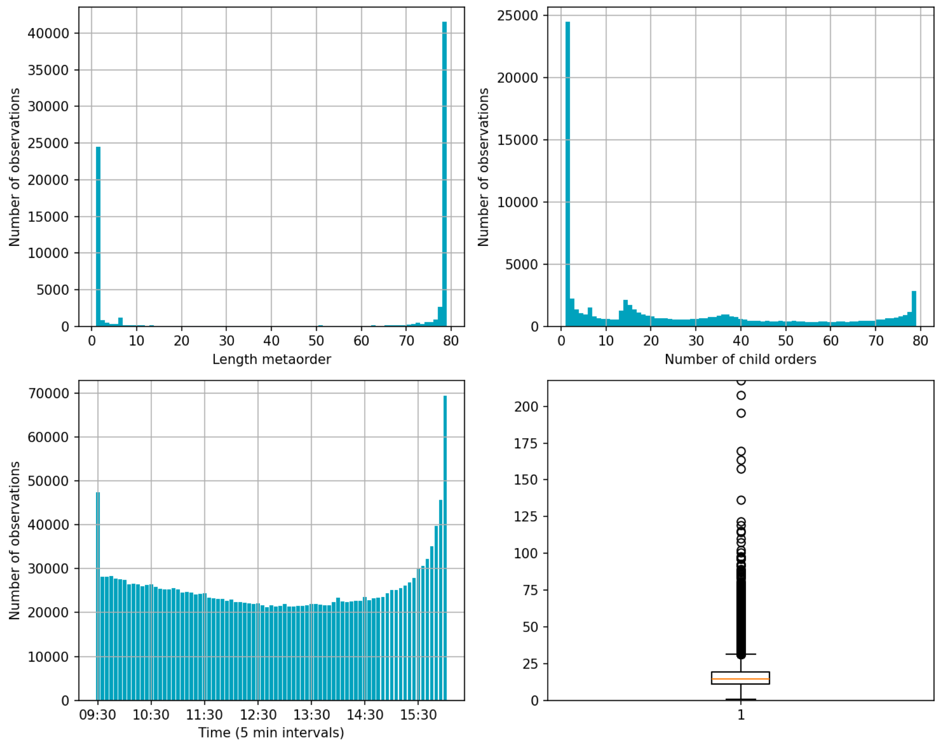

Figure 2.2: Before data filtering, number of metaorders equal to 80393. (Upper left) histogram of the length of every metaorder. (Upper right) histogram of the number of child orders per metaorder. (Lower left) amount of child orders in each intraday time bin. (Lower right) boxplot of the intraday 5 min volatility in bps.

In the boxplot of Figure 2.2 we see that there are many outliers in the intraday volatility estimates. This could be due to some measurement errors or due to some very extreme market conditions. Either way, these very high intraday volatility estimates do not represent normal market behaviour and adds noise to the data sample. Therefore, we exclude all trading days from the data sample where the intraday 5 min volatility is above 70 bps. This corresponds to a daily volatility of 6.2%, which equal to 98.9% annualized.

Furthermore, we see in the upper right plot that we have a large number of metaorders (≈ 25000) which only have a length of one. These metaorders are mostly corresponding to cash flows. Cash flows are orders that are executed because a client wants to withdraw cash from its account. The mandate for these cash flow orders is that they need to be executed as close to the closing price as possible, because the client always receives the closing price. As a result, these cash flows are most of the time executed during the closing auction, which means that we can only observe them in one time bin. This adds noise to the data and therefore we decided to exclude them from the data sample by setting the minimum length of an order to 3 time bins (15 minutes). Note however that this is a proxy since we can not flag the cash flows individually.

Moreover, we filter out all metaorders which are in size smaller than 0.01% of the daily volume of the stock. We do this to reduce the noise in the data set. To summarize, the filters we use to obtain the final data set are given in Table 2.1.

Table 2.1: Data filters used to improve the quality of the data.

Sort filter	Filters
Number of tradings bins (T)	79 (5 min)
Maximum intraday 5 min volatilty	70 bps
Minimum length of order	15 min
Minimum ordersize per metaorder	0.01%
Number of metaorders	48005

The data filters given in Table 2.1 result in the diagnostic plots in Figure C.1 and can be found in Appendix C. In this figure we visualize the average trading behaviour after applying the filters. The original data sample included 80393 metaorders, this means that our current sample includes around 60% of the original data. For all our analysis in this thesis we use the cleaned data set.

2.2.2. Analysis of order flow

Empirical studies have shown that returns in financial markets exhibit minimal or no autocorrelation, implying that price movements are largely random and not easily predictable based on past returns. This observation is consistent with the Efficient Market Hypothesis (see, Fama (1970)), which states that stock prices fully reflect all available information. However, for the order flow of institutional investors this is different.

To investigate the presence of autocorrelation in the order flow at a child order level, we analyze two key dimensions: trade signs (i.e., buy, sell or no trade) and the number of shares traded. The trade sign is an integer indicator that captures the direction of the child order, while the number of shares traded quantifies, the child order's magnitude. By examining these dimensions, we aim to find statistical patterns in the trading behaviour.

For this analysis, we construct time series of child order signs (1 for buy, 0 for no trade and -1 for sell) and number of traded shares per child order. The autocorrelation function (ACF) is then computed for these time series separately. By averaging the autocorrelation functions across all metaorders, we capture the average order flow dynamics at the child order level.

In the left and middle plot of Figure 2.3, we compute the average autocorrelation functions for trade signs and trade volumes over a random sample of metaorders[1]. In the most right plot we show the autocorrelation function of the corresponding log returns.

The ACF for trade signs shows minimal autocorrelation. Almost all lags fall within the confidence interval, except lag one. A possible reason for this could be that we are working with integer values, 1 for buy, 0 for not trading and -1 for selling. Consequently, if we do not trade in every interval and the distribution of trading and not trading during an metaorder is somewhat random, we get this result.

In contrast, the ACF for the number of shares traded per child order shows more persistence, with significant autocorrelation in first few lags. This aligns with our earlier observation that buy-side firms often break metaorders into smaller pieces to mitigate their market impact. Note that the ACF decays like a power-law or exponential function. This is an important observation as we will see in Chapter 3.

Additionally, the analysis includes the ACF of log returns, which we plot in the right hand side of Figure 2.3. This ACF demonstrates no significant autocorrelation for lags larger than 1, aligning with the widely accepted notion that log returns follow a random walk behaviour and are not easily predictable. However, at the first lag, we notice some autocorrelation indicting a short momentum in the returns on a 5 minute frequency.

[1] Different random samples do not significantly change the ACFs.

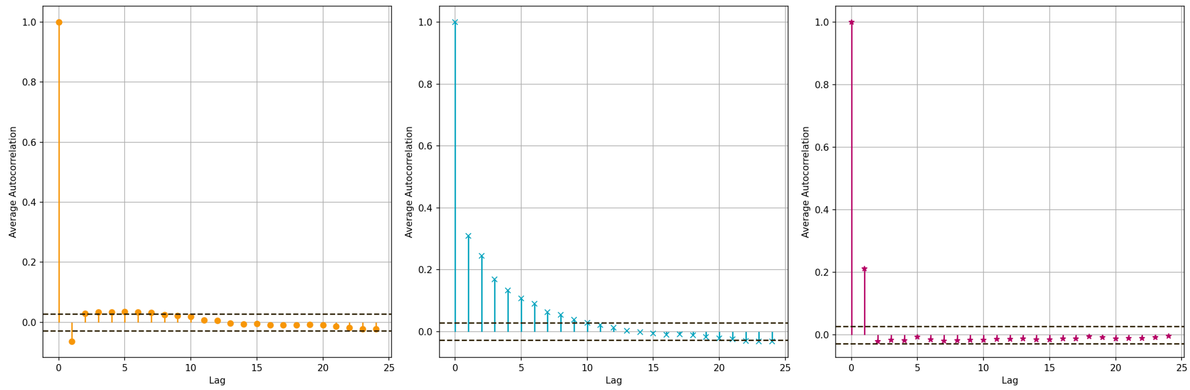

Figure 2.3: (Left) average autocorrelation function of order signs child orders. (Middle) average autocorrelation function of number of shares child orders. (Right) average autocorrelation function of log returns. All autocorrelation functions are based on a random sample of 6000 metaorders.

2.3. Preliminaries on market impact

The economic definition of market impact is that trading causes adverse price movements in the stock price that otherwise not would have happened. In a simplified setting, we can decompose the mid-price S of a stock, into:

$$S = I + P,$$

where I is the market impact caused by trading and P is the unobserved or unaffected price process. The unobserved price process is driven by the action of other market participant and external factors. An illustrative example of this decomposition is given in Figure 2.4. In this example we display the execution of a 2.3% Average Daily Volume (ADV) buy order in a certain stock and the resulting market impact calculated using a market impact model. The amount of shares traded in every 5 minute interval of the day is given in the the upper figure. The resulting market impact, i.e. $I = S - P$, the unaffected price and mid-price are given in the lower figure.

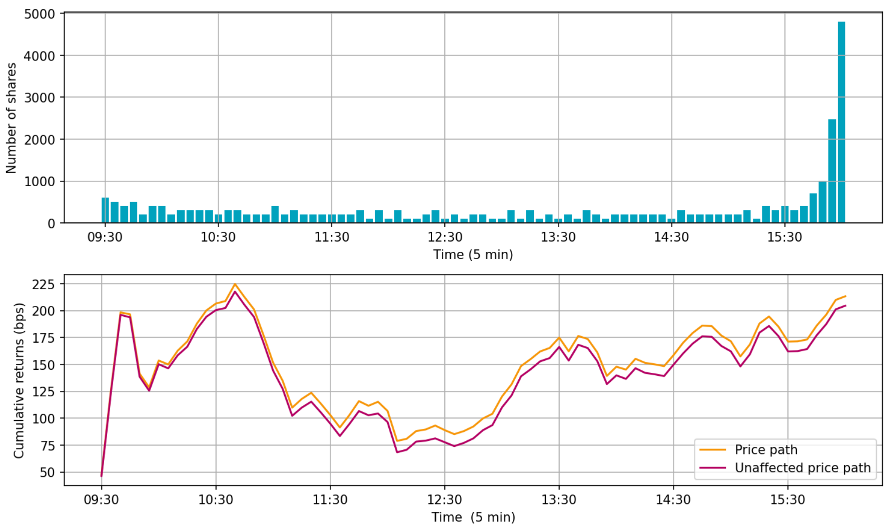

Figure 2.4: (Upper figure) trading schedule in number of shares during an 2.3% ADV buy order. (Lower Figure) illustration of the decomposition of cumulative log returns in the mid-price and the unaffected price path both given in basis points.

Accurately modeling market impact is critical for asset managers for several reasons. Primarily, it allows for better understanding how trading affects price returns. This way traders can decompose trading cost into several factors such as market impact cost, alpha cost, spread cost and other types of cost. This is key in evaluating the effectiveness of trading strategies and identifying areas for improvement. Secondly, by understanding market impact, traders can design optimal execution strategies that minimize market impact when trading large positions. Lastly, regulatory frameworks require asset managers to demonstrate best execution practices (see e.g. Wagner and Edwards (1993)), and accurate market impact models provide the necessary documentation and analysis to show that trades were executed in a manner that minimizes costs and maximizes value for clients.

In the seminal paper by Almgren and Chriss (2001), the authors categorize market impact into two main types: temporary and permanent. Temporary impact refers to the immediate price movement caused by the execution of the trade. In contrast, permanent impact represents the lasting change in price, reflecting the new information perceived by the market. In reality, the permanent impact is hardest to measure. This is because after the execution of an order the variance scales linearly with time and is amplified by the volatility of the asset, which means that the signal of the order becomes very weak and therefore hard to properly measure (see e.g. Brokmann et al. (2015), Bucci et al. (2018), and references therein).

Bouchaud et al. (2003), Obizhaeva and Wang (2013) and Bouchaud et al. (2009) proposed to split temporary impact further into two categories: instantaneous impact and transient impact. On the one hand, instantaneous impact has no memory effect and only effects the current trade. It is for example caused by crossing the spread and the need to take away liquidity deeper in the order book. On the other hand, transient impact refers to the impact trajectory between instantaneous and permanent. This form of impact decays overtime which is due to the resilience effect of prices as we see later in this section.

To mitigate the adverse effects of market impact, traders often split large metaorders into smaller child orders (see Figure 2.4). This strategy, known as ”order slicing”, offers several advantages. Smaller orders are less likely to move the market significantly because they can more easily be matched with available liquidity, thereby reducing the temporary impact and the likelihood of revealing the traders trading intentions. Additionally, by spreading orders over time, traders can exploit favorable market conditions, i.e. executing child order when liquidity is higher and volatility is lower or take advantage of short momentum.

In the coming two subsections, we empirically investigate two important properties of market impact: its concavity with respect the order size and the transient nature of market impact. These properties are important to consider when making modelling choices as we will see in the next chapter.

2.3.1. Concavity of market impact

Order size is the main driver of market impact. It is well documented that the market impact of a metaorder scales proportional to square root of the order size. (see e.g., Tóth et al. (2011) Bacry et al. (2015), Bucci et al. (2018)). However, in this thesis we are mostly interested in the relation between the order size of a child order and the resulting market impact.

To be able to compare different stocks and days in one analysis, we apply multiple normalizations. We start by normalizing the log returns by subtracting the fair value and divide by the intraday volatility. With fair value we mean the part of the returns we can attribute to market and sector movements and for this we use Robeco's internal fair value model. To normalize the child order sizes by volume we define two types: a child order size normalized by ADV and a child order size normalized by the total volume traded in the associated intraday time bin, also called the participation rate.

The reason why we consider the participation rate, is that the intraday volume curve is not constant over the day. Volume tent to be much higher near the end of the trading day and during the opening as can be seen in Figure 2.5. Consequently, leading to less market impact when trading the same amount of shares near the close compared to other parts of the day. Therefore, to be able compare different child orders equally over the day we divide by the total volume traded by the market in a 5 min interval.

To investigate the impact of a child order on the mid-price, we consider the volatility and fair value corrected average signed log return in a intraday time bin. We bucket the normalized child orders in 1000 quantile buckets and make a scatterplot of their mean. To find a pattern in the scatter plot we fit a power-law and a linear function to the data points with the package *scipy.optimize.curve_fit* in Python. In the left hand side of Figures 2.6 we plot the average signed normalized log returns versus the child

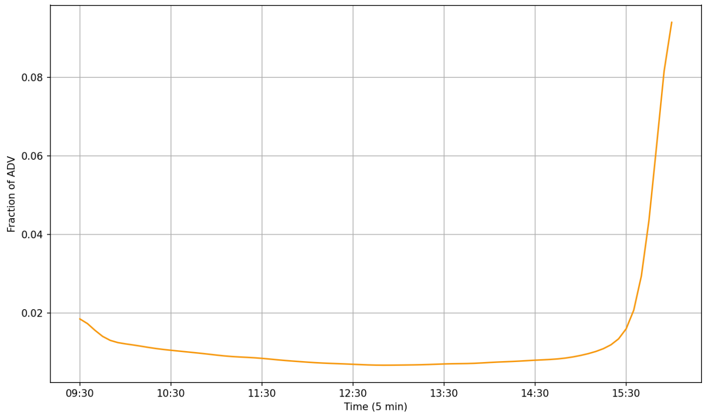

Figure 2.5: Intraday volume curve based on the average volume curves of all the stock in the MSCI US index in 2023.

orders normalized by ADV and in the right hand side figure the average signed normalized log returns versus participation rate on a log-log scale.

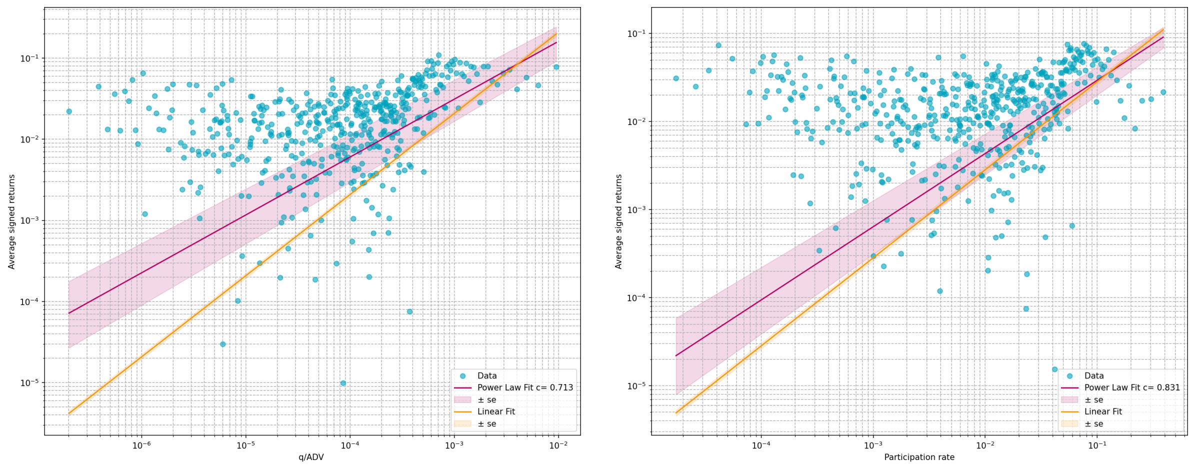

Figure 2.6: Scatterplot of expected signed log return plotted versus the mean of every quantile bucket for different normalizations of the child orders. We use 1000 buckets and the fitted functions are a linear and power law function. (Left) the normalization by ADV. (Right) the normalization by the intraday volume.

We see in the left hand side of Figure 2.6 that a power-law function with concavity parameter 0.71 fits the data best. Conversely, the relation between the average signed normalized log returns and the participation rate is less concave (concavity parameter equal to 0.83) and moves closer to a linear relation. To link this directly to a functional form of how the instantaneous impact scales with the normalized order size is more subtle. Due to the large autocorrelation in order flow, it is hard to obtain the isolated instantaneous impact of a child order without the interference of the impact from previous trades. Therefore, we can only use the relative difference in concavity between the two normalizations in future analysis.

2.3.2. Transient nature of market impact

The transient nature of market impact is observed on a trade level but mostly on a metaorder level. On a child order level this is hard to empirically measure because the large autocorrelation in order flow makes it difficult to isolate the impact of individual child orders. However, on a metaorder level this is easier to measure and visualize.

To visualize the transient property of market impact on metaorder level, we consider the volatility and fair value corrected average cumulative signed log return versus the relative length of a metaorder. To be able to compare different metaorders with different lengths, we use in our analysis the relative length, i.e we set the start of the execution to 0, the end of the execution to 1, such that a relative length of 2 means one order length after the end of the execution. In Figure 2.7 we plot the normalized average cumulative signed log return during the execution up to one order length after the execution. We average over 3350 metaorders with an average duration of 45 minutes. Moreover, all metaorder have similar length.

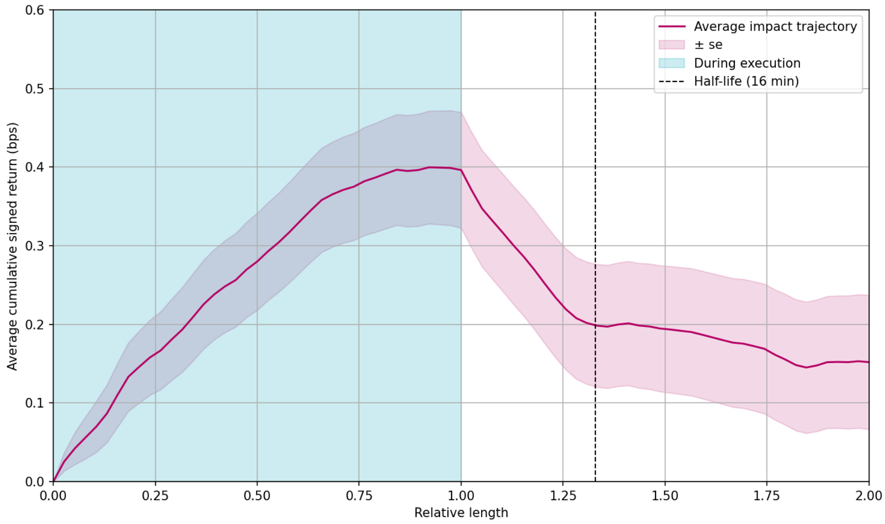

Figure 2.7: Average signed cumulative log returns versus the relative length of a set of metaorders. We use 3350 metaorders with an average duration of 45 min.

In Figure 2.7 we see that the average impact trajectory during the execution is a concave function of time. Directly after the execution of the metaorder the impact reverts and decays like a power-law or exponential function. This clearly highlights the transient nature of market impact after the execution. The figure also includes the half-life of the impact trajectory, defined as the time after execution when half of the maximum impact has reverted. For this sample of orders the half-life is equal to 16 minutes. Additionally, one order length after the end of execution, 38% of the maximum impact remains. This suggests a two-stage decay process: an initial rapid and short decay followed by a slow and long decay.

We must note that this analysis is only based on a small sample of the metaorders ($\approx 7\%$). The reason for this is that for this subset only, we can measure the decay for more than 1 order length after the end of the execution. This is because the data contains only intraday returns and the average order length is 360 minutes.

In summary, we analysed multiple characteristics of order data. We have seen that our order flow shows a large autocorrelation for the first few lags and decays like a power-law or exponential function, while we found little to no autocorrelation in the returns. Furthermore, we have shown how market impact of our child order scales with respect to different normalized order sizes and that the impact trajectory after the execution of a metaorder reverts back over time. To explain and model these empirical observations, propagator models play an essential role as we see in the next chapter.

3

The propagator model

This chapter begins by presenting the mathematical setup of this thesis. Next, we review two important propagator models that serve as foundational elements in the literature. Following this, we introduce the new variant of the linear propagator model which combines a time-varying liquidity process and general decay kernels. We investigate conditions under which the linear model is free of price manipulation and other irregularities. Additionally, we outline a two-step optimization approach to calibrate the propagator model to trading data. Finally, we discuss the numerical results of the calibration.

3.1. Mathematical setup

We start by defining some key processes and common notation that we use throughout the thesis. Let $T > 0$ and consider a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{P})$. We define an admissible trading strategy as follows:

Definition 3.1.1. *(Admissible trading strategy) Define $\Pi = (Q_t)_{t\in[0,T]}$ as a trading strategy that describes the number of shares held by the trader at each time $t \in [0,T]$, where Q_0 is the initial amount of shares traded such that $Q_0 > 0$ for a buy order and $Q_0 < 0$ for a sell order. Its variation dQ describes the trades.*

The trading strategy Q is considered admissible if it satisfies the following conditions:

(a) Q is càdlàg ("right continuous with left limits") and $\mathcal{F}_t$-adapted for all $t \in [0,T]$;

(b) Q has finite and $\mathbb{P}$-a.s bounded total variation on any finite interval $[0,T]$;

(c) There exists $T > 0$ such that $Q_t = 0$ $\mathbb{P}$-a.s for all $t \geq T$.

Furthermore, the strategy can be classified as follows:

- *If $Q \in \mathcal{D}$, then it is a deterministic admissible strategy.*
- *If $Q \in \mathcal{Q}$, then it is a stochastic admissible strategy.*

The practical meaning of the admissibility conditions in the definition above is as follows: The first condition ensures that the trading strategy relies only on information available up to time t and allows for the possibility of jumps. For the second condition, we note that a general trading strategy Q might alternate between buys and sells, and can therefore be decomposed into a nonincreasing sell strategy $(X_t)_{t\in[0,T]}$ and a nondecreasing buy strategy $(Y_t)_{t\in[0,T]}$, such that $Q = X + Y$. To allow for such a representation, Q must have bounded total variation, which results in the second condition.

The last condition states that the trading strategy concludes within a finite time frame, aligning with realistic trading scenarios. Furthermore, the last condition implies that the total variation of Q is bounded by $X_0 + |Y_0|$. The $\mathbb{P}$-almost sure bound on total variation implies that the total quantities of both buy and sell orders are constrained. Given that the number of shares for any stock is finite, this assumption is economically reasonable.

Notice that an admissible trading strategy can exhibit jumps. Therefore, we define a jump as follows:

Definition 3.1.2. *(Jump) Let Q be a process that exhibits jumps. Then we denote a jump by $\Delta Q_t = Q_t - Q_{t-}$, where $Q_{t-} = \lim_{s \uparrow t} Q_s$.*

To ensure clarity, we define the following process spaces:

Definition 3.1.3. *(Process spaces)*

- *Let $\mathcal{Z}$ be the space of continuous martingales.*
- *Let $\mathcal{S}$ be the space of semi-martingales.*

All processes in these spaces are adapted to the filtration $(\mathcal{F}_t)_{t\in[0,T]}$ and take values in $\mathbb{R}$. Note that both $\mathcal{Q} \subseteq \mathcal{S}$ and $\mathcal{Z} \subseteq \mathcal{S}$.

Within this framework we define the following price processes: we define the unobserved price process by $P \in \mathcal{S}$ and the observable mid-price by $S \in \mathcal{S}$ such that the market impact I, caused by trading a strategy Q, is equal to:

$$I = S - P.$$

We should note that the observable price process S depends on the trading strategy Q, while the unobserved price process P is independent of the trading strategy Q. The formulation of the observable mid-price S allows the stock price to become negative with non-zero probability. However, since we only consider short time scales, this non-zero probability is negligible (see e.g., Almgren and Chriss (1999) and Almgren and Chriss (2001) for discussion).

In the most general case, the dynamics of the unobserved price process $P \in \mathcal{S}$ can be decomposed in:

$$dP_t = -d\alpha_t + dZ_t,$$

for all $t \in [0, T]$, where $Z \in \mathcal{Z}$ is a continuous martingale that accounts for noise introduced by other market participants or external factors. Furthermore, $\alpha \in \mathcal{S}$ is the short-term alpha signal of the stock and defined as $\alpha_t = \mathbb{E}[P_T - P_t \mid \mathcal{F}_t]$ for all $t \in [0, T]$. This process is often modeled using an Ornstein–Uhlenbeck process (see Uhlenbeck and Ornstein (1930)) or via machine learning techniques.

In many cases, we start by formulating a model in discrete time before we extend it to continuous time. To make the derivations more comprehensive, we use the following notation in discrete time:

Definition 3.1.4. *(Time discretization) Let $T > 0$ be length of a trading period. Then we discretize time $t \in [0, T]$ by an integer N, such that:*

$$t_n^N = \frac{n}{N}T = n\Delta t^N,$$

and

$$Q_n^N = Q_{t_n^N}.$$

The same notation we use for other processes. Moreover we define the operator Δ_n such that:

$$\Delta_n Q^N = Q_n^N - Q_{n-1}^N,$$

and describes the increments of the discrete variable Q^N.

3.2. Foundational propagator models

In this section, we review two propagator models which serve as foundational elements in propagator model literature: the Bouchaud model introduced by Bouchaud et al. (2003) and the Obizhaeva and Wang (OW) model by Obizhaeva and Wang (2013). The first model is motivated by empirical observations of market microstructure, where the second is based on a simplified mathematical description of the limit order book.

3.2.1. Bouchaud's model

Bouchaud et al. (2003) propose the first version of the propagator model in trade time, a different time measure that counts time in the arrival of market orders. Therefore, in this subsection, we use slightly different notation than introduced in the previous section.

The model is based on the empirical observation that the sign of market orders have long-term autocorrelation and decays like a power-law function and price returns do not (see e.g., Bouchaud et al. (2003), Bouchaud et al. (2006), Brokmann et al. (2015)). Although their observations are on trade level, in Section 2.2.2 we have made a similar observation for the order flow of our child orders on a 5 minute frequency.

Bouchaud's model is based on the reasoning that if the impact of market orders would be constant and permanent, price returns would show significant autocorrelation. This contradicts empirical observations and violates the principle of market efficiency. To illustrate their reasoning, we consider a naive market impact model, in which every trade causes a constant and permanent impact. In this example we follow the reasoning of Bouchaud et al. (2018) which uses a linear version of the model proposed in Bouchaud et al. (2003).

Let's denote n by the arrival of the n-th market order, and define the change in the mid-price between the arrival of two market orders as follows:

$$r_n = S_{n+1} - S_n.$$

Furthermore, assume that each trade has a mean permanent impact of G. Then, the stock price dynamics are:

$$r_n = G\omega_n + \epsilon_n,$$

where ω_n is the sign of the market order and ϵ_n is a noise term, which captures price changes not related to trading (cancellations). We assume that ϵ_n are i.i.d random variables with mean zero and unit variance. Given the mid-price S_0 at some initial time, the mid-price S_n at some further trade time n can be written as:

$$S_n = S_0 + G\sum_{m=1}^{n} \omega_m + \sum_{m=1}^{n} \epsilon_m. \tag{3.1}$$

It is clear from the above model that whatever happened at all previous trade times, is now permanently incorporated in the price at trade time n. To investigate the impact of a trade under this model, we define the response function $\mathcal{R}$. The response function is the expected signed impact on the mid-price and is given by:

$$\mathcal{R}(h) := \mathbb{E}[\omega_n \cdot (S_{n+h} - S_n)],$$

where h is an arbitrary lag of trade time. Then when the signs of the trades are independent random variables with mean zero and unit variance, the authors show that the naive model predicts that the expected impact at lag h is constant and therefore permanent as is expected:

$$\mathcal{R}(h) = G.$$

Moreover, the authors show that the mid-price under this model is a diffusion process as is observed in financial markets for short lags h. However, this naive model ignores the important empirical observation that the order-sign series ω_n are strongly autocorrelated. If we incorporate the autocorrelation structure of order signs, i.e. $\rho(h) := \mathbb{E}[\omega_n \omega_{n+h}]$, into the naive model, we get the following autocorrelation of the returns:

$$\mathbb{E}[r_n r_{n+h}] = G^2 \rho(h),$$

for arbitrary $h > 0$. This clearly shows that the returns become autocorrelated as well since it depends on $\rho(h)$. This does not comply with empirical observations and violates the market efficiency principle, because this makes price returns predictable.

Therefore, Bouchaud et al. (2003) propose a generalisation of the naive model. In this model, the impact of market orders is not constant and permanent, but rather a function of trade time . This describes how the impact of a market order propagates through time. Consequently, the authors propose the following model for the stock price:

$$S_n = S_0 + \sum_{m=1}^{n} G(n-m)\omega_m \ln(Q_m) + \sum_{m=1}^{n} \epsilon_m,$$

where $G(\cdot)$ is the kernel function and describes the time dependent impact of a market order. This function is assumed to be fixed and non-random that only depends on the time differences. Furthermore, $\ln(Q_n)$ is the natural logarithm of the aggregated fill quantity at trade time n. This term comes from their observation that the impact scales with the natural logarithm of the aggregated fill quantity.

To find a suitable functional form for the kernel function, the authors calculate the response function under this model for an arbitrary lag $h > 0$, and show that it is equal to:

$$\mathcal{R}(h) = \mathbb{E}[\ln(Q_m)]G(h) + \sum_{m<h} G(h-m)\rho(h) + \sum_{m>0} [G(h+m) - G(m)]\rho(h).$$

The above computation reveals that the response function's behavior is influenced by the ACF of order signs (see dependence on $\rho(h)$). The authors show that when the ACF of the order signs decays following a power-law, the response function grows proportional to $h^{1-2\gamma}$. Such a growth pattern implies a significant amplification of the impact as the lag h increases, which contradicts the anticipated decrease in impact over time.

Therefore, the authors propose that the kernel $G(\cdot)$ should decay over time to counterbalance the amplification effect induced by the order sign's ACF. To ensure this counterbalancing effect, the authors suggest a power-law decay as kernel:

$$G(t-s) = \frac{\zeta}{(\zeta + t - s)^{\gamma}}, \qquad 0 < \gamma < 1 \quad \text{and} \quad \zeta > 0,$$

They conclude that if the kernel function $G(\cdot)$ decays according to the above function, the mid-prices under this model display diffusive behavior, aligning with the principle of market efficiency over short timescales.

3.2.2. The Obizhaeva and Wang model

The original Obizhaeva and Wang (OW) model is proposed in the paper Obizhaeva and Wang (2013), however we summarize the model as presented in Webster (2023) because it gives a short and concise description of the model.

In contrast with Bouchaud's propagator model, the OW model is a continuous time market impact model that can be used to model the impact of only one trader. Therefore, we consider a stochastic admissible trading strategy $Q \in \mathcal{Q}$. In addition to the processes defined in Section 3.1, we define $\tilde{S}$ as the execution price. It includes the market impact as well as instantaneous transaction cost and we define the instantaneous transaction cost by $\bar{S} = \tilde{S} - S$.

The OW model is based on some market micro-structure assumptions of the limit order book:

(a) A trader only executes an order at the bid or the ask price. This means that, we ignore the bid-ask spread and denote S by the observed ask price. The model is analog of trading at the bid.

(b) The order book is proportional to the Lebesgue measure when going deeper into the limit order book. This implies that the limit order book is block shaped, i.e. every quoted price on the ask(bid) has the same volume.

To introduce the OW model, we first define the model in discrete time using the notation in Definition 3.1.4 before we extend it to continuous time. We start by taking a detailed look at what happens when a trader executes a fill under these market microstructure assumptions.

Let $\frac{1}{\lambda}$ be a positive constant and equal to the volume quoted on the ask. When a trader executes an order of size $\Delta_n Q^N > 0$ at time t_{n-1}^N, it walks trough the limit order book before reaching the price:

$$S_{n-1}^N + \lambda \Delta_n Q^N.$$

Then the average execution price per share is:

$$\tilde{S}^N_{n-1} = S^N_{n-1} + \frac{\lambda}{2}\Delta_n Q^N.$$

This means that the linear instantaneous transaction cost are equal to $\tilde{S}^N_n = \frac{\lambda}{2}\Delta_n Q^N$. The order book remains proportional to the Lebesgue measure for all prices above the lowest ask price and the new ask price becomes:

$$S^N_n = S^N_{n-1} + \lambda \Delta_n Q^N.$$

We assume that $\Delta_n S^N \Delta_n Q^N \to 0$, such that the trade does not adversely select the limit order book. We illustrate an order execution under the OW model in Figure 3.1.

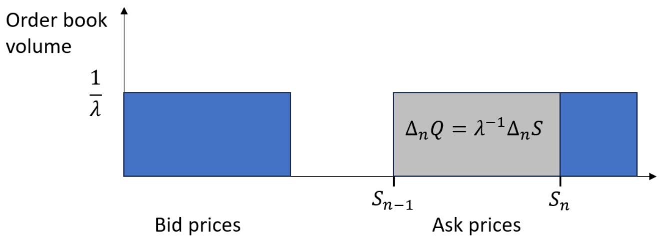

Figure 3.1: Illustration of the execution of a fill under the OW model.

Finally, Obizhaeva and Wang (2013) assume that the price dislocation on the ask(bid) reverts exponentially over time such that the market impact becomes:

$$\Delta_n I^N = -\beta I^N_{n-1}\Delta t^N + \lambda \Delta_n Q^N.$$

The authors show that when $N \to \infty$, the above difference equation converge uniformly in probability (see, Protter (2005)) to the continuous time Stochastic Differential Equation (SDE):

$$dI_t = -\beta I_t dt + \lambda dQ_t.$$

For constant values of the liquidity parameters $\beta, \lambda > 0$, we can solve the SDE explicitly. We solve the SDE using the integrating factor method and assume without loss of generality that $I_0 = 0$. Using the integrating factor $e^{\beta t}$, we find that the solution is equal to:

$$I_t = \int_0^t \lambda e^{-\beta(t-s)} dQ_s.$$

This means that we can write the observable stock price under OW model as:

$$S_t = \int_0^t \lambda G(t-s) dQ_s + P_t,$$

with $G(t-s) = e^{-\beta(t-s)}$ and where P is the unobserved price process.

One of the major differences between the OW model and Bouchaud's original propagator model is the scaling of the traded quantity. In the OW model, instantaneous market impact scales linearly with traded quantity and logarithmically in Bouchaud's model.

Later Bouchaud et al. (2009) review empirical properties of Bouchaud's propagator model and find evidence that the impact of the traded quantity scales as a power-law function $f(x) \propto x^c$ with $c \in [0.2, 0.5]$. This model is also referred to as the locally concave Bouchaud model.

The locally concave Bouchaud model with an exponential kernel can be written as the following difference equation:

$$\Delta_n I^N = -\beta I^N_{n-1}\Delta t^N + \lambda f(\Delta_n Q^N),$$

where $f(\cdot)$ is a differentiable function and concave on $[0, \infty)$.

The transformation above suggest that we by choosing different combinations of the instantaneous impact function $f(\cdot)$ and the decay kernel $G(\cdot)$, we can find different forms of the propagator model. This leads to a general version of the propagator model which we introduce in the next section.

3.3. The linear propagator model

The Bouchaud and OW propagator models have a similar structure because in both cases the impact of a trade propagates trough time by means of a decay kernel. However, their are also some notable differences. In this section, we bring these two models together by introducing a general version of the propagator model. The general model serves as foundation for the new variant of the linear propagator model we propose here. In this model we use clock-time and we consider the actions of only one trader.

We assume in this section that there is no short-term alpha signal, such that for all $t \in [0, T]$, the unobserved price is given by:

$$P_t = S_0 + Z_t,$$

where $Z \in \mathcal{Z}$ is a continuous martingale. This means that, as long as the trader does not participate in the market, the prices are determined by the actions of other market participants or external factors.

To define the general propagator model, we consider a stochastic admissible trading strategy $Q \in \mathcal{Q}$ but in addition we require the strategy to be continuous, i.e. the strategy can have no jumps. Furthermore, we let $\dot{Q}$ represent the trading rate (e.g., 20 shares per second), such that $dQ_t = \dot{Q}_t\, dt$. Later, when we introduce the new variant of the linear propagator model, allow the admissible trading strategy to have jumps again.

Using this mathematical framework, we are able state the general propagator model introduced by Gatheral (2010):

$$S_t = S_0 + \int_0^t f\left(\dot{Q}_s\right) G(t-s) ds + Z_t, \tag{3.2}$$

where $f(\dot{Q}_t)$ is the instantaneous market impact function of a trade $\dot{Q}_t$ at time t and $G(t-s)$ represents the decay of the market impact after the execution of every trade. We prefer to present the model in this form because writing it as a SDE is mathematically involved.

The general propagator model is a generalisation of Bouchaud's model and the OW model. If we let $f(\cdot)$ be a differentiable and concave on the interval $[0, \infty)$ and $G(t-s) = \frac{\zeta}{(\zeta+t-s)^{\gamma}}$, then we find a continuous time variant of locally concave propagator model:

$$S_t = S_0 + \int_0^t \frac{\zeta}{(\zeta + (t-s))^{\gamma}} f(\dot{Q}_s) ds + Z_t.$$

However, if we let $f(\dot{Q}_t) = \lambda \dot{Q}_t$ and $G(t-s) = e^{-\beta(t-s)}$, we retrieve a variant of the OW propagator model[1]:

$$S_t = S_0 + \int_0^t \lambda e^{-\beta(t-s)} dQ_s + Z_t,$$

for $\lambda, \beta > 0$. Furthermore, Gatheral (2010) shows that the general propagator model is also a generalisation of the model proposed by Almgren et al. (2005) .

Notice that the integral representation of the market impact I can be interpreted as a convolution between the functions $f(\cdot)$ and $G(\cdot)$:

$$I_t = [f * G](t) = \int_0^t f(\dot{Q}_s) G(t-s) ds,$$

[1] This version does not allow for jumps in the trading strategy while the origional OW model does.

where $*$ denotes the convolution operator. Using this representation, market impact in the propagator model can be viewed as the accumulated effect of all previous trades. By "accumulated," we mean that the impact of individual trades can be superimposed through convolution to form the total market impact of an order. For an illustrative example see Figure 3.2. This figure illustrates how the impacts of individual child orders are stacked on top of each other and how the impact decays after each child order.

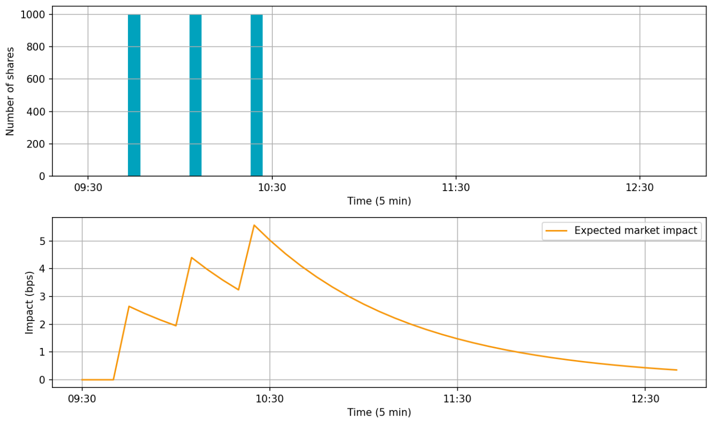

Figure 3.2: (Upper figure) an arbitrary trading strategy in number shares. (Lower figure) the expected market impact during and after an order calculated using the OW model.

As mentioned earlier the market impact I can be decomposed in multiple components. We distinguish between 3 types of market impact:

- *Instantaneous impact*: the instantaneous impact of a trade at time t is given by $f(\dot{Q}_t)G(0)$. This type of impact only affect impact cost of the trade at time t and not any subsequent orders.
- *Permanent impact*: the permanent impact is given by $f(\dot{Q}_t)G(\infty)$, where $G(\infty) := \lim_{t\uparrow\infty} G(t)$. This type of impact affects all trades equally.
- *Transient impact*: everything between instantaneous and permanent impact is called transient market impact.

Gatheral (2010) shows that the general propagator model allows for price manipulation when the instantaneous impact function $f(\cdot)$ is non-linear in the trading rate. We formally introduce the concept of price manipulation in the next section but a market impact model that permits price manipulation can be seen as an asset-pricing model which allows for arbitrage opportunities, making it essential to exclude. Furthermore, the general propagator model does not have tractable solutions in the optimal execution problem (see e.g., Curato et al. (2017)). Therefore, we shift our attention to linear models in the trading rate.

In Figure 2.6 of Section 2.3.1, we observed that child orders normalized by the intraday volume curve exhibit a more linear relationship compared to those normalized by the average daily volume. This observation supports the use of the participation rate in a linear model. Furthermore, normalizing by the intraday volume curve allows the model to account for the varying liquidity throughout the trading day, a critical factor influencing market impact.

Consequently, we propose a new variant of the linear propagator model by incorporating a time-dependent liquidity process. Therefore, we introduce the liquidity process $\Theta = (\Theta_t)_{t\in[0,T]}$ such that $f(\cdot) \propto \Theta$. We assume the liquidity process to be deterministic, positive, continuously differentiable and equal to:

$$\Theta_t = \frac{\sigma \cdot \kappa_t}{V_t} = \frac{\sigma \cdot \kappa_t}{\text{ADV } \cdot v_t}, \qquad \text{for all } t \geq 0,$$

where σ is the intraday volatility, and v_t the intraday volume curve.

In the linear model, we assume that the admissible trading strategy is a stochastic process $Q \in \mathcal{Q}$ which can exhibit jumps as defined in Definition 3.1.2. We define the new variant of linear propagator model with a time-dependent liquidity process and general decay kernel as follows:

$$S_t = S_0 + \int_0^t \Theta_s G(t-s) dQ_s + Z_t, \tag{3.3}$$

where $G : [0,\infty) \to [0,\infty)$ is a measurable function serving as the decay kernel and $G(0)$ is equal to the instantaneous impact of a trade. We assume that $G(\cdot)$ is bounded, to ensure that the integral does not diverge. Note that, $\frac{dQ_t}{V_t}$ is the traders participation rate for all $t \in [0,T]$.

Lets have a closer look at some properties of the integral representation of the market impact in Equation (3.3). Given that $Q \in \mathcal{Q} \subseteq \mathcal{S}$, it can be represented as a semi-martingale. However, since we assumed that Q has finite and $\mathbb{P}$-a.s bounded total variation, the local martingale part of Q vanishes. Consequently, the integral representation of the market impact is actually a Lebesgue-Stieltjes integral, which is deterministic.

Given that the data is one a 5 minute frequency and certain derivations are more straightforward in discrete time, we also define the discrete time version of the linear propagator model with a time-varying liquidity process using the notation outlined in Definition 3.1.4:

$$S_n^N = S_0^N + \sum_{m=1}^{n} \Theta_m^N G(n-m) \Delta_m Q^N + Z_n^N, \tag{3.4}$$

where S_n^N is the mid-price just before time t_n^N and Θ_n^N represents the liquidity in the interval $[t_{n-1}^N, t_n^N)$.

3.4. Admissible set of kernels and liquidity processes

In this section, we derive admissible sets for the decay kernel $G(\cdot)$ and the liquidity process Θ using no-arbitrage principles. In the propagator model we can influence the stock price by trading in the stock. Therefore, market participants might be able to manipulate the price in their favour by trading certain strategies. In market impact literature this is also known as price manipulation (see Huberman and Stanzl (2004)) and it is important to exclude. Before we formally introduce price manipulation, we derive the expected cost of trading under the linear propagator model.

To define the cost of an admissible trading strategy under the linear model, we first define the cost for a continuous strategy. For this we use the implementation shortfall originally proposed by Perold (1988). This means that the cost of a continuous trading strategy without jumps is equal to:

$$\int_0^T (S_t - S_0) dQ_t = \int_0^T \int_0^t \Theta_s G(t-s) dQ_s dQ_t + \int_0^T Z_t dQ_t$$

Since $Q \in \mathcal{Q}$, we allow the trading strategy to have jumps. Therefore, we also need to account for the cost of these jumps. We assume that when there is a jump trade of size ΔQ_t the price moves from S_{t-} to $S_t = S_{t-} + \Theta_t \Delta Q_t G(0)$, where $G(0)$ is the instantaneous impact. From this it follows that the average execution price of a single jump trade is then equal to:

$$S_{t-} + \frac{\Theta_t G(0)}{2} \Delta Q_t,$$

such that the associated cost of a single jump trade is:

$$\left(S_{t-} + \frac{\Theta_t G(0)}{2} \Delta Q_t \right) \Delta Q_t.$$

Notice that the term involving $S_{t-} \Delta Q_t$ is already accounted for in the cost calculation for continuous trading. Therefore, the total cost of all the jump trades is equal to:

$$\frac{G(0)}{2} \sum_{t \in \mathscr{J}} \Theta_t (\Delta Q_t)^2,$$

where $\mathscr{J}$ is the collection of all time points where Q jumps. The total cost of an admissible strategy $Q \in \mathcal{Q}$ is then the sum of the continuous trading cost and the costs of all the jump trades, which is then equal to:

$$C(\Pi) = \int_0^T (S_t - S_0) dQ_t + \frac{G(0)}{2} \sum_{t \in \mathscr{J}} \Theta_t (\Delta Q_t)^2. \tag{3.5}$$

In the following Lemma we calculate the expected cost of an admissible strategy. We use the reasoning of Lemma 2.3 from the paper Gatheral et al. (2012) but adapt it to our model specifications.

Lemma 3.4.1. *The expected cost of an admissible strategy $Q \in \mathcal{Q}$ is equal to:*

$$\begin{aligned}
\mathbb{E}[C(\Pi)] &= \mathbb{E}\left[\int_0^T (S_t - S_0) dQ_t + \frac{G(0)}{2} \sum_{t \in \mathscr{J}} \Theta_t (\Delta Q_t)^2\right] \\
&= \mathbb{E}\left[\frac{1}{2}\int_0^T \int_0^T \widetilde{\Theta}(t,s) G(|t-s|) dQ_s dQ_t\right],
\end{aligned} \tag{3.6}$$

where we define:

$$c(\Pi) := \int_0^T \int_0^T \widetilde{\Theta}(t,s) G(|t-s|) dQ_s dQ_t, \tag{3.7}$$

with:

$$\widetilde{\Theta}(t,s) = \mathbb{1}_{[0,t)}(s)\Theta_s + \mathbb{1}_{[t,T]}(s)\Theta_t.$$

Proof. First note that:

$$\begin{aligned}
C(\Pi) &= \int_0^T (S_t - S_0) dQ_t + \frac{G(0)}{2} \sum_{t \in \mathscr{J}} \Theta_t (\Delta Q_t)^2, \\
&= \int_0^T \int_0^t \Theta_s G(t-s) dQ_s dQ_t + \frac{G(0)}{2} \sum_{t \in \mathscr{J}} \Theta_t (\Delta Q_t)^2 + \int_0^T Z_t dQ_t.
\end{aligned}$$

We observe that:

$$\begin{aligned}
\int_0^T \int_0^t \Theta_s G(t-s) dQ_s dQ_t &= \int_0^T \int_0^t \Theta_s G(|t-s|) dQ_s dQ_t \\
&= \frac{1}{2}\int_0^T \int_0^t \Theta_s G(|t-s|) dQ_s dQ_t + \frac{1}{2}\int_0^T \int_t^T \Theta_t G(|t-s|) dQ_s dQ_t \\
&\quad - \frac{G(0)}{2} \sum_{t \in \mathscr{J}} \Theta_t (\Delta Q_t)^2,
\end{aligned}$$

To ensure causality we have:

$$= \frac{1}{2}\int_0^T \int_0^T (\mathbb{1}_{[0,t)}(s)\Theta_s + \mathbb{1}_{[t,T]}(s)\Theta_t) G(|t-s|) dQ_s dQ_t - \frac{G(0)}{2} \sum_{t \in \mathscr{J}} \Theta_t (\Delta Q_t)^2.$$

Combining gives:

$$C(\Pi) = \frac{1}{2}\int_0^T \int_0^T \widetilde{\Theta}(t,s) G(|t-s|) dQ_s dQ_t + \int_0^T Z_t dQ_t,$$

where:

$$\widetilde{\Theta}(t,s) = \mathbb{1}_{[0,t)}(s)\Theta_s + \mathbb{1}_{[t,T]}(s)\Theta_t.$$

It remains to show that $\mathbb{E}\left[\int_0^T Z_t dQ_t\right] = -Z_0\,\mathbb{E}[Q_0] = 0$. Using Ito's product rule we find (see e.g., Shreve et al. (2004)):

$$\int_0^T Z_t\,dQ_t = Z_T Q_T - Z_0 Q_0 - \int_0^T Q_t\,dZ_t$$

Since Z is a continuous martingale and Q has finite and $\mathbb{P}$-a.s bounded total variation, the stochastic integral $\int_0^T Q_t\,dZ_t$ is a martingale starting from zero and has an expected value of zero. Additionally, since $Q_T = 0$ $\mathbb{P}$-a.s:

$$\begin{aligned}\mathbb{E}\left[\int_0^T Z_t\,dQ_t\right] &= \mathbb{E}\left[Z_T Q_T - Z_0 Q_0 - \int_0^T Q_{t-}\,dZ_t\right]\\ &= \mathbb{E}\left[0 - Z_0 Q_0 - 0\right]\\ &= -Z_0\mathbb{E}[Q_0]\end{aligned}$$

Since $Z_0 = 0$, we find the required result. □

In the next subsection we use this lemma to exclude price manipulation from the model.

3.4.1. Excluding price manipulation

In this subsection, we formally introduce the concept of price manipulation. Furthermore, to prevent price manipulation in the linear propagator model, we derive conditions for the decay kernel $G(\cdot)$ and the liquidity process Θ to satisfy.

To define the notion of price manipulation, we first introduce the concept of a round trip trade. A round trip trade is a strategy Q such that:

$$\int_0^T dQ_t = 0.$$

The integral above can be interpreted as the total net change in the trader's position over the interval from 0 to T.

Price manipulation has multiple definitions with different strengths (see, Definition 1,2 and 3 in Huberman and Stanzl (2004)). To exclude price manipulation in the propagator model we use its weakest form. Gatheral (2010) defines this form as dynamic arbitrage:

Definition 3.4.1. (Dynamic arbitrage) *Dynamic arbitrage is a round trip trade such that the expected cost of trading is negative.*

Using the notion of dynamic arbitrage we define price manipulation as follows:

Definition 3.4.2. (Price manipulation) *A model admits price manipulation if and only if it admits dynamic arbitrage, i.e. there exist a round trip trade such that the expected cost of trading is negative.*

Gatheral (2010) investigates when dynamic arbitrage is possible under the general propagator model and finds that it depends on specific functional forms of the instantaneous impact function $f(\cdot)$ and the decay kernel $G(\cdot)$. For example in Lemma 4.1, he concludes that if the impact function is nonlinear in the trading rate and the decay kernel is an exponential function, price manipulation is always possible.

We deduce from Equation (3.6) that the linear propagator model is free of price manipulation whenever $c(\Pi) \geq 0$ (refer to Equation (3.7)). For a constant liquidity process, this can be characterized by an extension of Bochner's theorem (see Bochner (1932)) and is proofed by Gatheral et al. (2012) in Proposition 2.6. We summarize their result in the following theorem which we present without proof.

Theorem 3.4.1. *Assume the liquidity process Θ to be constant. If the decay kernel $G(\cdot)$ is positive definite then $c(\Pi) \geq 0$, and if $G(\cdot)$ is strictly positive definite then $c(\Pi) > 0$, where $c(\Pi)$ is defined in Equation (3.7). In other words, for every positive definite decay kernel, the linear propagator model is free of price manipulation.*

The theorem above provides a general characterisation on the decay kernel such that the linear propagator model with a constant liquidity process does not admit any price manipulation in the sense of dynamic-arbitrage. Note that a class of strictly positive definite functions are given by the class of bounded non-increasing convex functions.

In general, the liquidity process dependents on the intraday volume curve. Because the intraday volume changes over time (volume tends to be higher near the open and just before the close), price manipulation opportunities could occur. For example, one could move the price up by buying during an ill-liquid period and selling during a liquid period and therefore possibly making a round-trip trade with negative cost. To prevent this from happening we define some restrictions on the liquidity process Θ and in particular, the volume process $V = (V_t)_{t\in[0,T]}$.

To the best of our knowledge, for general decay kernels in continuous time this has not been solved. However, for an exponential linear propagator model with time-dependent liquidity process, it is possible to find a condition. In the following theorem we follow the reasoning of Isichenko (2021) but adapt it to our model specifications.

Theorem 3.4.2. *Consider a continuous admissible trading strategy $Q \in \mathcal{Q}$. Consider the linear propagator model with an exponential decay kernel and assume that the liquidity process Θ is a positive continuous differentiable and deterministic process. Then the model is free of price manipulation when the following inequality is satisfied:*

$$\frac{\dot{\Theta}_t}{\Theta_t} > -2\beta, \qquad \forall t \geq 0, \tag{3.8}$$

where $\dot{\Theta}_t/\Theta_t$ is the percentage change of the process at time t and β is the exponential decay parameter. In particular, when $\Theta_t = (\lambda\sigma)/V_t$, with $\lambda > 0$ and $\sigma > 0$ then:

$$\frac{\dot{V}_t}{V_t} < 2\beta, \qquad \forall t \geq 0,$$

Proof. Consider the dynamics of the market impact I in the linear propagator model with an exponential kernel, which can be written as the SDE:

$$dI_t = -\beta I_t dt + \Theta_t dQ_t$$

We start with expressing the dynamics of the impact in the trading rate $\dot{Q}$:

$$\begin{aligned} dI_t &= -\beta I_t dt + \Theta_t dQ_t \\ \dot{Q}_t &= \frac{\dot{I}_t + \beta I_t}{\Theta_t}. \end{aligned}$$

Then the cost of trading can be rewritten by repeatedly integrating by parts:

$$\begin{aligned} \int_0^T S_t dQ_t &= \int_0^T I_t \dot{Q}_t dt \\ &= \int_0^T I_t \frac{\dot{I}_t + \beta I_t}{\Theta_t} dt \\ &= \int_0^T \frac{I_t^2}{2\Theta_t^2}(\dot{\Theta}_t + 2\beta\Theta_t) dt. \end{aligned}$$

For the cost of trading to be positive, we need above integral to stay positive:

$$\begin{aligned} &\dot{\Theta}_t + 2\beta\Theta_t > 0 \\ &\frac{\dot{\Theta}_t}{\Theta_t} > -2\beta, \end{aligned}$$

for all $t \geq 0$. In particular, when $\Theta_t = (\lambda\sigma)/V_t$, with $\lambda > 0$ and $\sigma > 0$ then:

$$\dot{\Theta}_t = \lambda\sigma \frac{d}{dt}\left(\frac{1}{V_t}\right) = \frac{-\lambda\sigma\dot{V}_t}{V_t^2},$$

and:

$$\begin{aligned} \frac{-\lambda\sigma\dot{V}_t}{V_t^2}\frac{V_t}{\lambda\sigma} &> -2\beta \\ \frac{\dot{V}_t}{V_t} &< 2\beta, \end{aligned}$$

for all $t \geq 0$. □

From Theorem 3.4.1 and Theorem 3.4.2, we conclude that even when the decay kernel is positive definite, for some instances of the liquidity process the model allows for price manipulation. Furthermore, for general decay kernels, no restriction on the liquidity process have been derived yet. To work around this issue, we switch to discrete time version of the model to find a restriction on the liquidity process and decay kernel combined, to prevent price manipulation.

Therefore, consider the discrete time linear propagator model from Equation 3.4. The expected cost of trading under this model is derived in the following lemma for which the proof can be found in Appendix A.

Lemma 3.4.2. *Consider the discrete time version of the linear propagator model with time-dependent liquidity process:*

$$S_n^N = S_0^N + \sum_{m=1}^{n} \Theta_m^N G(n-m)\Delta_m Q^N + Z_n^N.$$

Then for a discrete admissible trading strategy $\Pi^N = Q^N$, the expected cost of trading in discrete time equals:

$$\mathbb{E}[C(\Pi^N)] = \mathbb{E}\left[\frac{1}{2}\sum_{n=1}^{N}\sum_{m=1}^{N} \Theta_{\min(n,m)}^N G(|n-m|)\Delta_m Q^N \Delta_m Q^N\right], \tag{3.9}$$

which in matrix-vector notation is equal to:

$$\mathbb{E}[C(\Pi^N)] = \mathbb{E}\left[\mathbf{q}^T\mathbf{\Phi}\mathbf{q}\right] = \mathbb{E}\left[\mathbf{q}^T\left(\frac{1}{2}\widetilde{\mathbf{\Theta}} \odot \mathbf{G}\right)\mathbf{q}\right], \tag{3.10}$$

where $q_n = \Delta_n Q^N$, $\widetilde{\Theta}_{n,m} = \Theta_{\min(n,m)}$, $G_{n,m} = G(|n-m|)$ and the operator $\odot$ means the Hadamard product between the matrices $\widetilde{\mathbf{\Theta}}$ and $\mathbf{G}$. Notice that $\widetilde{\mathbf{\Theta}}$ and $\mathbf{G}$ are both symmetric matrices. For simplicity we define $\mathbf{\Phi} = 1/2 \cdot \mathbf{\Theta} \odot \mathbf{G}$, which is a symmetric matrix as well.

Remark 3.4.1. *Please note that in the discrete-time case, we assume that we trade uniformly in a time interval such we pay the average instantaneous impact. However, this could easily be adjusted by adding a multiple of the diagonal to the matrix $\mathbf{\Phi}$. This is also a way to regularize the matrix $\mathbf{\Phi}$.*

Using the expected cost of trading derived in lemma above, we are able to find a sufficient condition such that the discrete model is free of price manipulation. This result is formalized in the following theorem.

Theorem 3.4.3. *Consider the discrete time version of the linear propagator model:*

$$S_n^N = S_0^N + \sum_{m=1}^{n} \Theta_m^N G(n-m)\Delta_m Q^N + Z_n^N.$$

Then for a discrete admissible trading strategy $\Pi^N = Q^N$, the model is free of price manipulation when the expected cost of trading is non-negative. That is when the matrix $\mathbf{\Phi}$ is positive definite, i.e. for all $\mathbf{q} \in \mathbb{R}^N$:

$$\mathbf{q}^T\mathbf{\Phi}\mathbf{q} \geq 0.$$

Proof. Follows directly from the definition of positive definiteness. □

Using Theorem 3.4.3 we define the set of admissible liquidity processes such that the model is free of price manipulation.

Definition 3.4.3. (Admissible liquidity processes) *Consider the discrete time version of the linear propagator model with a positive definite decay kernel. Then we define the set of admissible liquidity processes $\mathscr{L}$ by all processes Θ^N such that $\mathbf{\Phi} = 1/2 \cdot \mathbf{\Theta} \odot \mathbf{G}$ is positive definite, i.e. the set:*

$$\mathscr{L} := \{\Theta^N | \mathbf{\Phi} = 1/2 \cdot \mathbf{\Theta} \odot \mathbf{G} \textit{ positive definite } \}$$

Theorem 3.4.3 also provides a usefully check for price manipulation when we have a positive definite decay kernel and an arbitrary liquidity process Θ^N. In the next chapter we use the derived expected cost of trading to find optimal trading strategies.

3.4.2. Transaction-triggered price manipulation

In this subsection we move back to the continuous time model and consider a different type of price manipulation. Alfonsi et al. (2012) investigate the viability of a linear propagator model beyond price manipulation, and have discovered a new class of irregularities referred to as transaction-triggered price manipulation:

Definition 3.4.4. (Transaction-triggered price manipulation) *A market impact model admits transaction-triggered price manipulation when the expected cost of trading of a sell (buy) strategy can be decreased by intermediate buy (sell) trades.*

The mandate of an institutional investors typically prohibits the use of intermediate opposite-side trading. Furthermore, alternating between large buy and sell trades to influence prices in your favour, might even be seen as market manipulation, which is illegal. Therefore, transaction-triggered price manipulation is important to exclude. To investigate when transaction-triggered price manipulation occurs and what conditions are necessary to exclude it, the authors consider a discrete version of the propagator model. Gatheral et al. (2012) and Alfonsi and Schied (2013) generalized the results of Alfonsi et al. (2012) to continuous time. Because this is the most general form of the model we discuss these results in more detail.

We consider the linear propagator model but we take a constant liquidity process Θ. To the best of our knowledge, no results on transaction-triggered price manipulation have been proven in case of a time dependent liquidity process. To dive into transaction-triggered price manipulation we shortly turn to the optimal execution problem.

The optimal execution problem can be formalized in many different ways and we formally introduce it in the next section. At this point we are not seeking to obtain optimal strategies which can be used in industrial execution algorithms. Therefore, Gatheral et al. (2012) note that it is enough to consider a simplified version of the optimal execution problem to characterize irregular solutions. In particular, we consider the problem in which we minimize the expected cost for a risk-neutral trader that buys or liquidates Q_0 number of shares within a given time frame $\mathbb{T} = [0, T]$. This problem is not well-defined when the model admits price manipulation. We therefore assume the decay kernel to be positive definite.

An immediate consequence of Equation (3.7) is that every admissible strategy that minimizes the expected cost must have a path in $\mathcal{D}$. This means that it is sufficient to only consider deterministic admissible strategies to characterize irregular solutions. This significantly simplifies the derivations. Therefore, we assume for the remaining of this chapter that every admissible trading strategy $Q \in \mathcal{D}$ is deterministic.

The first main results of Gatheral et al. (2012) is Theorem 2.11 and it states that the optimal strategies are solutions of a generalized Fredholm integral equation of the first kind. We summarize the result in the following theorem which we present without proof.

Theorem 3.4.4. *Suppose that the decay kernel $G(\cdot)$ is positive definite. Then there exist an optimal solution $Q^* \in \mathcal{D}$ that minimizes $c(\Pi)$ if and only if there exist a constant φ such that Q^* solves the* *generalized Fredholm integral equation:*

$$\int G(|t-s|)dQ_s^* = \varphi \qquad \forall t \in \mathbb{T}.$$

In this case, $C(Q^*) = \varphi Q_0$

The above theorem gives us a characterization of the optimal solution for minimizing the expected cost. Using this characterization Gatheral et al. (2012) proves their second main result. We summarize the result in the following theorem, which we present without prove.

Theorem 3.4.5. *Let the decay kernel $G(\cdot)$ be convex, non-increasing and non-constant. Then there exist a unique optimal strategy $Q^* \in \mathcal{D}$. Moreover, the optimal solution Q^* is free of transaction-triggered price manipulation. This means that the optimal strategy does not alternate between buying and selling.*

Using this theorem we define a set of what we call admissible kernels:

Definition 3.4.5. (Set of admissible kernels) *Consider the decay kernel $G(\cdot)$ of the linear propagator model. Then we define the set of admissible kernels by all the decay kernels for which the linear propagator model with constant liquidity process, is free of price manipulation and transaction-triggered price manipulation. That is, all decay kernels $G(\cdot)$ in the set:*

$$\mathscr{G} := \{G(\cdot)|\textit{Convex, non-increasing and non-constant}\} \tag{3.11}$$

Note that the decay functions:

$$G(t-s) = \frac{\zeta}{(\zeta+t-s)^\gamma} \qquad \text{and} \qquad G(t-s) = \lambda e^{-\beta(t-s)} \qquad \gamma, \beta > 0$$

are in the admissible set of kernels. The structure of the admissible set of decay kernels suggest that there also exist nonparametric kernels. With nonparametric we mean kernels which satisfy the constraints in the set $\mathscr{G}$ without specifying their functional form.

3.5. Calibrating to trading data

In this section we use a two-step approach to calibrate the linear propagator model to intraday trading data and project a general kernel onto the set of admissible kernels. This method is similar to the one discussed by Neuman et al. (2023) but we adapt it to our model specifications.

We consider the situation in which we executes a metaorder into smaller child orders. We let a trader's inventory during the execution of the metaorder be given by Q^N, which describes the number of shares held by the trader at each point in time. Moreover, we assume that there is no short-term alpha signal and we subtract the fair value [2] from the unaffected price process and the observed mid-price such that:

$$S_n^N = I_n^N + P_n^N = S_0^N + I_n^N + Z_n^N,$$

where S^N represents the fair value corrected mid-price cumulative log return in each intraday time bin n. Additionally, we set $\Theta_n^N = \sigma / V_n^N$ for all $n \in [1, T]$, which describes the liquidity in every time bin n, and we assume that the martingale Z^N is equal to $Z_n = \sum_{m=1}^n \sigma\, \epsilon_m$, where ϵ are i.i.d standard normal random variables.

To calibrate the model we use the discrete time linear propagator model defined in Equation 3.4 :

$$S_n^N = S_0^N + \sum_{m=1}^n \Theta_m^N G(n-m)\Delta_m Q^N + Z_n^N,$$

$$S_n^N = S_0^N + \sum_{m=1}^n \frac{\sigma}{V_m^N} G(n-m)\Delta_m Q^N + \sigma \sum_{m=1}^n \epsilon_m^N,$$

$$S_n = S_0 + \sigma \sum_{m=1}^n G(n-m) U_m + \sigma \sum_{m=1}^n \epsilon_m,$$

[2]We use Robeco's internal fair value model

where we omit the superscript N and define $U_n = \Delta_n Q/V_n$ to simplify the notation. Furthermore, we let S_0 be the the fair value corrected mid-price cumulative log-return just before our first trade.

To calibrate the model, we use intraday trading data which is divided into $n = 1, \cdots, T$ time bins. We assume that each metaorder is executed within a single day, and our participation rate in each time bin from the start to the end of the execution is given by the sequence $\{U_m : m = 1, \cdots, M\}$. During the execution of the order, it is possible that we may not participate in certain time bins, resulting in $U_m = 0$ or those time bins m where no trading occurs.

Suppose we have executed $k = 1, \cdots, K$ metaorders on not necessary the same stock. Then we have K different price trajectories. We can then write the cumulative log returns for a given metaorder k in a specific intraday time bin n after the start of the execution as follows:

$$S_n^{(k)} = S_0^{(k)} + \sigma^{(k)} \sum_{m=1}^{n} G_{n-m}^{(k)} U_m^{(k)} + \sigma^{(k)} \sum_{m=1}^{n} \epsilon_m^{(k)}$$

$$\frac{1}{\sigma^{(k)}} \left(S_n^{(k)} - S_0^{(k)} \right) = \sum_{m=1}^{n} G_{n-m}^{(k)} U_m^{(k)} + \sum_{m=1}^{n} \epsilon_m^{(k)}$$

$$\frac{1}{\sigma^{(k)}} \left(S_n^{(k)} - S_0^{(k)} \right) = \sum_{m=1}^{n} U_{(n+1)-m}^{(k)} G_{m-1}^{(k)} + \sum_{m=1}^{n} \epsilon_m^{(k)} \qquad \text{where} \quad G_{n-m}^{(k)} U_m^{(k)} = U_{n-m}^{(k)} G_m^{(k)}$$

The formal proof of the equality $G_{n-m}^{(k)} U_m^{(k)} = U_{n-m}^{(k)} G_m^{(k)}$ can be found in Appendix A. For every $n \in [1, M]$, the equation above can we written in matrix-vector notation in the following way:

$$\mathbf{y}_k = \mathbf{U}_k \mathbf{g}_k + \boldsymbol{\epsilon}_k,$$

where $\mathbf{g}_k = [G_0^{(k)}, \cdots G_{M-1}^{(k)}]^\top \in \mathbb{R}^M$. Furthermore, the other matrix and vectors are equal to:

$$\mathbf{y}_k = \begin{bmatrix} \frac{1}{\sigma^{(k)}} \left(S_1^{(k)} - S_0^{(k)} \right) \\ \vdots \\ \frac{1}{\sigma^{(k)}} \left(S_M^{(k)} - S_0^{(k)} \right) \end{bmatrix} \in \mathbb{R}^M, \quad \mathbf{U}_k = \begin{bmatrix} U_1^{(k)} & 0 & \cdots & 0 \\ U_2^{(k)} & U_1^{(k)} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ U_M^{(k)} & U_{M-1}^{(k)} & \cdots & U_1^{(k)} \end{bmatrix} \in \mathbb{R}^{M \times M}, \quad \boldsymbol{\epsilon}_k = \begin{bmatrix} \epsilon_1^{(k)} \\ \vdots \\ \sum_{m=1}^{M} \epsilon_M^{(k)} \end{bmatrix} \in \mathbb{R}^M,$$

where $\mathbf{U}_k$ is a lower-triangular Toeplitz matrix.

We aim to find a universal kernel for a set of metaorders. This means that the decay kernel vector $\mathbf{g}_k$ should be independent of the metaorder k. However, we cannot simply drop the dependence on k because each metaorder has a different length $M^{(k)}$, and the dimensions of the matrix-vectors should align to find a universal kernel.

To address this, we specify the length of the kernel vector $\mathbf{g}$ beforehand. In our case, we set the length of the kernel vector equal to the number of intraday time bins T, i.e., $\mathbf{g} = [G_0, \cdots, G_{T-1}]^\top \in \mathbb{R}^T$. Consequently, all other matrix-vectors must match the dimension of the decay kernel. We achieve this by appending $T - M$ zeros to the sequence $\{U_m : m = 1, \cdots, M\}$. The matrix-vector notation then changes in the following way:

$$\mathbf{U}_k = \begin{bmatrix} U_1^{(k)} & 0 & \cdots & \cdots & \cdots & \cdots & \cdots & 0 \\ U_2^{(k)} & U_1^{(k)} & 0 & \cdots & \cdots & \cdots & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots & \vdots & \vdots & \vdots \\ U_M^{(k)} & \cdots & U_2^{(k)} & U_1^{(k)} & 0 & \vdots & \vdots & \vdots \\ 0 & U_M^{(k)} & \cdots & U_2^{(k)} & U_1^{(k)} & 0 & \vdots & \vdots \\ \vdots & 0 & \ddots & \vdots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \vdots & \ddots & U_M^{(k)} & \cdots & U_2^{(k)} & U_1^{(k)} & 0 \\ 0 & 0 & \cdots & 0 & U_M^{(k)} & \cdots & U_2^{(k)} & U_1^{(k)} \end{bmatrix} \in \mathbb{R}^{T\times T},$$

$$\mathbf{y}_k = \begin{bmatrix} \frac{1}{\sigma^{(k)}} r_1^{(k)} \\ \vdots \\ \frac{1}{\sigma^{(k)}} r_n^{(k)} \\ \vdots \\ \frac{1}{\sigma^{(k)}} r_T^{(k)} \end{bmatrix} \in \mathbb{R}^T, \qquad \boldsymbol{\epsilon}_k = \begin{bmatrix} \epsilon_1^{(k)} \\ \vdots \\ \sum_{i=1}^n \epsilon_i^{(k)} \\ \vdots \\ \sum_{i=1}^T \epsilon_i^{(k)} \end{bmatrix} \in \mathbb{R}^T,$$

where $\mathbf{U}_k$ remains a lower-triangular Toeplitz matrix. Additionally, we let $r_1^{(k)}, \cdots, r_T^{(k)}$ represent the fair value and volatility corrected cumulative mid-price log returns in each time bin n after the start of execution. We use normalized log returns to ensure comparability of stocks both across the cross-section and over different days. Note that if $M < T$, we have that $r_M^{(k)}, \cdots, r_T^{(k)}$ are all equal to zero.

Using this matrix-vector representation, we use a two-step approach to find the optimal admissible decay kernel $\mathbf{g}^* \in \mathscr{G}$:

1. First, solve the following least-squares problem to find the optimal non-projected kernel:

$$\widetilde{\mathbf{g}} = \arg\min_{\mathbf{g}} \sum_{k=1}^{K} \left\|\mathbf{y}_k - \mathbf{U}_k \mathbf{g}\right\|^2. \tag{3.12}$$

   This optimization problem is a convex quadratic optimization problem with an analytical solution given in Theorem 3.5.1.

2. Second, project the analytical solution onto the set of admissible kernels by solving:

$$\mathbf{g}^* = \arg\min_{\mathbf{g}\in\mathscr{G}} ||\mathbf{g} - \widetilde{\mathbf{g}}||^2. \tag{3.13}$$

   We solve this optimization problem using *scipy.optimize.curve_fit* in Python for a parametric kernel. For a nonparametric kernel, we use the "SLSQP" solver within the *scipy.optimize.minimize* package to enforce necessary constraints.

   The optimal admissible decay kernel of length T is then given by $\mathbf{g}^* = [G_0^*, \cdots, G_{T-1}^*]^\mathsf{T}$. Since this kernel function acts as a convolution kernel, we can measure the decay up to T time bins after the execution of the last child order.

This approach is advantageous because it is efficient: the first and computationally the most intensive step has an analytical solution, which significantly speeds up the process. Additionally, the analytical solution provides valuable insights into the underlying structure of the data.

The analytical solution of the least-squares problem of Equation (3.12) is given in the theorem below.

Theorem 3.5.1. (Analytical solution least-squares) *For an arbitrary decay kernel $\mathbf{g}$, the least-squares problem:*

$$\min_{\mathbf{g}} \sum_{k=1}^{K} ||\mathbf{y}_k - \mathbf{U}_k \mathbf{g}||^2,$$

has the following analytical solution:

$$\widetilde{\mathbf{g}} = (\mathbf{W})^{-1}\sum_{k=1}^{K}\mathbf{U}_k^T\mathbf{y}_k,$$

where:

$$\mathbf{W} = \sum_{k=1}^{K}\mathbf{U}_k^T\mathbf{U}_k.$$

Moreover, the standard errors of the solution are given by:

$$se(\widetilde{\mathbf{g}})_n = \sqrt{\hat{\sigma}^2\left(\mathbf{W}^{-1}\right)_{n,n}}, \tag{3.14}$$

for all $n = 1, \cdots T$, with

$$\hat{\sigma}^2 = \frac{1}{K-p-1}\sum_{k=1}^{K}||\mathbf{y}_k - \mathbf{U}_k\widetilde{\mathbf{g}}||^2,$$

where p is the number of predictors which is in our case equal to T.

Proof. The analytical solution of the least-squares problem can be found in the following way:

$$\min_{\mathbf{g}}\sum_{k=1}^{K}||\mathbf{y}_k - \mathbf{U}_k\mathbf{g}||^2 = \min_{\mathbf{g}}\sum_{k=1}^{K}(\mathbf{y}_k - \mathbf{U}_k\mathbf{g})^{\mathsf{T}}(\mathbf{y}_k - \mathbf{U}_k\mathbf{g}),$$

Taking derivative w.r.t $\mathbf{g}$ and setting to zero then gives:

$$\begin{aligned}
\frac{\partial}{\partial\mathbf{g}}\left(\sum_{k=1}^{K}(\mathbf{y}_k - \mathbf{U}_k\mathbf{g})^{\mathsf{T}}(\mathbf{y}_k - \mathbf{U}_k\mathbf{g})\right) &= 0 \\
\sum_{k=1}^{K}\mathbf{U}_k^{\mathsf{T}}\mathbf{y}_k - \sum_{k=1}^{K}\mathbf{U}_k^{\mathsf{T}}\mathbf{U}_k\mathbf{g} &= 0 \\
\sum_{k=1}^{K}\mathbf{U}_k^{\mathsf{T}}\mathbf{U}_k\mathbf{g} &= \sum_{k=1}^{K}\mathbf{U}_k^{\mathsf{T}}\mathbf{y}_k.
\end{aligned}$$

Lastly, solving for $\mathbf{g}$ then gives the optimal solution:

$$\begin{aligned}
\widetilde{\mathbf{g}}_N &= \left(\sum_{k=1}^{K}\mathbf{U}_k^{\mathsf{T}}\mathbf{U}_k\right)^{-1}\sum_{k=1}^{K}\mathbf{U}_k^{\mathsf{T}}\mathbf{y}_k, \\
&= (\mathbf{W})^{-1}\sum_{k=1}^{K}\mathbf{U}_k^{\mathsf{T}}\mathbf{y}_k.
\end{aligned}$$

Because the objective function is convex and quadratic, the solution is a global minimizer.

To find the associated standard errors, we start with calculating the covariance matrix of the solution. In case of our this is equal to:

$$\begin{aligned}
\mathrm{Cov}(\widetilde{\mathbf{g}}) = \mathbb{E}\left[(\widetilde{\mathbf{g}} - \mathbb{E}[\widetilde{\mathbf{g}}])(\widetilde{\mathbf{g}} - \mathbb{E}[\widetilde{\mathbf{g}}])^{\mathsf{T}}\right] &= \mathbb{E}\left[\left(\mathbf{W}^{-1}\sum_{k=1}^{K}\mathbf{U}_k^{\mathsf{T}}\boldsymbol{\epsilon}_k\right)\left(\mathbf{W}^{-1}\sum_{k=1}^{K}\mathbf{U}_k^{\mathsf{T}}\boldsymbol{\epsilon}_k\right)^{\mathsf{T}}\right] \\
&= \sigma^2\mathbf{W}^{-1}.
\end{aligned}$$

Then the standard errors of the least-squares problem are given by:

$$\text{se}(\widetilde{\mathbf{g}})_n = \sqrt{\hat{\sigma}^2 \text{Cov}(\widetilde{\mathbf{g}})_{n,n}} = \sqrt{\hat{\sigma}^2 \, (\mathbf{W}^{-1})_{n,n}},$$

for all $n = 1, \cdots, T$. Furthermore, the variance $\hat{\sigma}^2$ has to be estimated using the sample variance of the prediction error in the following way:

$$\hat{\sigma}^2 = \frac{1}{K - p - 1} \sum_{k=1}^{K} ||\mathbf{y}_k - \mathbf{U}_k \widetilde{\mathbf{g}}||^2,$$

where p is the number of predictors which is in our case equal to T. □

3.6. Numerical results calibration

In this section, we calibrate the linear propagator model to intraday trading data using the two-step approach outlined in the previous section. First, we find the non-projected decay kernel and then project it onto the admissible set of kernels. Next, we evaluate the kernels based on predefined metrics. Finally, we compare the performance of the linear propagator model with industry standards.

In this section, we evaluate the parameter uncertainty using parametric standard errors as defined in Equation (3.14) and nonparametric bootstrapped standard errors (see, Efron and Tibshirani (1994)). Both confidence intervals should be of similar magnitude.

Furthermore, to evaluate the model on performance and accuracy, we use respectively the Root Mean Squared Error (RMSE):

$$\text{RMSE} = \sqrt{\frac{1}{T \cdot K + 1} \sum_{i=1}^{T \cdot K} (y_i - \hat{y}_i)^2},$$

and the centered R-squared:

$$R_c^2 = 1 - \frac{\sum_{i=1}^{T \cdot K} (y_i - \hat{y}_i)^2}{\sum_{i=1}^{T \cdot K} (y_i - 0)^2}.$$

We assume that the price dynamics, in the absence of trading, follow an arithmetic Brownian motion. Therefore, we use the centered R-squared with the mean set to zero, which we refer to simply as the R-squared.

In the calculation of the RMSE and the R-squared we combine all observations and predictions for every metaorder k such that we get one number for each measure instead of one for every metaorder k separately.

3.6.1. The analytical kernel estimate

Solving the least-squares problem defined in Equation (3.12) results in the non-projected kernel, as shown on the left side of Figure 3.3. In this figure, we plot the analytical solution along with the corresponding confidence interval. On the right side of Figure 3.3, we display the bootstrapped version of the analytical solution using 20 bootstrap samples.

We see in Figure 3.3 that the instantaneous impact $G(0)$ is around 2.49 bps as fraction of the the intraday 5 min volatility. To put this into perspective, suppose the 5 min volatility is 15 bps and we participate 10% in a 5 min interval, we make $2.49 \cdot 0.10 \cdot 15 = 3.73$ bps instantaneous market impact. Moreover, we find that the half-life is only 2 time bins. The half-life represents the duration after execution over which the impact has decayed to half of its maximum impact. So in our case the non-projected kernel estimate suggests that after 10 minutes the impact has already decayed to half of its maximum.

In addition, we see in the figure that after $50 \cdot 5 = 250$ min, the impact has reverted to almost zero. However, we should note that after time bin 50, the measurements become unreliable as can be seen in the bootstrapped estimate . The reason for this is that there are not a lot of orders in the data set for

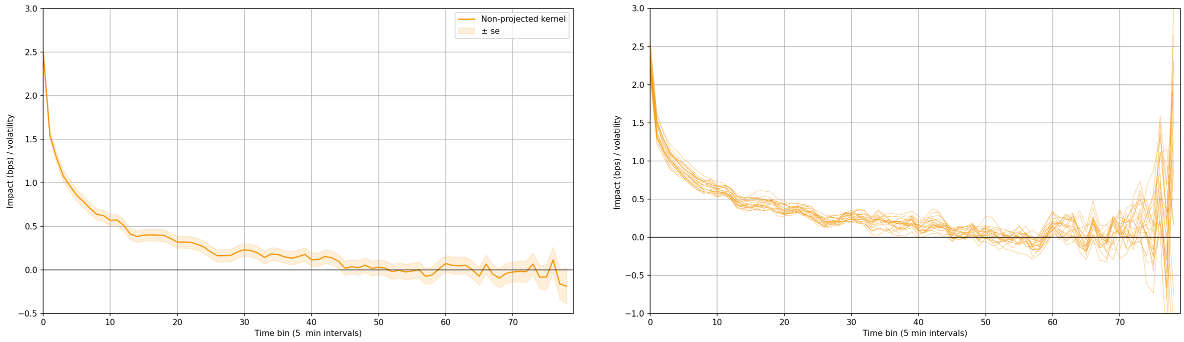

Figure 3.3: (Left) the analytical solution of the least-squares problem in Equation (3.12) based on the calibration sample defined in Table 2.1. (Right) the bootstrapped non-projected kernel estimates using 20 bootstraps.

which we can measure the decay for more than 50 intervals after the execution of a child order (see, Figure C.1).

When comparing the decay of the non-projected kernel with the decay of the total impact of a metaorder, as visualized in Figure 2.7, we observe some similarities. Both exhibit a two-stage decay: an initial short, fast decay followed by a long, slow decay. Additionally, the half-life of the decay is similar in both cases. However, a notable difference is that the non-projected kernel shows no evidence of permanent impact, as it fully reverts back to zero, whereas the decay in Figure 2.7 suggests a non-zero limit. Notice, that the analytical solution is not globally convex, monotonic decreasing and non-negative. Therefore, we still have to project this function on the set of admissible kernels.

To estimate the non-projected kernel we use a least-squares approximation. To make sure that the least-square approximation gives the best unbiased solution to the problem, some assumptions need to be satisfied (see Stock and Watson (2020)). Which regression assumptions apply in this special case and whether they are satisfied is further investigated in Appendix B.1.

The predictor of the least-square regression is the participation rate of every child order. Therefore, it is likely that the model performance depends on the participation rate of a metaorder as well. Here we define the participation rate during the order as the number of shares traded during the whole order divided by the total of shares traded by the market during the execution of the metaorder. To investigate how the model performances changes with respect to this predictor, we make a scatterplot of the RMSE vs the participation rate during the order. We use the RMSE because it is more robust for outliers than the R-squared for individuals metaorders. Since we have many observations we use a linear spline regression (see Hastie (1986)) with the associated confidence intervals to find a pattern. The scatterplot with the linear spline regression for the participation rate per metaorder is present in Figure 3.4.

In Figure 3.4, we see that the RMSE decreases when the participation rate increases. The strongest decrease is after a participation rate of 0.001 We can explain this observation as follows. If the participation rate per order increases, then it is more likely that we notably influence the price by taking away liquidity from the market. This results in a higher signal to noise ratio and therefore also in a better model fit.

An important explicit assumption in the linear propagator model is that the system is time invariant. This means that the shape of the decay kernel does not change over time. Properly testing this assumption is challenging due to the identification problem. This problem refers to the difficulty in distinguishing between the instantaneous impact and the decay of impact in a sequence of child orders. We leave this issue for future research.

3.6.2. Projecting on the admissible set

We proceed by projecting the analytical solution onto the set of admissible kernels by solving the optimization problem in Equation (3.13). We consider two functions for the decay kernels: an exponential function and a power-law function. However, we use a slight variation of the power-law function be-

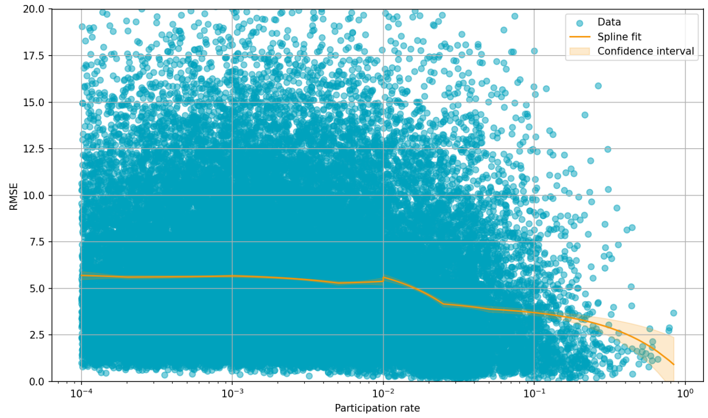

Figure 3.4: Scatterplot of participation rate and RMSE with linear spline regression in the knots: $\{0.0002, 0.001, 0.005, 0.01, 0.025, 0.05, 0.01\}$.

cause it is a better fit for our data. The decay kernels we use are defined as follows:

$$G_1(t;\zeta,\gamma) = \frac{\zeta}{(1+t)^{\gamma}}, \quad \zeta,\gamma > 0 \quad \text{and} \quad G_2(t;\lambda,\beta) = \lambda e^{-\beta t}, \quad \lambda,\beta > 0.$$

By construction, these functions are within the set of admissible kernels, so we do not need to impose any constraints on the optimization problem. We solve this optimization problem for the parameters ζ,γ and λ,β using a numerical solver in Python. We use *scipy.optimize.curve_fit*, which is ideal for fitting nonlinear functions to data. We do not apply any bounds or constraints and initialize the solver with $[1, 0.1]$. To calculate the standard errors for the parameters, we use the covariance matrix of the estimates, which is directly returned by the solver.

To find the nonparametric decay kernel, we need to impose constraints on the optimization problem to ensure it belongs to the set of admissible kernels. For this, we use the ”SLSQP” solver within the *scipy.optimize.minimize* function in Python, as it is well-suited for handling constrained optimization problems. We initialize the solver with the analytical solution of the least-squares problem. To determine the confidence intervals for the nonparametric kernel, we bootstrap the estimates.

Solving the projections for the parametric and nonparametric kernels result in Figure 3.5. In these figures we plot the analytical solution of the least-squares problem with their admissible projections. The bootstrapped nonparametric kernel can be found is Figure C.2 in Appendix B. A detailed overview in terms of parameters and accuracy of these projections are present in Table 3.1. In this table we give the parameter sets, half-lifes and standard errors.

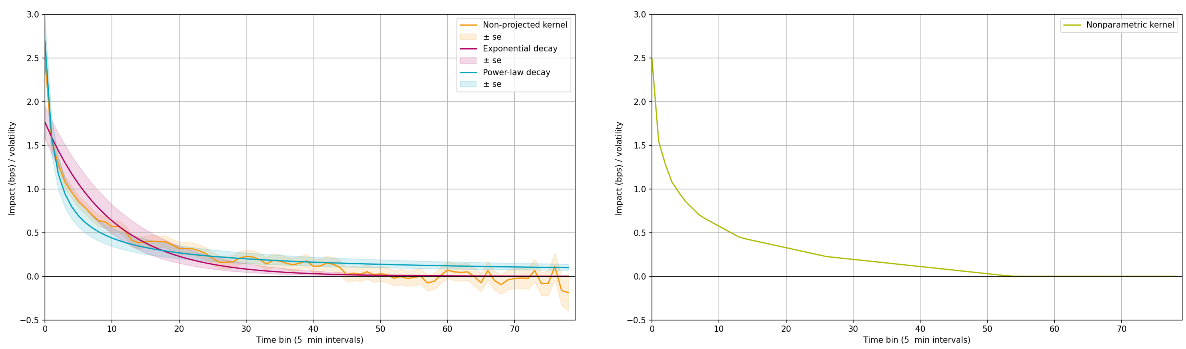

Figure 3.5: (Left) the analytical solution of the least-squares problem plotted with the estimated exponential and power-law decay function with the corresponding confidence intervals. (Right) the estimated nonparametric decay kernel.

From Table 3.1 we find that the exponential kernel has the lowest standard errors for the parameter estimates. Even though, the instantaneous and the half-life deviates most from the analytical solution.

Table 3.1: Parameters, standard errors (se) and half-lifes for the analytical solution, the exponential, power-law and nonparametric decay kernel estimates

Function	Estimates	se	half-life (min)
Non-projected	$\widetilde{G}_N(0) = 2.4908$	-	10
Exponential	$\lambda = 1.7635$	0.0738	34
Exponential	$\beta = 0.1022$	0.0063	34
Power-law	$\zeta = 2.6913$	0.1015	8
Power-law	$\gamma = 0.7578$	0.0242	8
nonparametric	$G_N^*(0) = 2.4908$	-	10

Furthermore, from Figure 3.5 we find that the power-law kernel seems the best fit in the first few buckets. The instantaneous impact and the half-life are closed to the non-projected kernel estimate but after the fifth bucket it starts to deviate. Also, we should note that the power-law kernel does not fully revert to zero within the given time frame, where the others does. The nonparametric kernel seems the best overall fit for the analytical solution. It perfectly matches the analytical solution and belongs to the set of admissible kernels. However, this is also what is expected because it has the most degrees of freedom to fit the analytical solution.

3.6.3. Performance evaluation

To measure the predictive power of the model with the estimated decay kernels we look at the in/out-of-sample RMSE and in/out-of-sample R-squared. Therefore, we perform a 5 fold moving window cross validation and calculate the average RMSE and the average R-squared with corresponding standard errors for the different time windows. The results are present Table 3.2.

Table 3.2: In-sample and out-of-sample RMSE and R-squared calculated based on the moving window cross validation framework. All metrics are calculated for the model with exponential, power-law and nonparametric decay kernels.

Function	In-sample RMSE	s.e	Out-sample RMSE	s.e
Exponential	6.5176	0.1956	6.5185	0.1957
Power-law	6.5175	0.1956	6.5185	0.1957
nonparametric	6.5174	0.1964	6.5184	0.1957
Function	In-sample $R^2(\%)$	s.e (10^{-3})	Out-sample $R^2(\%)$	s.e (10^{-3})
Exponential	0.217	0.1987	0.212	0.2008
Power-law	0.219	0.2134	0.210	0.2186
nonparametric	0.222	0.2084	0.214	0.2141

From Table 3.2 we find that the overall in-sample performance is best for the model with nonparametric decay kernel. Again this is as expected since it has more degrees of freedom compared to parametric kernels. In addition, the model with the nonparametric decay kernel is also the best out-sample fit. This is maybe a bit surprising since we would expect that it might over-fit on the training sets. In general all evaluation metrics are really close which makes it impossible to conclude which kernel is the best performing one.

The R-squared for the model on the whole data set is 0.24%. An R-squared of 0.24% means that we can only explain 0.24% of the variance around its mean, which is in this case zero. During most orders our signal to noise ratio is very small since our average execution time 360 minutes and the average participation rate of our child orders is 0.047%. To get an idea how much R-squared we can expected with our data, we run a Monte Carlo simulation. We assume that the market behaves like the linear propagator model and we simulate cumulative log returns from this model for every order. We do this 256 times per order and compare the simulated paths with the predicted paths. The result is an empirical distribution for the R-squared and is given in Figure 3.6.

In Figure 3.6 we see that the calculated R-squared based on the actual cumulative log returns (point estimate) falls within the empirical distribution of the simulated data and is almost equal to the mean. This means that the R-squared is as expected based on the low signal to noise ratio of the data. We have done this Monte Carlo simulation because their are no performance resultd available in the lit-

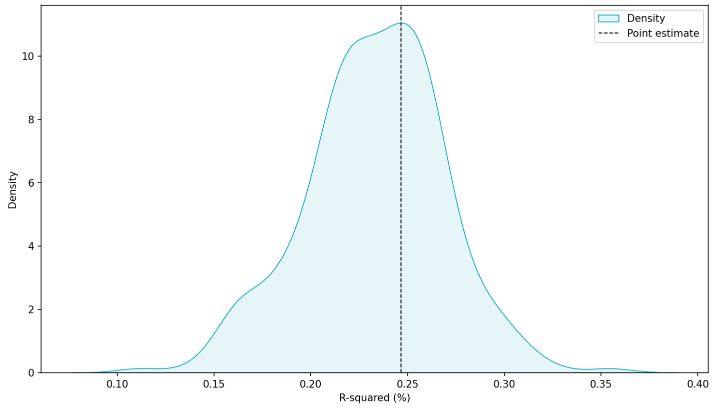

Figure 3.6: Empirical distribution of the R-squared based of a Monte Carlo simulation of the linear propagator model with simulated unaffected price paths. Per metaorder we use 256 realizations. The point estimate is the calculated R-squared based on the actual cumulative log returns.

erature. The only result which gives a bit of inside of the model performance of a propagator model calibrated proprietary order data is by Hey, Bouchaud, et al. (2023). In the paper the authors calibrated the OW-model on CFM's metaorder data and found an R-squared of approximately 0.5% (See Figure 3(c) of Hey, Bouchaud, et al. (2023)). However, we should note that CFM's dataset has other characteristics.

The new variant of the propagator model is a linear model based on the participation rate. In Appendix B.2, we demonstrate that this model achieves a higher R-squared compared to a linear propagator model using the number of shares normalized by ADV. Additionally, we show that it performs comparably to a locally concave model that also uses the number of shares normalized by ADV.

We conclude this section by showing that the expected market impact and expected cost are in line with results from the literature. We start with calculating the expected market impact and the decay of a metaorder. This gives an idea how the model predicts the market impact trajectory and how we can use this model further in transaction cost analysis. To calculate the expected cost of a trading strategy we use the expected implementation shortfall (IS) to make the cost comparable across models. The IS for the linear propagator model is defined by (see Lemma 3.4.1):

$$\mathbb{E}[IS(\text{Linear propagator model})] = \int_0^T \int_0^t \frac{\sigma}{V_t} G(t-s) dQ_s dQ_t,$$

and we benchmark it against the power-law model for market impact Bouchaud et al. (2018). This model is often used by practitioners to get a pre-trade estimate of the implementation shortfall and is defined as:

$$\mathbb{E}[IS(\text{power-law model})] = \alpha \cdot \sigma \cdot \left(\frac{Q_0}{ADV}\right)^{\delta},$$

where σ is the annualized volatility, Q_0 is the order size and ADV is the average daily volume. Furthermore, α and δ are parameters than need to be calibrated to data [3].

For illustration, we calculate the expected impact and expected IS of a 5% average daily volume (ADV) order using a VWAP strategy. We assume that the annualized volatility is equal to 21%, and we use the power-law decay kernel. The result is displayed in Figure 3.7. In Figure 3.7 we see that the maximum impact of a 5% ADV order is approximately 16 bps and that the expected IS is approximately 13 bps. The expected IS using the power-law model is equal to 15 bps. This is a difference of 2 bps with the expected IS using the propagator model.

Furthermore, we see in the figure that the impact has not fully reverted before the end of the next trading day. This is also what is expected when looking at the shape of the power-law kernel. However, we

[3] We use internally calibrated parameters

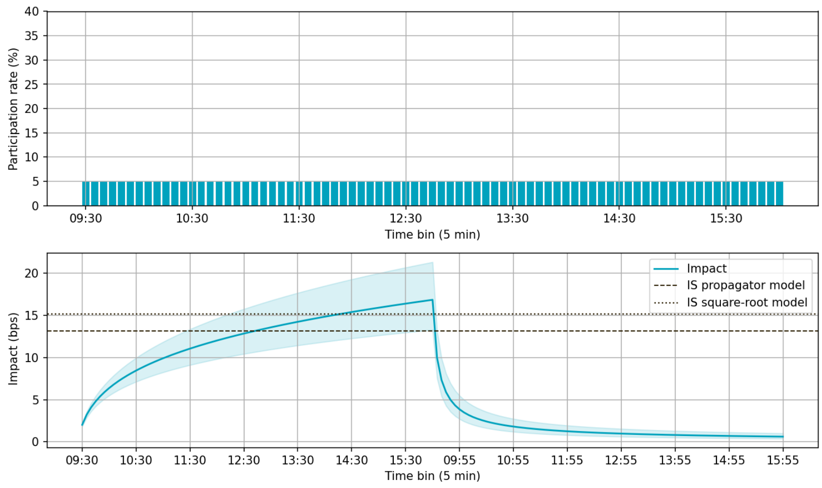

Figure 3.7: (Upper figure) evolution of the participation rate during a VWAP trading strategy which takes a whole day to execute. (Lower Figure) expected impact trajectory during and after execution of the metaorder with associated confidence interval. The black dotted lines represents the expected IS of the execution.

should note that this is just an extrapolation because the kernel is only estimated on intraday data. Furthermore, how market impact behaves during the closing and overnight session is still an open question and we leave it to further research.

To show how the linear propagator model behaves in terms of expected IS on a metaorder level, we calculate the expected IS for different order sizes. To calculate the IS for the linear propagator model we use VWAP strategies which take the whole day to execute. Moreover, we assume a 21% annualized volatility, which correspond with an average US stock. To benchmark expected IS we also plot the square-root model with internally calibrated parameters and fit a spline regression to the data. This is displayed in Figure 3.8

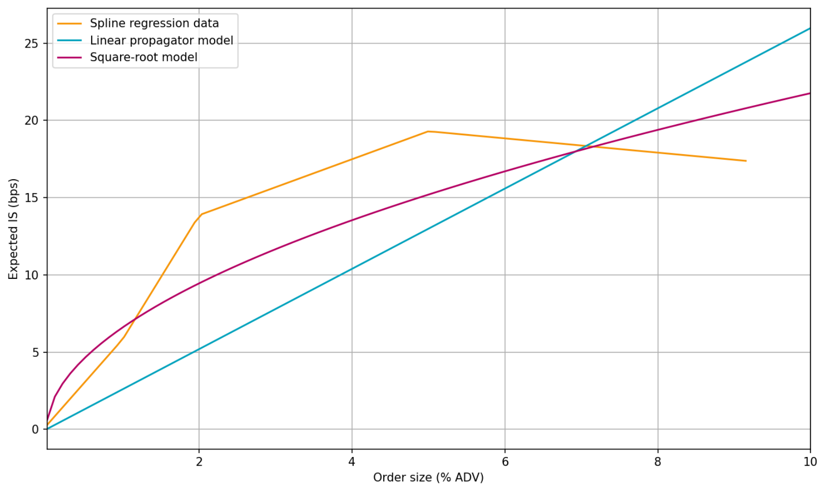

Figure 3.8: Expected implementation shortfall of a metaorder executed over the whole day using a VWAP strategy for multiple order sizes as percentage of ADV and 21% annualized volatility. As benchmark we use the calibrated power-law model and a spline regression to represent the data.

We see in the Figure 3.8 that also on the metaorder level the propagator model has a linear relation between cost and order size. Comparing it with the square-root model, we find that the propagator model underestimates the cost for smaller order sizes and over estimates for order sizes above 7% ADV because the data shows a concave relation as already mentioned. To tackle this issue Alfonsi et al. (2010) introduced the globally concave propagator model, also known as the AFS model.

4

Optimal intraday execution strategies

This chapter begins with an introduction to optimal control problems and applying this framework to the OW model. Subsequently, we derive the optimal execution problem for the new introduced linear propagator model with general decay kernels and time-varying liquidity process. Furthermore, we extend the optimal execution problem to a mean-variance optimization problem and incorporate short-term alpha signals. We solve all discrete-time versions of the optimal execution problems analytically and illustrate the results by means of realistic examples.

4.1. Optimal control problems

Optimal control problems are present in various domains of quantitative finance, with portfolio optimization problems, such as the one solved by Merton (1975), serving as classical examples. The general goal of an optimal control problem is to maximize (minimize) an expected profit (cost) function by identifying a strategy that influences the dynamics of an underlying stochastic process.

For example in Merton's problem, an agent seeks to maximize his expected (discounted) wealth by trading in a risky asset and a risk-free bank account. The agent actions affect his wealth, but the asset dynamics also impact this wealth. Thus, the optimal allocation depends on both the agent's wealth process and the asset dynamics combined, creating a complex optimization problem.

Mathematically, optimal control problems are formulated as follows. Let $(Q_t)_{t\geq 0}$ denote the control process and $(X_t)_{t\geq 0}$ the state variable influenced by Q. For instance, X could be modeled by an Ito diffusion process:

$$dX_t = \mu(t, X_t, Q_t)dt + \sigma(t, X_t, Q_t)dW_t,$$

where W represents Brownian motion. The optimal control problem, in its general form, is given by:

$$\sup_{Q\in\mathcal{A}} \mathbb{E}\left[H(X_T) + \int_0^T F(t, X_t, Q_t)dt\right],$$

where $\mathcal{A}$ is a set of admissible controls, H is a terminal reward or penalty depending on the final state of X, and F is a running reward or penalty possibly dependent on time, the state variable, and the control variable. The objective function may also be subject to constraints.

The interpretation of this general formulation is that an agent aims to maximize the terminal reward H and the running reward F by determining an optimal control process Q. The agent's actions influence the dynamics of the state variable X, meaning past actions affect future states, necessitating an adaptive strategy to account for this feedback.

Various methods exist to solve optimal control problems. The simplest approach involves transforming the problem into a pointwise (myopic) optimization problem, which is solvable via numerical or analytical techniques. However, only a subset of optimal control problems can be addressed this way.

More sophisticated methods include the dynamic programming principle (DPP) and the associated Hamilton-Jacobi-Bellman (HJB) equation, also known as the dynamic programming equation (DPE) (see, e.g., Bertsekas (2012), Cartea et al. (2015)). This approach converts the optimal control problem into a series of time-indexed optimization problems, from which a DPP can be derived. The infinitesimal form of DPP leads to a DPE, whose solution provides a tentative solution to the original control problem.

Other techniques involve the Pontryagin maximum principle (see, e.g., Yong and Zhou (2012)), backward stochastic differential equations, and reinforcement learning. In this chapter, we focus on transforming optimal control problems into pointwise optimization problems.

4.2. Optimal execution for the OW model

To introduce the optimal execution problem, we revisit the OW model from Section 3.2.2. Following the reasoning of Webster (2023), we formulate the optimal execution problem from the perspective of a trader aiming to maximize their profit and loss position.

We first formulate the problem in discrete time before extending it to continuous time. We use the notation defined in Definition 3.1.4 and the processes described in Section 3.1 and 3.2.2 but for the remainder of this chapter we only consider deterministic admissible trading strategies $Q \in \mathcal{D}$. Additionally, we introduce two new processes. The trader's cash position is defined as the stochastic process C^N, which satisfies:

$$\Delta_n C^N = -\tilde{S}^N_{n-1}\Delta_n Q^N,$$

where $\tilde{S}^N_{n-1}$ is the execution price of a trade, $\Delta_n Q^N$ is the number of shares traded and C^N_0 is the initial cash position. Furthermore, we define a trader's fundamental profit and loss (P&L) position by the process X^N, which satisfies:

$$X^N_n = P^N_n Q^N_n + C^N_n,$$

where P^N is the unobserved price process. Using these processes we are able to define the discrete self-financing equation. The discrete fundamental P&L process X^N satisfies the following self-financing equations:

$$\begin{aligned}
\Delta_n X^N &= X^N_n - X^N_{n-1} = P^N_n Q^N_n + C^N_n - P^N_{n-1}Q^N_{n-1} + C^N_{n-1} \\
&= Q^N_{n-1}\Delta_n P^N + P^N_{n-1}\Delta_n Q^N + \Delta_n S^N \Delta_n Q^N - \Delta_n C^N \\
&= Q^N_{n-1}\Delta_n P^N + (P^N_{n-1} - \tilde{S}^N_{n-1})\Delta_n Q^N + \Delta_n S^N \Delta_n Q^N \\
&= Q^N_{n-1}\Delta_n P^N - I^N_{n-1}\Delta_n Q^N - \bar{S}^N_{n-1}\Delta_n Q^N + \Delta_n S^N \Delta_n Q^N.
\end{aligned}$$

Note that these self-financing equations are different from the Black-Scholes self-financing equations (see, Hull (2016)), in the sense that we have lifted the frictionless trading assumption. This means the observable price and the execution price are not the same.

Every term in the final version of the self-financing equation has a specific interpretation. Lets have a closer look:

(a) $Q^N_{n-1}\Delta_n P^N$ represents the change in value of the inventory.

(b) $-I^N_{n-1}\Delta_n Q^N$ is the price impact of the trade $\Delta_n Q^N$

(c) $-\bar{S}^N_{n-1}\Delta_n Q^N$ is the transaction cost the trader has to pay or receives.

(d) $\Delta_n S^N \Delta_n Q^N$ is the adverse selection of the trading book.

To find the fundamental P&L process under the OW model we include the OW market microstructure assumptions: the linear instantaneous transaction cost and no adverse selection assumption. This means that $\bar{S}^N_{n-1}\Delta_n Q^N = \frac{\lambda}{2}(\Delta_n Q^N)^2$ and that $\Delta_n S^N \Delta_n Q^N \to 0$. The fundamental P&L process under the OW model then becomes:

$$\Delta_n X^N = Q^N_{n-1}\Delta_n P^N - I^N_{n-1}\Delta_n Q^N - \frac{\lambda}{2}(\Delta_n Q^N)^2.$$

Webster (2023) shows that as $N \to \infty$, the discrete P&L converges to the continuous time limit:

$$dX_t = Q_t dP_t - I_t dQ_t - \frac{\lambda}{2} d[Q,Q]_t,$$

where $d[Q,Q]_t$ is the quadratic variation of Q. Note that this term is non-zero since the trading strategy can exhibit jumps.

Recall that the impact process under the OW model with a deterministic strategy $Q \in \mathcal{D}$ is given by the Ordinary differential equation (ODE):

$$dI_t = -\beta I_t dt + \lambda dQ_t.$$

Then substituting this into the quadratic variation term, we find that the simplified fundamental P&L process:

$$dX_t = Q_t dP_t - I_t dQ_t - \frac{1}{2} d[Q,I]_t,$$

where $d[Q,I]_t$ denotes the quadratic covariation between the processes Q and I. Using the fundamental P&L process defined above, we are able to define the optimal execution problem for the OW model.

Suppose we would like to trade a position $Q_0 \neq 0$ such that $Q_T = 0$, then the goal is to maximise the fundamental P&L process X by choosing a deterministic admissible strategy $\Pi = (Q_t)_{t\in[0,T]} \in \mathcal{D}$. Then we define the optimal execution problem under the OW model as:

$$\sup_{Q\in\mathcal{D}} \mathbb{E}\left[\int_0^T Q_t dP_t - \int_0^T I_t dQ_t - \frac{1}{2}[I,Q]_T\right] \quad (4.1)$$
$$\text{s.t } Q_T = 0$$

This control problem can be solved explicitly by using an important insight from Fruth et al. (2014). They exploit the one-to-one map between the control variable Q and the state variable I in the OW model. The authors prove in Lemma 8.6 that the optimal control problem can be solved by inverting this relationship. We summarize the result in the following theorem which we present with proof.

Theorem 4.2.1. (Solution to the OW optimal execution problem) *Let the unaffected price process P be equal to the continuous martingale $Z \in \mathcal{Z}$. Furthermore, consider a deterministic admissible strategy $Q \in \mathcal{D}$ and the OW optimal execution problem:*

$$\sup_{Q\in\mathcal{D}} \mathbb{E}\left[\int_0^T Q_t dP_t - \int_0^T I_t dQ_t - \frac{1}{2}[I,Q]_T\right] \quad (4.2)$$
$$s.t\ Q_T = 0$$

Then the optimal trading strategy is given by:

$$dQ_t^* = \begin{cases} -\frac{\beta Q_0}{2+\beta T} dt, & \forall t \in (0,T) \\ -\frac{Q_0}{2+\beta T}, & t = 0, t = T \end{cases}$$

Proof. We start by rewriting the dynamics of the OW model:

$$dQ_t = \frac{1}{\lambda}(\beta I_t dt + dI_t).$$

Then by substituting this into the control problem we find for the expression in the expectation:

$$\begin{aligned}\int_0^T Q_t dP_t - \int_0^T I_t dQ_t - \frac{1}{2}[I,Q]_T &= \int_0^T Q_t dP_t - \frac{1}{\lambda}\int_0^T I_t(\beta I_t dt + dI_t) - \frac{1}{2\lambda}[I,I]_T \\ &= \int_0^T Q_t dP_t - \frac{1}{\lambda}\int_0^T \beta I_t^2 dt - \frac{1}{\lambda}\int_0^T I_t dI_t - \frac{1}{2\lambda}[I,I]_T.\end{aligned}$$

Then using Ito's lemma (see e.g., Shreve et al. (2004)) with $f(t, I_t) = \frac{1}{2}I_t^2$ and without loss of generality we let $I_0 = 0$, we find that the integral $\int_0^T I_t dI_t$ equals:

$$\begin{aligned}\int_0^T I_t dI_t &= \frac{1}{2}I_T^2 - \frac{1}{2}I_0^2 - \frac{1}{2}\int_0^T d[I]_t \\ &= \frac{1}{2}I_T^2 - \frac{1}{2}[I,I]_T\end{aligned}$$

Substituting the above into the original expression and simplifying, we obtain:

$$\int_0^T Q_t dP_t - \frac{1}{\lambda}\int_0^T \beta I_t^2 dt - \frac{1}{2\lambda}I_T^2,$$

and the transformed optimal control problem becomes:

$$\begin{aligned}&\sup_I \mathbb{E}\left[\int_0^T Q_t dP_t - \frac{1}{\lambda}\int_0^T \beta I_t^2 dt - \frac{1}{2\lambda}I_T^2\right] \\ &\text{s.t} \quad I_T + \int_0^T \beta I_t dt = -\lambda Q_0,\end{aligned}$$

By assumption we have that the unaffected price process P is a continuous martingale Z. Therefore, the optimal control problem simplifies in:

$$\begin{aligned}&\inf_I \frac{1}{\lambda}\mathbb{E}\left[\int_0^T \beta I_t^2 dt + \frac{1}{2}I_T^2\right] \\ &\text{s.t} \quad I_T + \int_0^T \beta I_t dt = -\lambda Q_0,\end{aligned}$$

because integrating w.r.t a martingale has expectation zero. The above control problem is much easier to solve than the original problem because we can solve it using the Lagrange multiplier method. Introducing a Lagrange multiplier $\nu \in \mathbb{R}$ for the constraint, the Lagrangian for the problem is formulated as

$$\mathcal{L}(I_t, \nu) = \int_0^T \beta I_t^2\, dt + \frac{1}{2}I_T^2 - \nu\left(I_T + \int_0^T \beta I_t\, dt + \lambda Q_0\right).$$

We then find the optimal trajectory by taking partial derivatives of $\mathcal{L}$ with respect to I_t for all $t \in (0,T)$ and I_T and setting them to zero. This yields the conditions:

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial I_t} &= 2\beta I_t - \nu\beta = 0, \\ \frac{\partial \mathcal{L}}{\partial I_T} &= I_T - \nu = 0.\end{aligned}$$

From the above conditions, we deduce that the optimal trajectory is given by:

$$\begin{aligned}I_t^* &= \frac{\nu}{2}, \quad \text{for } t \in (0,T), \\ I_T^* &= \nu.\end{aligned}$$

The linear constraint is then used to solve the Lagrange multiplier ν. The constraint can be rewritten into:

$$\begin{aligned}\nu + \int_0^T \beta\frac{\nu}{2}dt &= -\lambda Q_0 \\ \nu + \beta\frac{\nu}{2}T + \lambda Q_0 &= 0.\end{aligned}$$

Solving for ν results in:

$$\nu = -\frac{2\lambda Q_0}{1+\beta T}.$$

Finally, using the one-to-one mapping between I and Q given by the dynamics $dQ_t = \frac{1}{\lambda}(\beta I_t dt + dI_t)$, we obtain the optimal strategy for Q_t for all $t \in (0, T)$:

$$\begin{aligned} dQ_t^* &= \frac{1}{\lambda}(\beta I_t^* dt + dI_t^*) \\ &= -\frac{\beta Q_0}{2+\beta T} dt. \end{aligned}$$

□

The optimal solution of the OW optimal execution problem calculated in Theorem 4.2.1 is visualized in Figure 4.1. In this figure we visualize the block trades at the beginning and and the end of the order with two dots and the continuous trading with a solid line. In this example we take the whole day to execute 10000 shares.

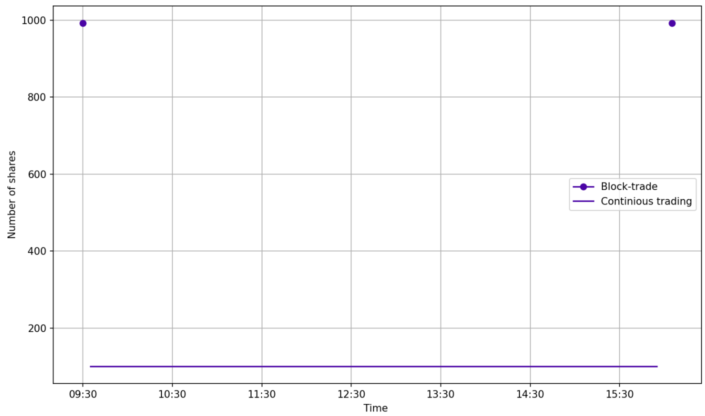

Figure 4.1: Optimal trading strategy under the continuous-time OW model. $Q_0 = 10000$, $dt = 5$ min and we trade from the start till the end of the trading day.

4.3. Optimal execution for the linear propagator model

In this section, we formally introduce the optimal execution problem for the new variant of linear propagator model incorporating a time-dependent liquidity process. We begin with the discrete-time formulation before we extend it to continuous time. We illustrate the theoretical findings by means of a realistic example.

To derive a general optimal control problem, we consider the discrete self financing-equations again from previous section:

$$\Delta_n X^N = Q_{n-1}^N \Delta_n P^N - I_{n-1}^N \Delta_n Q^N - \bar{S}_{n-1}^N \Delta_n Q^N + \Delta_n S^N \Delta_n Q^N.$$

To derive the control problem for the linear model, we need to make some market microstructure assumptions. The first assumption we make is that there is no adverse selection in the order book, i.e. $\Delta_n S^N \Delta_n Q^N \to 0$. Furthermore, since we consider the linear propagator model with time dependent liquidity process, we assume that the average execution price is:

$$\bar{S}_{n-1}^N = S_{n-1}^N + \frac{G(0)}{2} \Theta_n^N \Delta_n Q^N,$$

such that the instantaneous transaction cost equal to:

$$\bar{S}_{n-1}^N \Delta_n Q^N = \frac{G(0)}{2} \Theta_n^N (\Delta_n Q^N)^2.$$

Note that this complies with the calculation of the cost of trading in Section 3.4. The fundamental P&L process X^N under the linear propagator model becomes:

$$\Delta_n X^N = Q_{n-1}^N \Delta_n P^N - I_{n-1}^N \Delta_n Q^N - \frac{G(0)}{2} \Theta_n^N (\Delta_n Q^N)^2.$$

By taking $N \to \infty$, we find the continuous time limit:

$$dX_t = Q_t dP_t - I_t dQ_t - \frac{G(0)}{2} d[\Theta Q, Q]_t.$$

The goal is to maximize the traders P&L process. Using the definition of quadratic variation, and inserting the linear propagator model with time-varying liquidity, we rewrite this into:

$$\mathbb{E}\left[\int_0^T Q_t dP_t - \int_0^T \int_0^t \Theta_s G(t-s) dQ_s dQ_t - \frac{G(0)}{2} \sum_{t \in \mathscr{J}} \Theta_t (\Delta Q_t)^2\right],$$

where $\mathscr{J}$ is the set of all time points where Q jumps. In the second and third term we recognize the cost due to respectively continuous and discrete trading. Using Lemma 3.4.1, the general control problem becomes:

$$\begin{aligned} &\sup_{Q \in \mathcal{D}} \mathbb{E}\left[\int_0^T Q_t dP_t - \frac{1}{2} \int_0^T \int_0^T \widetilde{\Theta}(t,s) G(|t-s|) dQ_s dQ_t\right], \\ &\text{s.t.} \quad \int_0^T dQ_t = Q_0, \end{aligned} \tag{4.3}$$

where we take the supremum over the set of deterministic admissible strategies. In the coming subsections we extend this objective function such that we can incorporate a risk measure and include a alpha signal.

If we assume that the unaffected price process is equal to a martingale $Z \in \mathcal{Z}$, maximising the traders P&L process is equivalent to minimizing the expected cost of trading:

$$\begin{aligned} &\sup_{Q \in \mathcal{D}} \mathbb{E}\left[\int_0^T Q_t dZ_t - \frac{1}{2} \int_0^T \int_0^T \widetilde{\Theta}(t,s) G(|t-s|) dQ_s dQ_t\right], \\ &= \inf_{Q \in \mathcal{D}} \mathbb{E}\left[\frac{1}{2} \int_0^T \int_0^T \widetilde{\Theta}(t,s) G(|t-s|) dQ_s dQ_t\right], \end{aligned}$$

which is equal to the expected cost of trading by Lemma 3.4.1. Therefore, the optimal execution problem to minimize the expected cost of trading when trading a position $Q_0 \neq 0$ such that $Q_T = 0$ becomes:

$$\begin{aligned} &\inf_{Q \in \mathcal{D}} \frac{1}{2} \int_0^T \int_0^T \widetilde{\Theta}(t,s) G(|t-s|) dQ_s dQ_t \\ &\text{s.t.} \quad \int_0^T dQ_t = Q_0. \end{aligned} \tag{4.4}$$

Solving this problem in continuous time is mathematically complex. Therefore, we revert to the discrete time linear propagator model from Equation 3.4. This approach enables us to find an analytical optimal execution strategy using the Lagrange multiplier method, which we present in the following theorem.

Theorem 4.3.1. *Consider the discrete time linear propagator model from Equation 3.4 for which we derived the expected cost of trading in Lemma 3.4.2. Then the corresponding discrete time optimal execution problem to minimize the expected cost of trading a position $Q_0^N \neq 0$ such that $Q_T^N = 0$ is equal to:*

$$\begin{aligned} &\min_{\mathbf{q} \in \mathcal{D}} \mathbf{q}^T \mathbf{\Phi} \mathbf{q} \\ &\textit{s.t.} \quad \mathbf{1}^T \mathbf{q} = Q_0^N. \end{aligned}$$

This is a quadratic optimization problem with optimal solution equal to:

$$\mathbf{q}^* = \frac{Q_0^N \boldsymbol{\Phi}^{-1}\mathbf{1}}{\mathbf{1}^T\boldsymbol{\Phi}^{-1}\mathbf{1}}.$$

Proof. Let $\mathcal{L}(\mathbf{q}, \nu)$ be the Lagrangian for this problem, given by:

$$\mathcal{L}(\mathbf{q}, \nu) = \mathbf{q}^T\boldsymbol{\Phi}\mathbf{q} + \nu(Q_0^N - \mathbf{1}^T\mathbf{q}).$$

Taking the gradient of $\mathcal{L}$ with respect to $\mathbf{q}$ and setting it to zero yields:

$$2\boldsymbol{\Phi}\mathbf{q} - \nu\mathbf{1} = 0.$$

Solving for $\mathbf{q}$ gives us:

$$\mathbf{q} = \frac{\nu}{2}\boldsymbol{\Phi}^{-1}\mathbf{1}.$$

Applying the constraint $\mathbf{1}^T\mathbf{q} = Q_0^N$, we have:

$$\mathbf{1}^T\left(\frac{\nu}{2}\boldsymbol{\Phi}^{-1}\mathbf{1}\right) = Q_0^N,$$

which simplifies to:

$$\nu = \frac{2Q_0^N}{\mathbf{1}^T\boldsymbol{\Phi}^{-1}\mathbf{1}}.$$

Substituting this value of ν back into the expression for $\mathbf{q}$, we obtain the optimal solution:

$$\mathbf{q}^* = \frac{Q_0^N \boldsymbol{\Phi}^{-1}\mathbf{1}}{\mathbf{1}^T\boldsymbol{\Phi}^{-1}\mathbf{1}}.$$

This solution is a global minimizer since the objective function is quadratic and convex and $\boldsymbol{\Phi}$ is symmetric positive definite (SPD). □

To demonstrate that the OW optimal execution problem coincides with the discrete time optimal execution problem for the linear propagator model with an exponential kernel and a constant liquidity process, we compare the solutions. We use a small grid size, $dt = T/N = 79/1000$, and trade $Q_0^N = 10000$ shares over the course of the entire day. The solution for the OW model is calculated using Theorem 4.2.1, while the solution for the discrete linear propagator model is calculated using Theorem 4.3.1. The resulting optimal strategies are present in Figure 4.2, in which we see that both solutions to the optimal execution problems align perfectly.

To visualize the solutions to the optimal execution problem with a time-varying liquidity process in this and the forthcoming sections, we assume that for all $n \in [1, T]$:

$$\Theta_n^N = \frac{\sigma}{V_n^N} = \frac{\sigma(5 \text{ min})}{\text{ADV} \cdot v_n^N},$$

where $\sigma(5 \text{ min})$ is the average 5 min intraday volatility in bps, ADV the average daily volume of a stock and we let v^N be the intraday volume curve. As predictor for v^N we take the average volume curve of all stocks in the MSCI US index in 2023, which can be found in Figure 2.5 of Section 2.3.1. However, in practice more sophisticated methods can be considered as well.

Furthermore, we let the average daily volume (ADV) be equal to 1000000 shares and assume that the intraday $\sigma(5 \text{ min}) = 15$ bps. This corresponds with an annualized volatility of 21%. Moreover, we trade a 5% ADV order which takes the whole day to execute and use the calibrated power-law kernel for illustration.

To make sure that the resulting model does not admit any price manipulation with this specific forecast for the liquidity process Θ^N, we use the result from Theorem 3.4.3. Using the calibrated power-law

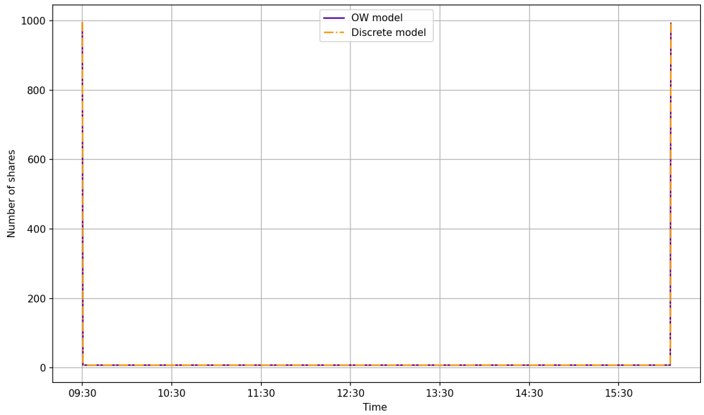

Figure 4.2: Optimal trading strategy under OW model and discrete linear propagator model with exponential kernel and constant liquidity process. $Q_0^N = 10000$ and $dt = 79/1000$.

kernel, we find that the matrix $\mathbf{\Phi}$ is positive definite which means that $\Theta^N \in \mathscr{L}$ and the model is free of price manipulation. If this were not the case, we could add a regularization term to the diagonal of the matrix $\mathbf{\Phi}$ to ensure it becomes positive definite. This regularization term represents additional costs on the instantaneous impact, such as spread costs (see Remark 3.4.1).

However, it is still possible that optimal solution admits transaction-triggered price manipulation, i.e. alternating between buy and sell trades to decrease the expected cost of trading. When this occurs, we use a solver to find a optimal schedule which is restricted to one way trading. The solver we use is *cvxpy* in Python. In addition, we restrict the solver to buy/sell more than 20% of the volume in a 5 min interval. We do this such that the restricted solution found by the optimizer is realistic and implementable by a trader. In summary, when we use a solver we add the following constraint to the optimization problem:

$$q_n \geq 0 \quad \text{and} \quad q_n \leq 0.2V_n^N, \quad \forall n \in [1, T].$$

Calculating the optimal execution strategy with the before mentioned parameters using Theorem 4.3.1 and the solver results in Figure 4.3. As benchmark we use a VWAP strategy and we plot the traders participation rate in every 5 minute interval.

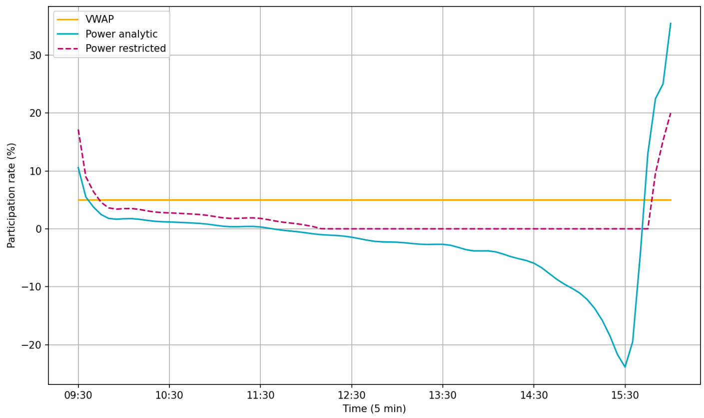

Figure 4.3: Optimal risk-neutral trading strategies with VWAP as benchmark. We trade a 5% ADV order with a power-law kernel and plot the traders participation rate in every 5 min interval.

We see in Figure 4.3 that the optimal analytical solution deviate quite a bit from the restricted solution. The analytical solution would like to take a large short position to buy it back near the end of the trading day. Asset managers are typically not allowed to follow such a strategy when execution an order. Therefore, we also calculate the optimal solution via a solver which is restricted to one way trading. However, this results in a sub-optimal solution in terms of cost compared to the analytical solution because we restrict the solution space. To find out how material this difference is we calculate the evolution of the impact and expected cost of trading for these strategies. These are displayed in in Figure 4.4.

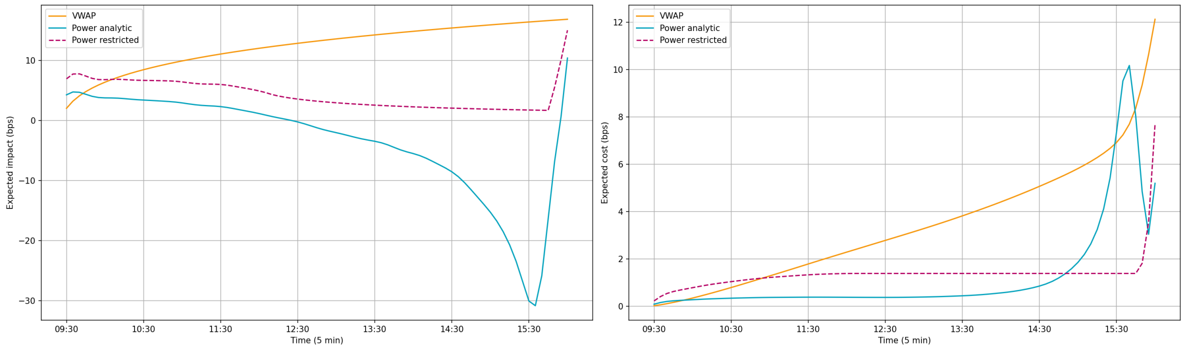

Figure 4.4: (Left) the evolution of the expected impact in bps for the optimal, restricted and VWAP strategies. (Right) the evolution of the expected cost in bps for the optimal, restricted and VWAP strategies.

In the left hand side of Figure 4.4 the evolution of the market impact in bps is plotted for the different strategies. We see that the VWAP and restricted solution both reach a similar final impact state while the analytical solution also has a negative impact because it take a large intermediate short position. In the right hand side plot of Figure 4.4 we see that the analytical solution has the lowest expected cost and VWAP the largest. Moreover, we conclude from this figure that the analytical solution admits transaction-triggered price manipulation because this strategy decreases its expected cost by trading in the opposite way.

To get a better understanding how material the difference is in terms of cost, we calculate the relative difference in the following table.

Table 4.1: Expected cost of trading a 5% ADV order and relative out-performance in % compared with a VWAP strategy for the power-law kernel.

Strategy	Expected cost (bps)	Out-performance (%)
VWAP	12.13	0
Optimal restricted	7.69	−36.639
Optimal analytical	5.19	−57.175

We see in Table 4.1 that the difference in cost is high. The restricted solution is already 37 % cheaper in terms of expected cost than the VWAP strategy and the analytical solution even more. If the restricted solution is really implementable remains a question because targeting 20% of the volume in the first and last time bins of the day might not be realistic.

Furthermore, we should note that comparing strategies on the expected cost alone is not fair. As with everything in financial markets, extra reward comes with extra risk, which in our case means, lower cost comes with a higher risk. Therefore, in the next section we dive into the cost/risk trade off of trading strategies.

4.4. Mean-variance optimal

In this section we extend the optimal execution problem by adding a risk measure to the objective function. In this setting the optimal execution problem can be viewed as a mean-variance optimization.

In general, a trader would like to minimize the expected cost but also tries to manage his risk exposure. For example, when a trader takes a long time to execute an order, the risk that the price moves against him increases. Therefore, a risk measure is important to consider in realistic trading settings.

There are multiple ways to manage the risk exposure in calculating optimal strategies (see e.g. Cartea et al. (2015) and Lehalle and Neuman (2019)). One way is to add the following penalty term to the objective function:

$$\sup_{Q\in\mathcal{D}} \mathbb{E}\left[\int_0^T Q_t dP_t - \frac{1}{2}\int_0^T\int_0^T \widetilde{\Theta}(t,s)G(|t-s|)dQ_s dQ_t - \psi\left(\int_0^T P_t dQ_t - \int_0^T \mathbb{E}[P_t]dQ_t\right)^2\right].$$

This penalty term is the variance of the cost trading and is derived in the lemma below.

Lemma 4.4.1. *Consider the total cost of a deterministic admissible strategy $Q \in \mathcal{D}$ under the linear propagator model (refer to Equation 3.5). Then the variance of the total cost is equal to:*

$$Var(C(\Pi)) = \mathbb{E}\left[\left(\int_0^T P_t dQ_t - \int_0^T \mathbb{E}[P_t]dQ_t\right)^2\right].$$

In particular, if the unaffected price P is equal to the martingale $Z_t = \int_0^t \sigma_s dW_s$ where W is a Brownian motion and σ_t is a time dependent volatility. Then the variance of the total cost of trading is equal to:

$$Var(C(\Pi)) = \int_0^T \left(\int_t^T \sigma_s dQ_s\right)^2 dt. \tag{4.5}$$

For a discrete deterministic admissible strategy Π^N under the discrete linear propagator model, the variance of the total cost in matrix-vector notation equals:

$$Var(C(\Pi^N)) = \mathbf{q}^T\mathbf{\Sigma}\mathbf{q},$$

where:

$$\mathbf{\Sigma} = \mathbf{LSL}^T,$$

and $\mathbf{L}$ a matrix filled with ones on and below the diagonal and zeros above the diagonal, and $\mathbf{S}$ is a matrix with the variances σ_n^2 of each of the periods $n \in [1, T]$ on the diagonal. If the volatility is constant, i.e. $\sigma_n = \sigma$ for all $n \in [1, T]$, we have:

$$\mathbf{\Sigma} = \sigma^2 \cdot \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & 2 & 2 & \cdots & 2 \\ 1 & 2 & 3 & \cdots & 3 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 2 & 3 & \cdots & (T-1) \end{bmatrix}$$

Proof. Consider the total cost of an deterministic admissible strategy $Q \in \mathcal{D}$:

$$C(\Pi) = \int_0^T (S_t - S_0)dQ_t + \frac{G(0)}{2}\sum_{t\in\mathscr{J}} \Theta_t(\Delta Q_t)^2.$$

Using results from the proof of Lemma 3.4.1, we know that:

$$\begin{aligned} C(\Pi) &= \int_0^T (S_t - S_0)dQ_t + \frac{G(0)}{2}\sum_{t\in\mathscr{J}} \Theta_t(\Delta Q_t)^2, \\ &= \int_0^T P_t dQ_t + \frac{1}{2}\int_0^T\int_0^T \widetilde{\Theta}(t,s)G(|t-s|)dQ_s dQ_t. \end{aligned}$$

Then from the definition of variance we get:

$$
\begin{aligned}
Var(C(\Pi)) &= \mathbb{E}[(C(\Pi) - \mathbb{E}[C(\Pi)])^2] \\
&= \mathbb{E}\Bigg[\bigg(\int_0^T P_t dQ_t + \frac{1}{2}\int_0^T \int_0^T \widetilde{\Theta}(t,s) G(|t-s|) dQ_s dQ_t \\
&- \frac{1}{2}\int_0^T \int_0^T \widetilde{\Theta}(t,s) G(|t-s|) dQ_s dQ_t - \mathbb{E}\bigg[\int_0^T P_t dQ_t\bigg]\bigg)^2\Bigg], \\
&= \mathbb{E}\left[\left(\int_0^T P_t dQ_t - \int_0^T \mathbb{E}[P_t] dQ_t\right)^2\right].
\end{aligned}
$$

Let $P_t = \int_0^t \sigma_s dW_s$, then:

$$
\begin{aligned}
\mathbb{E}\left[\left(\int_0^T P_t dQ_t\right)^2\right] &= \mathbb{E}\left[\left(\int_0^T \int_0^t \sigma_s dW_s dQ_t\right)^2\right], \\
&= \mathbb{E}\left[\left(\int_0^T \left(\int_t^T \sigma_s dQ_s\right) dW_t\right)^2\right], \\
&= \int_0^T \left(\int_t^T \sigma_s dQ_s\right)^2 d[W]_t, \\
&= \int_0^T \left(\int_t^T \sigma_s dQ_s\right)^2 dt,
\end{aligned}
$$

where in the second to last line we used Ito's isometry and in the last line we used that the quadratic variation of a Brownian motion is equal to t.

For the discrete time case consider a deterministic admissible strategy Π^N under the discrete linear propagator model. Then using the notation in Definition 3.1.4, we discretize the above continuous version to obtain the variance of the total cost of trading:

$$
Var(C(\Pi^N)) = \sum_{n=1}^{T} \left(\sum_{m=n}^{T} \sigma_m \Delta_m Q^N\right)^2.
$$

Since the variance is bilinear in Q^N (see, Busseti and Lillo (2012)), we can express it in matrix-vector notation as follows:

$$
Var(C(\Pi^N)) = \mathbf{q}^T \mathbf{\Sigma} \mathbf{q},
$$

where we have that $q_n = \Delta_n Q^N$ for all $n = 1, \cdots, T$ and:

$$
\mathbf{\Sigma} = \mathbf{L}\mathbf{S}\mathbf{L}^T,
$$

and $\mathbf{L}$ a matrix filled with ones on and below the diagonal and zeros above the diagonal, and $\mathbf{S}$ is a matrix with the variances σ_n^2 of each of the periods $n \in [1, T]$ on the diagonal. If the volatility is constant, i.e. $\sigma_n = \sigma$ for all $n \in [1, T]$, we have

$$
\mathbf{\Sigma} = \sigma^2 \cdot \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & 2 & 2 & \cdots & 2 \\ 1 & 2 & 3 & \cdots & 3 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 2 & 3 & \cdots & (T-1) \end{bmatrix}
$$

This concludes the proof. □

Using the lemma above with a constant volatility, the optimal execution problem with the variance of the total cost of trading as risk measure becomes:

$$\begin{aligned}&\inf_{Q\in\mathcal{D}} \frac{1}{2}\int_0^T\int_0^T \widetilde{\Theta}(t,s)G(|t-s|)dQ_s dQ_t + \psi\int_0^T\left(\int_t^T \sigma dQ_s\right)^2 dt, \\ &\text{s.t.} \quad \int_0^T dQ_t = Q_0.\end{aligned} \tag{4.6}$$

where ϕ is a risk aversion parameter. Note, that this formulation of the optimal execution problem is equal to a quadratic mean-variance optimization. Again as in the previous section we switch to discrete time to find an analytical solution. Solving the problem in continuous time is mathematically complex.

The analytical solution of the discrete time variant is derived in the following proposition and the proof is a direct consequence of the proof of Theorem 4.3.1.

Proposition 4.4.1. *Consider the discrete time linear propagator model from Equation 3.4 for which we derived the expected cost of trading in Lemma 3.4.2 and the variance of the cost in Lemma 4.4. Then the discrete version of the optimal execution problem in Equation (4.6) can be formalized in matrix-vector notation as follows:*

$$\begin{aligned}&\min_{\mathbf{q}\in\mathcal{D}} \mathbf{q}^T(\mathbf{\Phi}+\psi\mathbf{\Sigma})\mathbf{q} \\ &s.t. \quad \mathbf{1}^T\mathbf{q} = Q_0^N,\end{aligned}$$

This is a quadratic optimization problem with optimal solution equal to:

$$\mathbf{q}^* = \frac{Q_0^N(\mathbf{\Phi}+\psi\mathbf{\Sigma})^{-1}\mathbf{1}}{\mathbf{1}^T(\mathbf{\Phi}+\psi\mathbf{\Sigma})^{-1}\mathbf{1}}.$$

Proof. Direct consequence of Theorem 4.3.1, since $\mathbf{\Phi}+\psi\mathbf{\Sigma}$ is symmetric. □

The optimal strategies for a risk aware trader are visualized in Figure 4.5. In this figure we plot the traders participation rate in every 5 minute interval for the analytical solution and the solution calculated with the solver. Because we are risk-averse with respect to the variance of our expected trading cost, which is driven by the volatility of the asset, we are front loading our strategies compared to risk neutral strategies. This means that we want to achieve a higher order completion in the beginning of the order compared to the risk neutral one.

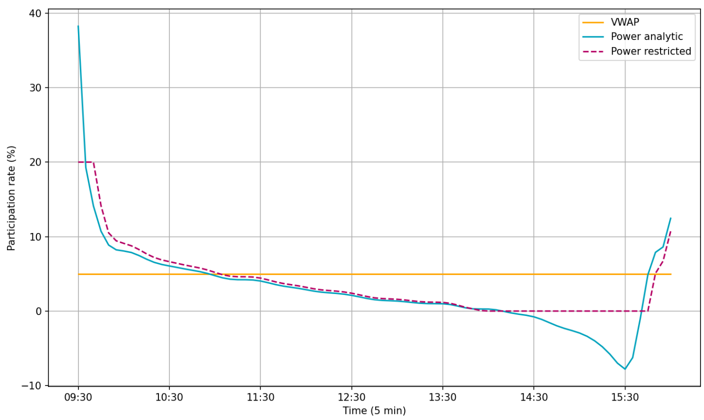

Figure 4.5: Optimal risk-aware trading strategies with VWAP as benchmark. We trade a 5% ADV order with a power-law kernel and plot the traders participation rate in every 5 min interval.

We see in Figure 4.5 that the analytical solution and the restricted solution are more similar than in the risk neutral case. However, the analytical solution still wants to sell some shares in the end of the day to buy it back just before the closing auction. Furthermore, we see that both strategies are front loading their strategies. This comes with a cost because we deviate from the optimal risk-neutral trading strategy. In Figure 4.6 we visualize the evolution of the impact and the cost of the strategy.

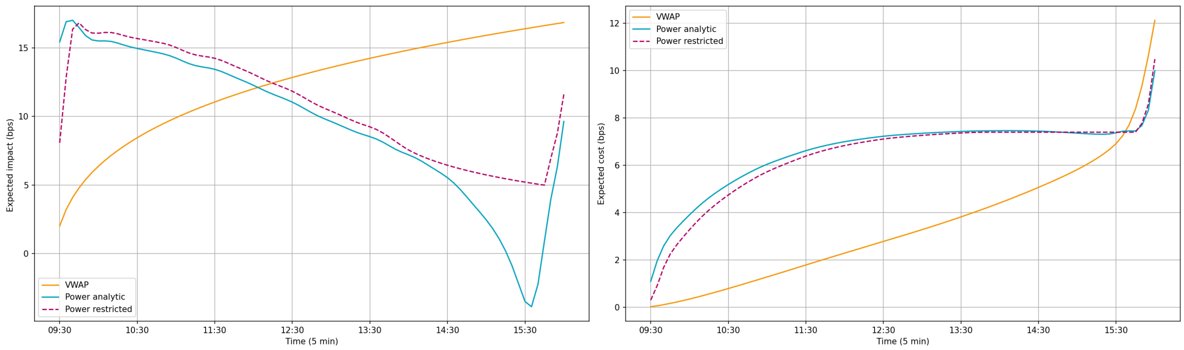

Figure 4.6: (Left) The evolution of the expected impact in bps for the optimal, restricted and VWAP strategies. (Right) The evolution of the expected cost in bps for the optimal, restricted and VWAP strategies.

If we compare Figures 4.4 and 4.6, we find that the risk aware strategy makes more impact in the beginning and less in the end and has a higher total cost than the risk neutral strategy. Since we are able to calculate the variance of the cost of trading, we can construct an efficient frontier between the standard deviation of the cost and the expected cost of trading. In Figure 4.7 we plot the efficient frontier for the analytical and restricted solutions.

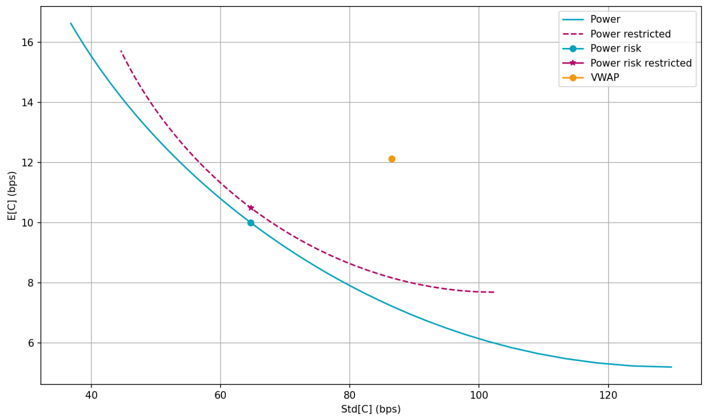

Figure 4.7: Efficient frontier of the analytical and restricted solutions for a 5% ADV order. As benchmark, we plot the VWAP strategy and the risk aware strategies we calculated in Figure 4.5.

In Figure 4.7, we see that there is a trade-off between the expected cost of trading an the standard deviation of the cost. Additionally, we see that the restricted solutions are indeed sub-optimal to the analytical solutions, for the same standard deviation the analytical solution has a lower cost. However, the analytical solutions are not implementable. Moreover, we find that for the same standard deviation as a VWAP strategy we are able to choose an optimal strategy with a lower cost. The efficient frontier enables a trader to choose a point on the frontier and retrieves the optimal strategy a desired combination of risk and cost.

4.5. Optimal strategies with alpha signals

In this section we extend the objective function further by including a short-term alpha signal. An alpha signal is an estimation of how the unaffected price changes in the near further. Mathematically, we define an alpha signal as a stochastic process that models:

$$\alpha_t = \mathbb{E}[P_T - P_t | \mathcal{F}_t].$$

Using this definition of an alpha signal, the dynamics of the unaffected price process $P \in \mathcal{S}$ is a semimartingale and equal to:

$$dP_t = -d\alpha_t + dZ_t,$$

where we assume that $Z \in \mathcal{Z}$ is a continuous martingale and equal to $Z_t = \mathbb{E}[P_T|\mathcal{F}_t]$ for all $t \in [0, T]$. Furthermore, without loss of generality we assume $Z_0 = 0$.

To incorporate the alpha signal, we consider the general objective function:

$$\sup_{Q\in\mathcal{D}} \mathbb{E}\left[\int_0^T Q_t dP_t - \frac{1}{2}\int_0^T\int_0^T \widetilde{\Theta}(t,s)G(|t-s|)dQ_s dQ_t - \psi\left(\int_0^T P_t dQ_t - \int_0^T \mathbb{E}[P_t]dQ_t\right)^2\right],$$

and we use the following identity:

Lemma 4.5.1. *Let $P \in \mathcal{S}$ be the unaffected price process which has the following dynamics:*

$$dP_t = -d\alpha_t + dZ_t,$$

where $\alpha_t = \mathbb{E}[P_T - P_t|\mathcal{F}_t]$, $Z \in \mathcal{Z}$ is a martingale and $Q \in \mathcal{D}$ a deterministic admissible trading strategy. Then

$$\mathbb{E}\left[\int_0^T Q_t dP_t\right] = \mathbb{E}\left[\int_0^T \alpha_t dQ_t\right]$$

Proof. By Ito's product rule we have:

$$\begin{aligned}
\alpha_T Q_T &= \int_0^T \alpha_t dQ_t + \int_0^T Q_t d\alpha_t + [\alpha, Q]_T, \\
0 &= \int_0^T \alpha_t dQ_t - \int_0^T Q_t dP_t + \int_0^T Q_t dZ_t + [\alpha, Q]_T, \\
\int_0^T Q_t dP_t &= \int_0^T \alpha_t dQ_t + \int_0^T Q_t dZ_t + [\alpha, Q]_T.
\end{aligned}$$

Since Q is deterministic we have $[\alpha, Q]_T = 0$ and $\alpha_T Q_T = 0$ because $Q_T = 0$ $\mathbb{P}$-a.s. Taking expectations we find:

$$\mathbb{E}\left[\int_0^T Q_t dP_t\right] = \mathbb{E}\left[\int_0^T \alpha_t dQ_t\right],$$

which concludes the proof. □

Using the identity derived in the lemma above we find that the general objective function becomes:

$$\sup_{Q\in\mathcal{D}} \mathbb{E}\left[\int_0^T \alpha_t dQ_t - \frac{1}{2}\int_0^T\int_0^T \widetilde{\Theta}(t,s)G(|t-s|)dQ_s dQ_t - \psi\left(\int_0^T (Z_t - \alpha_t)dQ_t + \int_0^T \mathbb{E}[\alpha_t]dQ_t\right)^2\right].$$

This is the most general objective function we consider. It includes the cost of trading, a risk aversion term and a stochastic alpha signal. Abi Jaber and Neuman (2022) address a similar objective function in continuous time but without the liquidity process Θ and a different measure of risk. They demonstrate that the solution can be explicitly determined by solving a linear Volterra equation using an "operator-valued Riccati equation." For the complete details, we refer to their paper.

In optimal execution literature an alpha signal is often modelled as a mean reverting stochastic process. Therefore, we assume that the alpha signal follows a Ornstein–Uhlenbeck process:

$$d\alpha_t = -\kappa(\mu - \alpha_t)dt + \varsigma dW'_t.$$

The solution of this process is given by:

$$\alpha_t = \alpha_0 e^{-\kappa t} + \mu(1 - e^{-\kappa t}) + \varsigma\int_0^t e^{-\kappa(t-s)}dW'_s,$$

and

$$\mathbb{E}[\alpha_t] = \alpha_0 e^{-\kappa t} + \mu(1 - e^{-\kappa t}).$$

Assume that $Z_t = \int_0^t \sigma dW_s$ such that W and W' are independent Brownian motions. Then the objective functions becomes:

$$\begin{aligned}
&\sup_{Q\in\mathcal{D}} \mathbb{E}\left[\int_0^T \alpha_t dQ_t - \frac{1}{2}\int_0^T\int_0^T \widetilde{\Theta}(t,s)G(|t-s|)dQ_s dQ_t - \psi\left(\int_0^T (Z_t - \alpha_t)dQ_t + \int_0^T \mathbb{E}[\alpha_t]dQ_t\right)^2\right],\\
&= \sup_{Q\in\mathcal{D}} \mathbb{E}\Bigg[\int_0^T \alpha_t dQ_t - \frac{1}{2}\int_0^T\int_0^T \widetilde{\Theta}(t,s)G(|t-s|)dQ_s dQ_t - \psi\Bigg(\int_0^T Z_t dQ_t -\\
&\qquad \int_0^T \left(\alpha_0 e^{-\kappa t} + \mu(1-e^{-\kappa t}) + \varsigma\int_0^t e^{-\kappa(t-s)}dW_s'\right)dQ_t + \int_0^T \alpha_0 e^{-\kappa t} + \mu(1-e^{-\kappa t})dQ_t\Bigg)^2\Bigg],\\
&= \sup_{Q\in\mathcal{D}} \mathbb{E}\left[\int_0^T \alpha_t dQ_t - \frac{1}{2}\int_0^T\int_0^T \widetilde{\Theta}(t,s)G(|t-s|)dQ_s dQ_t - \psi\left(\int_0^T Z_t dQ_t + \int_0^T \varsigma\int_0^t e^{-\kappa(t-s)}dW_s' dQ_t\right)^2\right],\\
&= \sup_{Q\in\mathcal{D}} \mathbb{E}\Bigg[\int_0^T \alpha_t dQ_t - \frac{1}{2}\int_0^T\int_0^T \widetilde{\Theta}(t,s)G(|t-s|)dQ_s dQ_t - \psi_1\left(\int_0^T Z_t dQ_t\right)^2\\
&\qquad -\psi_2\left(\varsigma\int_0^T\int_0^t e^{-\kappa(t-s)}dW_s' dQ_t\right)^2\Bigg],
\end{aligned}$$

where without loss of generality we split the risk aversion parameter into two separate parameters. ψ_1 for the variance of the cost and ψ_2 for the risk aversion w.r.t the alpha signal. We set the risk aversion w.r.t alpha signal equal to zero such that the control problem becomes:

$$\begin{aligned}
&\sup_{Q\in\mathcal{D}} \mathbb{E}\left[\int_0^T \alpha_t dQ_t - \frac{1}{2}\int_0^T\int_0^T \widetilde{\Theta}(t,s)G(|t-s|)dQ_s dQ_t - \psi_1\int_0^T\left(\int_t^T \sigma_s dQ_s\right)^2 dt\right] \qquad (4.7)\\
&\text{s.t.} \quad \int_0^T dQ_t = Q_0.
\end{aligned}$$

The discrete-time version of the above problem we solve analytically. The solution is derived in the following theorem.

Theorem 4.5.1. *Consider the discrete time linear propagator model from Equation 3.4 for which we derived the expected cost of trading in Lemma 3.4.2 and the variance of the cost in Lemma 4.4. Then the discrete version of the optimal execution problem in Equation (4.7) can be formalized in matrix-vector notation as follows:*

$$\begin{aligned}
&\min_{\mathbf{q}\in\mathcal{D}} \mathbf{q}^T(\mathbf{\Phi} + \psi\mathbf{\Sigma})\mathbf{q} - \boldsymbol{\alpha}^T\mathbf{q}\\
&\textit{s.t.} \quad \mathbf{1}^T\mathbf{q} = Q_0^N,
\end{aligned}$$

where the vector $\boldsymbol{\alpha}$ contains the alpha signal which is in this case equal to $\mathbb{E}[\alpha_n]$ for all $n \in [1,T]$.

This is a quadratic optimization problem with optimal solution equal to:

$$\mathbf{q}^* = \frac{1}{2}(\mathbf{\Phi} + \psi\mathbf{\Sigma})^{-1}\left(\boldsymbol{\alpha} - \frac{\mathbf{1}}{\mathbf{1}^T(\mathbf{\Phi} + \psi\mathbf{\Sigma})^{-1}\mathbf{1}}\left(\mathbf{1}^T(\mathbf{\Phi} + \psi\mathbf{\Sigma})^{-1}\boldsymbol{\alpha} - 2Q_0^N\right)\right)$$

Proof. We solve the above problem using the Lagrange multiplier method. Let $\mathbf{\Omega} = (\mathbf{\Phi} + \psi\mathbf{\Sigma})$ The Lagrangian $\mathcal{L}(\mathbf{q}, \nu)$ is given by:

$$\mathcal{L}(\mathbf{q}, \nu) = \mathbf{q}^T\mathbf{\Omega}\mathbf{q} - \boldsymbol{\alpha}^T\mathbf{q} + \nu(\mathbf{1}^T\mathbf{q} - Q_0^N).$$

Calculating the partial derivatives and setting it to zero, we find:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{q}} = 2\boldsymbol{\Omega}\mathbf{q} - \boldsymbol{\alpha} + \nu\mathbf{1} = 0,$$
$$\frac{\partial \mathcal{L}}{\partial \nu} = \mathbf{1}^T\mathbf{q} - Q_0^N = 0.$$

Solving the first equation for $\mathbf{q}$, we find:

$$\mathbf{q} = \frac{1}{2}\boldsymbol{\Omega}^{-1}(\boldsymbol{\alpha} - \nu\mathbf{1}).$$

Substituting the above in the second equation we find:

$$\mathbf{1}^T\left(\frac{1}{2}\boldsymbol{\Omega}^{-1}(\boldsymbol{\alpha} - \nu\mathbf{1})\right) - Q_0^N = 0,$$
$$\mathbf{1}^T\boldsymbol{\Omega}^{-1}\boldsymbol{\alpha} - \nu\mathbf{1}^T\boldsymbol{\Omega}^{-1}\mathbf{1} - 2Q_0^N = 0.$$

Solving for ν, we find:

$$\nu = \frac{\mathbf{1}}{\mathbf{1}^T\boldsymbol{\Omega}^{-1}\mathbf{1}}\left(\mathbf{1}^T\boldsymbol{\Omega}^{-1}\boldsymbol{\alpha} - 2Q_0^N\right)$$

Substituting this in the expression for $\mathbf{q}$, we find that the optimal solution is equal to:

$$\mathbf{q}^* = \frac{1}{2}\boldsymbol{\Omega}^{-1}\left(\boldsymbol{\alpha} - \frac{\mathbf{1}}{\mathbf{1}^T\boldsymbol{\Omega}^{-1}\mathbf{1}}\left(\mathbf{1}^T\boldsymbol{\Omega}^{-1}\boldsymbol{\alpha} - 2Q_0^N\right)\right),$$

Substituting $\boldsymbol{\Omega} = \boldsymbol{\Phi} + \phi\boldsymbol{\Sigma}$ concludes the proof. The solution is a global minimizer since the objective function is convex and quadratic. □

The solution of the optimal execution problem where we include the risk aversion w.r.t alpha signal can be found in Corollary A.3.1 in Appendix A.

To visualize the optimal strategy in the presence of a short-term Ornstein–Uhlenbeck (OU) alpha signal, we simulate a OU process with $\alpha_0 = 50$ bps, $\kappa = 0.1$, $\mu = 0$ and $\varsigma = 3$ bps in a 5 min interval. A number of OU realization and their mean in displaced in the right hand side of Figure 4.8. This short-term alpha signal represents for example a lower opening of a stock where we expect it to revert to some mean value. The optimal strategy with the power-law kernel in the presence of this alpha signal is given in the left hand side of Figure 4.8. We also plot the restricted solution and the risk neutral solution as benchmark.

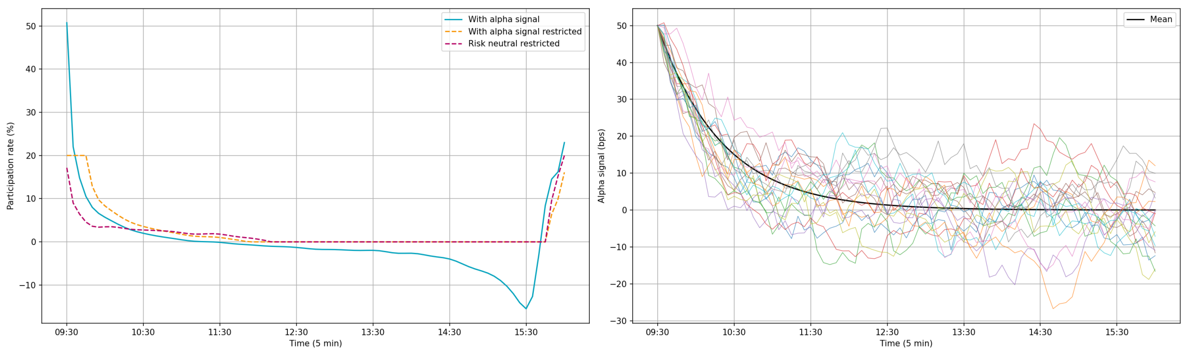

Figure 4.8: (Left) optimal alpha targeting trading strategies and the risk-neutral strategy as benchmark. (Right) 20 realisation of an Ornstein–Uhlenbeck (OU) alpha signal with parameters $\alpha_0 = 50$ bps, $\kappa = 0.2$, $\mu = 0$ and $\varsigma = 3$ bps in a 5 min interval .

We see in the left hand side of Figure 4.8 that the restricted and optimal solution front-load their strategies compared with the risk neutral one. This is also what it expected since in the beginning we can gain the most alpha. We should note that the unrestricted optimal solution again admit transaction-triggered price manipulation, because it sells some shares in the middle of the day to buy them back later.

Before the introduction of the alpha signal the expected cost where only driven by the impact cost. However, in the presence of an alpha signal, the expected cost are also driven the alpha decay. Therefore, in Figure 4.9 we plot the evolution of the market impact and give the decomposition of the cost.

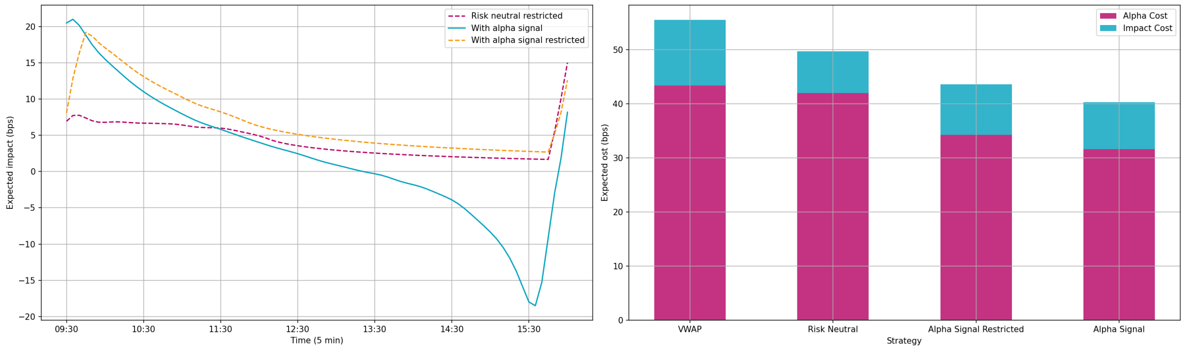

Figure 4.9: (Left) The evolution of the expected impact in bps for the alpha optimal, alpha restricted and risk-neutral strategy. (Right) the decomposition of the expected cost where we include VWAP as benchmark.

We see in the right hand side of Figure 4.9 that there is an trade-off between cost due to market impact and the cost due to the alpha decay. If we compare the risk neutral restricted and alpha optimal restricted than we see that the risk neutral restricted has a lower market impact cost but higher alpha cost. The unrestricted alpha optimal solution has the lowest total cost but is not comparable with the other strategies because it admits transaction-triggered price manipulation.

5

Optimal multiday portfolio rebalancing

In this chapter we extend the single day optimal execution setup to optimizing the execution strategy for a pair of sequential metaorders over two time horizons. This setup is motivated by the large auto-correlation in metaorders of asset managers. We start by defining the expected cost of trading a pair of sequential metaorder, where after we use this as object function to find optimal multiday strategies. Using these results, we quantify the hidden slippage of trading one stock on two consecutive days. Lastly, we run a simulation study to quantify the performance decrease when rebalancing two similar accounts the day after each other.

5.1. The expected cost of trading two adjacent metaorders

We start with quantifying the expected cost of trading for two adjacent metaorders. To the best of our knowledge this is a new research area and we mostly follow the reasoning of Bordigoni et al. (2022) but adapt it to our model specifications. The multiday setup is a interesting extension because single day optimal solutions have the property that they reach a high final impact state, which continues to decay afterwards. Therefore, this should be taken into consideration when calculating the expected cost of trading of the the second metaorder. An illustration of this set up is given in Figure 5.1.

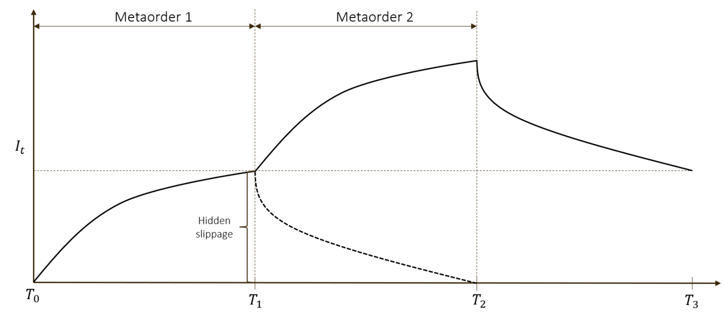

Figure 5.1: Illustration of the evolution of multiday expected impact of two adjacent metaorders.

For simplicity, we only consider the case where the execution of the first metaorder has to be completed before we start trading the second order. Moreover, we assume that we are always trading the same sign because this is in line with what we observe in the data. Lastly, we assume that during the overnight session, i.e. between the closing an opening auction on the next day, the impact does not decay. This is a strong assumption and should be taking into consideration when interpreting the results.

In this chapter, we use the framework of Section 3.3, where we consider the new variant of the linear propagator model with a time-varying liquidity process:

$$S_t = S_0 + \int_0^t \Theta_s G(t-s) dQ_s + Z_t,$$

where $Z \in \mathcal{Z}$ is a continuous martingale. Even though, we derive all theoretical results in continuous time, we solve the optimal execution problems using their discrete-time analogs with a Python solver.

The multiday setup is as follows: the first metaorder trades a quantity of $Q_{[0,1]}$ over the period $[T_0, T_1]$ and the second order trades a quantity of $Q_{[1,2]}$ over the period $[T_1, T_1]$. For the second order we differentiate between 2 cases: The case that $Q_{[1,2]}$ is known at T_0 and the case that $Q_{[1,2]}$ only becomes known at T_1, i.e. after the first metaorder is executed. In the latter, $Q_{[1,2]}$ is then treated as exogenous random variable for all time $t < T_1$, and independent of both the observed mid-price as the unaffected price process.

To derive the expected cost of trading for the two-period problem, we start by the total cost of trading for one period which is derived in Equation (3.5):

$$C(\Pi) = \int_0^T (S_t - S_0) dQ_t + \frac{G(0)}{2} \sum_{t \in \mathcal{J}} \Theta_t (\Delta Q_t)^2.$$

Extending the total cost function to a two-period cost function requires an appropriate benchmark price of the second trade. There are two possibilities: benchmark against the price at T_0, which we call the Two Trade One Cost (TTOC) problem. Or the price at T_1, which we call the Two Trade Separate Cost (TTSC) problem. The price at T_1 corresponds with the price just before we start executing the second metaorder. The price against which we benchmark depends on the moment the order size $Q_{[1,2]}$ becomes known.

Bordigoni et al. (2022) also defines the Two Trade Hybrid Cost approach (TTHC). In this framework, the authors divide the second order size $Q_{[1,2]}$ into a predictable part (known at T_0) which is benchmarked against S_0 and a random part (becomes known at T_1) which is benchmarked against S_{T_1}. We do not go into this and leave it for future research.

Suppose at time T_0, we know the order size $Q_{[1,2]}$ what we are going to trade between $[T_1, T_2]$, then we are in the TTOC scenario. In the following proposition, we derive the total cost of the TTOC problem.

Proposition 5.1.1. *The total expected cost of deterministic admissible strategies* $\Pi^{[0,1]} = (Q_t^{[0,1]})_{t \in [0, T_1]}$ *and* $\Pi^{[1,2]} = (Q_t^{[1,2]})_{t \in [T_1, T_2]}$ *for the TTOC is given by:*

$$\begin{aligned}
\mathbb{E}[C_{TTOC}(\Pi^{[0,1]}, \Pi^{[1,2]})] &= \mathbb{E}\Bigg[\int_0^{T_1} (S_t - S_0) dQ_t^{[0,1]} + \frac{G(0)}{2} \sum_{t \in \mathcal{J}_1} \Theta_t \left(\Delta Q_t^{[0,1]}\right)^2 + \int_{T_1}^{T_2} (S_t - S_0) dQ_t^{[1,2]} \\
&\quad + \frac{G(0)}{2} \sum_{t \in \mathcal{J}_2} \Theta_t \left(\Delta Q_t^{[1,2]}\right)^2 \Bigg] \\
&= \frac{1}{2} \int_0^{T_1} \int_0^{T_1} \widetilde{\Theta}(t,s) G(|t-s|)) dQ_s^{[0,1]} dQ_t^{[0,1]} + \frac{1}{2} \int_{T_1}^{T_2} \int_{T_1}^{T_2} \widetilde{\Theta}(t,s) G(|t-s|) dQ_s^{[1,2]} dQ_t^{[1,2]} \\
&\quad + \int_{T_1}^{T_2} \int_0^{T_1} \Theta_s G(t-s) dQ_s^{[0,1]} dQ_t^{[1,2]}
\end{aligned}$$

Proof. For deterministic strategies we know by Lemma 3.4.1 that the first expression in the expectation is equal to:

$$\mathbb{E}\left[\int_0^{T_1} (S_t - S_0) dQ_t^{[0,1]} + \frac{G(0)}{2} \sum_{t \in \mathcal{J}_1} \Theta_t \left(\Delta Q_t^{[0,1]}\right)^2\right] = \frac{1}{2} \int_0^{T_1} \int_0^{T_1} \widetilde{\Theta}(t,s) G(|t-s|) dQ_s^{[0,1]} dQ_t^{[0,1]}.$$

For the second expression, we have to look carefully at the representation of the stock price. This is because the stock price for $t > T_1$ exist of two parts: the decay of the previous trade and the impact of the new trade. For $t > T_1$, we write the stock price as:

$$S_t = \int_0^{T_1} \Theta_s G(t-s) dQ_s^{[0,1]} + \int_{T_1}^{T_2} \Theta_s G(t-s) dQ_s^{[1,2]} + Z_t,$$

where Z is a martingale. Substituting this into the total cost of trading for the second order we find:

$$\begin{aligned}
&\mathbb{E}\left[\int_{T_1}^{T_2} (S_t - S_0) dQ_t^{[1,2]} + \frac{G(0)}{2} \sum_{t \in \mathcal{J}_2} \Theta_t \left(\Delta Q_t^{[1,2]}\right)^2\right] \\
&= \int_{T_1}^{T_2} \left(\int_0^{T_1} \Theta_s G(t-s) dQ_s^{[0,1]} + \int_{T_1}^{T_2} \Theta_s G(t-s) dQ_s^{[1,2]}\right) dQ_t^{[1,2]} + \frac{G(0)}{2} \sum_{t \in \mathcal{J}_2} \Theta_t \left(\Delta Q_t^{[0,1]}\right)^2 \\
&= \frac{1}{2} \int_{T_1}^{T_2} \int_{T_1}^{T_2} \widetilde{\Theta}(t,s) G(|t-s|) dQ_s^{[1,2]} dQ_t^{[1,2]} + \int_{T_1}^{T_2} \int_0^{T_1} \Theta_s G(t-s) dQ_s^{[0,1]} dQ_t^{[1,2]},
\end{aligned}$$

where the last line follows from Lemma 3.4.1. Combining gives the desired result. □

The proposition above enables us to calculate the hidden cost or slippage in the expected trading cost of the second order as is illustrated by the dotted line in Figure 5.1. The hidden slippage for the second order is then equal to:

$$\int_{T_1}^{T_2} \int_0^{T_1} \Theta_s G(t-s) dQ_s^{[0,1]} dQ_t^{[1,2]},$$

and represents final impact and the decay of the first order. Total expected cost of trading for the second metaorder in the TTOC framework becomes:

$$\mathbb{E}[C_{TTOC}(\Pi^{[1,2]})] = \frac{1}{2} \int_{T_1}^{T_2} \int_{T_1}^{T_2} \widetilde{\Theta}(t,s) G(|t-s|) dQ_s^{[1,2]} dQ_t^{[1,2]} + \int_{T_1}^{T_2} \int_0^{T_1} \Theta_s G(t-s) dQ_s^{[0,1]} dQ_t^{[1,2]} \tag{5.1}$$

However, suppose that the order size of the second order becomes known that time T_1. Then we are in the TTSC scenario and we benchmark the cost of the second trade against the price at time T_1. We know from previous results that the price at T_1 depends on the impact of the trade in the period $[T_0, T_1]$. Therefore, we find:

$$\mathbb{E}[S_{T_1}] = \int_0^{T_1} \Theta_s G(T_1 - s) dQ_s^{[0,1]}$$

Using this result we find the the total expected cost of trading in the TTSC scenario:

Proposition 5.1.2. *The total expected cost of deterministic admissible strategies* $\Pi^{[0,1]} = (Q_t^{[0,1]})_{t \in [0, T_1]}$ *and* $\Pi^{[1,2]} = (Q_t^{[1,2]})_{t \in [T_1, T_2]}$ *for the TTSC is given by:*

$$\begin{aligned}
\mathbb{E}[C_{TTSC}(\Pi^{[0,1]}, \Pi^{[1,2]})] &= \mathbb{E}\Bigg[\int_0^{T_1} (S_t - S_0) dQ_t^{[0,1]} + \frac{G(0)}{2} \sum_{t \in \mathcal{J}_1} \Theta_t \left(\Delta Q_t^{[0,1]}\right)^2 + \int_{T_1}^{T_2} (S_t - S_{T_1}) dQ_t^{[1,2]} \\
&\quad + \frac{G(0)}{2} \sum_{t \in \mathcal{J}_2} \Theta_t \left(\Delta Q_t^{[1,2]}\right)^2\Bigg] \\
&= \frac{1}{2} \int_0^{T_1} \int_0^{T_1} \widetilde{\Theta}(t,s) G(|t-s|) dQ_s^{[0,1]} dQ_t^{[0,1]} + \frac{1}{2} \int_{T_1}^{T_2} \int_{T_1}^{T_2} \widetilde{\Theta}(t,s) G(|t-s|) dQ_s^{[1,2]} dQ_t^{[1,2]} \\
&\quad + \int_{T_1}^{T_2} \int_0^{T_1} \Theta_s G(t-s) dQ_s^{[0,1]} dQ_t^{[1,2]} - Q_{[1,2]} \int_0^{T_1} \Theta_s G(T_1 - s) dQ_s^{[0,1]}
\end{aligned}$$

Proof. Direct consequence of Proposition 5.1.1. □

The difference between the TTOC and the TTSC frameworks lie in the handling of the final impact state of the first metaorder. Because in the TTSC scenario the expected cost of the second order is benchmarked against S_{T_1}, the final impact state of the first metaorder and the benchmark price S_{T_1} cancels out. Therefore, the TTSC only includes the decay starting form zero. The result is that, when trading the same order size, the second metaorder has a lower expected cost than the first order due to the expected reversion. This should be carefully taken into consideration when using the TTSC framework.

5.2. Optimal multiday execution strategies

In this section, we propose different optimal multiday execution strategies that minimize the expected cost of trading two adjacent metaorders in different ways. These strategies are illustrated by means of a realistic example.

In TTOC problem, we benchmark the two adjacent metaorders against the arrival price of the first order. This means that the order size of the second trade is known at time T_0 but can not be executed till T_1. To execute two adjacent metaorders in the TTOC framework, we propose the following strategies:

- Separately optimal: Find strategies that minimizes the cost for each order separately.
- Impact optimal: Minimize cost for the first metaorder and include the final impact and decay of the first order as short-term alpha signal in the optimization of the second order.
- Combined optimal: Minimize the cost for both orders combined.

The reason why we consider three different strategies is that all of them deal differently with the impact decay of the first order. The first two strategies can also be used in the TTSC problem but not the last one. The combined optimization strategy is only possible when we know the order size of the second order at T_0.

Since the optimal execution problem for two adjacent metaorders is a new research direction, to the best of our knowledge, no results are yet proven to exclude price manipulation from the model. To exclude any possibilities of a round trip trade during the execution, we restrict the solver to only one way trading. In addition, we set the maximum participation rate to 20%, as in the previous section.

To find the optimal solutions for the 'separately optimal' strategy, one should solve the optimal execution problem in Equation (4.4) twice for both periods:

$$\begin{aligned}&\min_{Q\in\mathcal{Q}} \frac{1}{2}\int_0^T\int_0^T \widetilde{\Theta}(t,s)G(|t-s|)dQ_s dQ_t\\ &\text{s.t.} \quad \int_0^T dQ_t = Q_0, \quad Q_t \geq 0 \quad \text{and} \quad Q_t \leq 0.2V_t \quad \forall t \in [0,T].\end{aligned}$$

For the 'impact optimal' strategy, we propose a two-step approach: first, solve the optimal execution problem above for the first order. Then, use the final impact state of the first order and its subsequent decay in the optimization of the second order. In this approach, we minimize the expected cost of trading for the second metaorder as derived in Equation (5.1). The optimization problem for the second order in the 'impact optimal' strategy becomes:

$$\begin{aligned}&\min_{Q^{[1,2]}\in\mathcal{Q}} \frac{1}{2}\int_{T_1}^{T_2}\int_{T_1}^{T_2} \widetilde{\Theta}(t,s)G(|t-s|)dQ_s^{[1,2]}dQ_t^{[1,2]} + \int_{T_1}^{T_2}\int_0^{T_1}\Theta_s G(t-s)dQ_s^{[0,1]}dQ_t^{[1,2]} \qquad (5.2)\\ &\text{s.t.} \quad \int_0^T dQ_t^{[1,2]} = Q_{[1,2]}, \quad Q_t^{[1,2]} \geq 0 \quad \text{and} \quad Q_t^{[1,2]} \leq 0.2V_t \quad \forall t \geq 0.\end{aligned}$$

We solve the discrete version of this problem using the solver *scxpy* in Python.

For 'combined optimal' strategy one should use Proposition 5.1.1 as objective function and solve for $Q^{[0,1]}, Q^{[1,2]}$ combined. The minimization problem is then given by:

$$
\begin{aligned}
\min_{Q^{[0,1]},Q^{[1,2]}\in\mathcal{Q}} & \frac{1}{2}\int_0^{T_1}\int_0^{T_1}\widetilde{\Theta}(t,s)G(|t-s|))dQ_s^{[0,1]}dQ_t^{[0,1]}+\frac{1}{2}\int_{T_1}^{T_2}\int_{T_1}^{T_2}\widetilde{\Theta}(t,s)G(|t-s|)dQ_s^{[1,2]}dQ_t^{[1,2]} \\
& +\int_{T_1}^{T_2}\int_0^{T_1}\Theta_s G(t-s)dQ_s^{[0,1]}dQ_t^{[1,2]} \qquad (5.3)\\
\text{s.t.} \quad & \int_0^{T_1} dQ_t^{[0,1]}=Q_{[0,1]}, \quad \int_{T_1}^{T_2} dQ_t^{[1,2]}=Q_{[1,2]}, \quad Q_t^{[0,1]}\geq 0, \quad Q_t^{[1,2]}\geq 0,\\
& Q_t^{[0,1]}\leq 0.2V_t \quad \text{and} \quad Q_t^{[1,2]}\leq 0.2V_t \quad \forall t\geq 0.
\end{aligned}
$$

We solve the discrete version of this problem using the solver *"SLSQP"* within *scipy.optimize.minimize* in Python.

To illustrate these theoretical results, we calculate the optimal strategies using the before mentioned solvers in Python. For this we use the same settings as in the previous chapter. We use the same intraday volume curve as in Figure 2.5, assume $\sigma = 15$ bps in a 5 min interval and we let ADV be equal to 1000000 shares. Furthermore, both orders are 5% ADV orders and we use the power-law kernel for the calculations. The result in displayed in Figure 5.2.

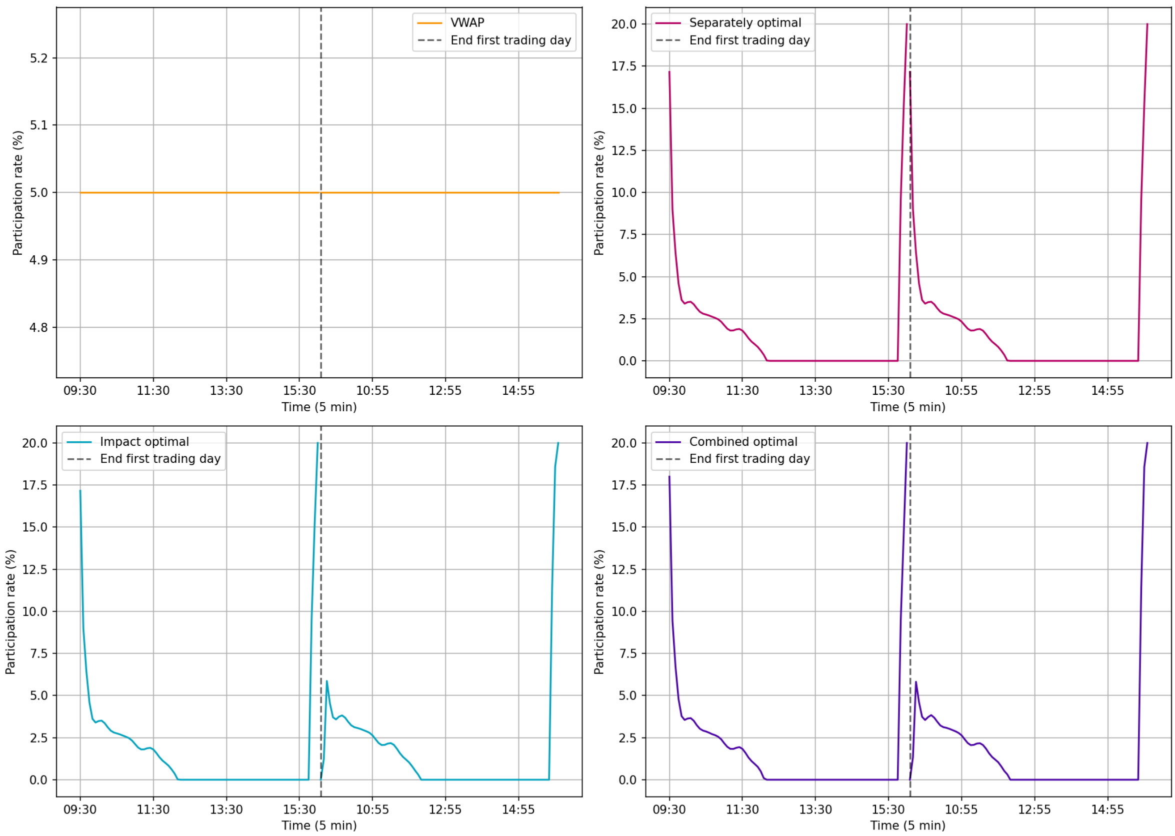

Figure 5.2: Optimal multiday execution strategies using a power-law decay kernel. (Upper left) VWAP strategy. (Upper right) Separately optimal strategy. (Lower left) Impact optimal strategy. (Lower right) Combined optimal strategy.

We see in Figure 5.2 that the combined optimal strategy has a lower participation rate in the last bin of the first day and does not trade in the first bin of the second day. The optimizer aims to let the impact for the first day decay before starting to trade the second order. A similar result holds for the 'impact optimal' strategy, where the optimal solution for the second order shows similar behaviour. The evolution of the expected impact and expected cost are displayed in Figure 5.3 and 5.4.

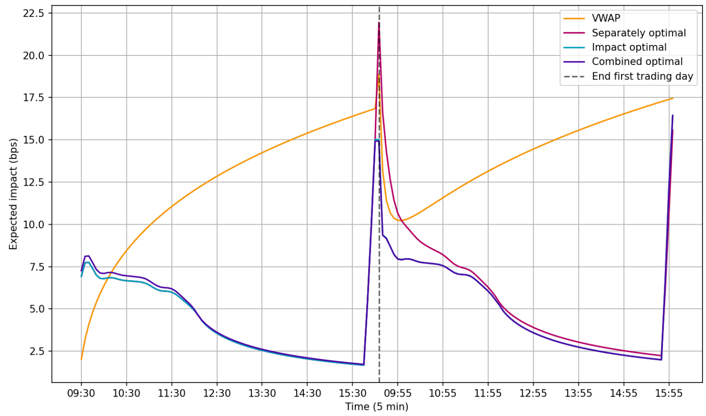

Figure 5.3: Evolution of multiday expected impact for different strategies using a power-law decay kernel.

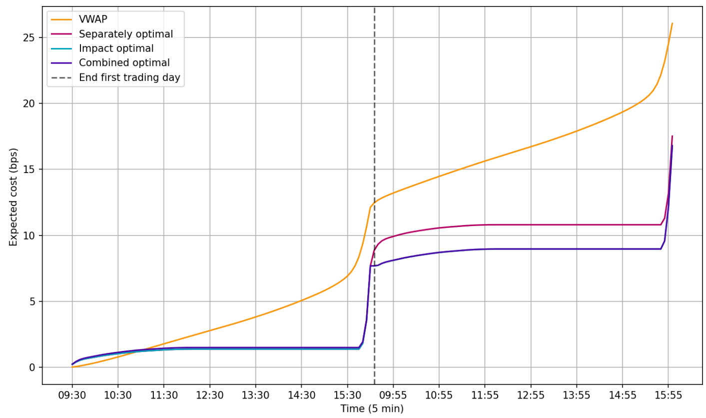

Figure 5.4: Evolution of multiday expected cost for different strategies using a power-law decay kernel.

In Figure 5.3 we see the evolution of the impact during the execution of the two adjacent metaorders. As the optimal solution for the combined and impact optimal strategies already suggested, the optimizer tries to minimize the impact at the end of the first day and the beginning of the second day. Since we do not trade on the third day, we not care about the high impact state at the end of the second day.

The evolution of the cost is displayed in Figure 5.4. In this figure we see that the VWAP strategy has the highest expected trading cost and the combined optimal the lowest. However, the 'impact optimal' strategy is very close to the combined optimal with only 0.5 bps difference.

To get a better understanding of the expected cost of each strategy and their relative difference, we compare them in Table 5.1. In the 'out-performance' column we compare the cost of each strategy with the cost of the VWAP strategy. Moreover, in the 'Q_2 vs Q_1' column we compare the cost of trading a strategy in the second period with trading the optimal strategy in the first period.

In Table 5.1 we make the following two important observations. In the 'Q_2 vs Q_1' column, we see that it is on average average 20% more expensive to trade an optimal strategy on the second day compared with an optimal strategy in the first period and even 80% when using a VWAP strategy. This is significant but as expected because we compare the cost with initial price S_0. The second important observation is that it is possible to decrease the cost significantly by trading market impact optimal strategies. Compared with VWAP, we can decrease the cost by 34%.

We should note that the combined optimal solution is not really a realistic strategy in the sense that in most cases we do not know the order size of the second trade a day before. Furthermore, it is very difficult for a trader to exactly follow the given optimal schedule for two days in a row.

However, the 'impact optimal' strategy seems a good alternative, it has almost the same cost as the

Table 5.1: Expected cost calculation for the first and second metaorder in the first columns. Relative difference in cost in the second and third column.

Strategy $Q_{[0,1]}$	Expected cost $Q_{[0,1]}$ (bps)	Out-performance (%)	
VWAP	12.12	0.0	
Separately Optimal	7.685	-36.639	
Impact Optimal	7.685	-36.639	
Combined Optimal	7.688	-36.614	
Strategy $Q_{[1,2]}$	Expected cost $Q_{[1,2]}$ (bps)	Out-performance (%)	$Q_{[1,2]}$ vs $Q_{[0,1]}$ (%)
VWAP	13.933	0.0	81.80
Separately Optimal	9.832	-29.430	27.94
Impact Optimal	9.116	-34.567	18.62
Combined Optimal	9.110	-34.611	18.54

combined optimal strategy but it allows for more flexibility. Moreover, since it is possible to incorporate the realized impact of previous trades, it also allows to intraday re-optimization. For example, a trader tries to follow an optimal participation schedule but halfway he realizes that the realized schedule deviates significantly from the pre-determined strategy. This optimization procedure allows to re-optimize with a different volatility and volume prediction such that the overall cost of the order w.r.t arrival price is minimized.

5.3. Optimal multiday portfolio rebalancing

In this section, we delve into the complexity of managing multiple investment accounts within a single strategy. Through a simulation study, we aim to quantify the performance decrease for an account that consistently trades a day after a similar account.

Major asset managers like Robeco offer a variety of investment strategies to their clients. For the quantitative investment side of Robeco these including factor strategies, conservative strategies, and enhanced indexing strategies. When a large institutional client decides to invest in one of these strategies, a new account is created, customized to the client's requirements but generally adhering to the chosen strategy.

To ensure that the accounts remain aligned with the strategies, Robeco performs monthly rebalancing. This process involves determining which stocks to buy or sell, guided by a stock ranking system. The ranking is primarily driven by long-term alpha signals (spanning more than three months) and is updated daily.

Robeco manages multiple accounts for each strategy. Consequently, accounts within the same strategy may need to be rebalanced on consecutive days. Due to the reliance on long-term alpha signals, the stock ranking is unlikely to change significantly overnight. This situation could lead to recommendations to buy or sell the same stocks on successive days. For this reason we compare all cost against S_0. As a result, due to the hidden slippage, we always obtain a higher expected cost of trading for the second order as is shown in Table 5.1.

When a strategy has only a limited number of accounts one could just put the rebalances as far a part as possible. However, when a strategy has many accounts (> 20), it is not possible to separate the rebalances anymore and one should consider more sophisticated approaches:

1. *Stock selection level*: do not rebalance the same stocks but use different stocks with almost the same ranking. Or include the impact of the previous execution as penalty in the objective function for the stock selection.
2. *rebalance schedule level*: randomize the days on which accounts are rebalanced.
3. *Execution level*: optimize the execution strategies by incorporating the impact from the previous trade.

In this section, we do a simulation study to see the effect of using the second and third approach in portfolio rebalancing. In particular, we quantify the hidden total cost when rebalancing two similar accounts on consecutive days. We leave the first approach for future research.

The simulation setting is as follows:

- *Universe*: the universe are 609 stock from the MSCI US index. Each stock has its own intraday volatility and ADV based on historical data from 2023 but use the same intraday volume curve which has the same shape as Figure 2.5.
- *Accounts*: We consider two accounts that rebalance on consecutive days. Both accounts rebalance 50 stocks, with 20 stocks overlapping between them. Additionally, we vary the number of overlapping stocks in subsequent analyses. During the two rebalances, we maintain constant intraday volatility and ADV.
- *Orders*: We sample the order sizes from a exponential distribution which is fitted to the observed order sizes. This means that for the non-overlapping stocks, each one gets his own order size and are different for both accounts. The order sizes for the overlapping stocks are also different but the same between the two accounts. The simulated order sizes are given in Figure 5.5. Note that both distributions look similar but that the average order size deviate a little. Moreover, with the black dotted line we indicate the average order size of the overlapping orders.
- *Scenarios*: We compare two cases in the simulation study. The first is the base case and represent the the expected trading cost when we separate the rebalances. That is when we have more than one day in between and this is our benchmark. The second scenario is when the accounts are rebalanced right after each other.
- *Execution strategies*: we use the execution strategies discussed in the previous section except the combined optimal strategy because this strategy is not implementable in reality.

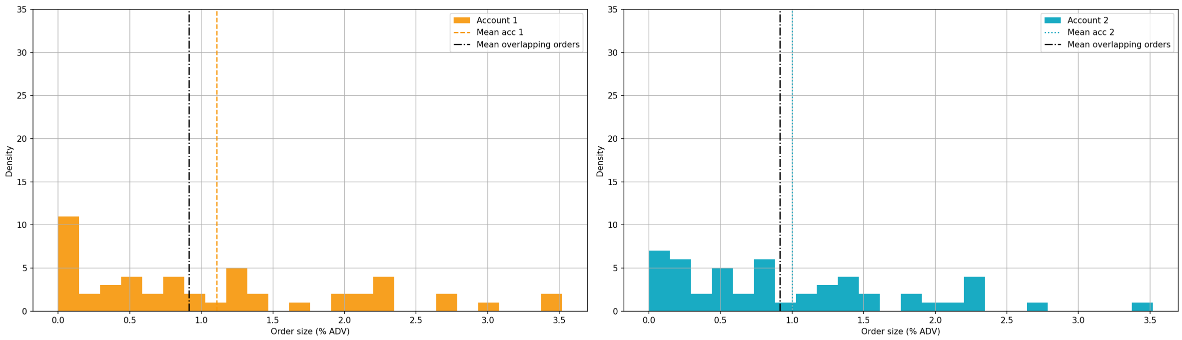

Figure 5.5: Histograms of 50 simulated order sizes for the rebalance simulation. The first account on the left and the second on the right. The orders are drawn from an exponential distribution fitted to our historical order sizes.

In the simulation study we assume that the stock-ranking does not change overnight due to long-term alpha signals such that we can use the TTOC framework to quantify the hidden slippage. Therefore, to calculate the cost of trading on the first day we use Equation (3.7). For trading on the second day, right after the execution of the first order we use Equation (5.1), since we need to account for the impact state and decay of the previous execution.

In Table 5.2 we calculate the cost of rebalancing the two accounts. The cost measure we use is the average expected cost of trading per stock in bps. The first column 'Expected Cost $[T_0, T_1]$ (bps) ' can be seen as the benchmark cost and represents trading on the first day or when we have separated the rebalance. The second column 'Expected Cost $[T_1, T_2]$ (bps) ' represents trading on the second day directly after the first rebalance. Lastly, the 'Difference' column shows how much more expensive trading on the second day is compared with trading an optimal strategy on the first day.

In Table 5.2, we see that the cost of the rebalance is lowest in the benchmark scenario, i.e. when we separate the rebalances such that there is minimal one day in between or if the rebalance happens on the first day. Furthermore, we see that if the rebalance happens on the second day, directly after the

Table 5.2: Expected cost analysis of rebalancing two accounts right after each other. Cost measure is the average expected cost of trading per stock in bps.

Account 1	Expected Cost $[T_0, T_1]$ (bps)	Expected Cost $[T_1, T_2]$ (bps)	Difference (%)
VWAP	2.870	3.008	65.83
Separately optimal	1.816	1.981	9.07
Impact optimal	1.816	1.923	5.91
Account 2	**Expected Cost $[T_0, T_1]$ (bps)**	**Expected Cost $[T_1, T_2]$ (bps)**	**Difference (%)**
VWAP	2.327	2.464	67.40
Separately optimal	1.472	1.637	11.19
Impact optimal	1.472	1.580	7.29

first rebalance we always incur a higher cost (≈ 0.1 to 0.2 bps). This also what is expected since we have 20 overlapping stocks that are traded on two consecutive days.

The expected cost are approximately 5% to 10% higher in case we use an optimal strategy and even 65% in case we use a VWAP strategy. Zooming in on the different execution strategies used, we find that VWAP is always the most expensive one. Therefore, it is possible to drastically decrease the expected cost by trading optimal strategies. This is also what is expected from Table 5.1. We should note however that we assume that during the overnight session, the impact of the previous order does not decay. This could lead to overestimating the cost of the second order.

In the above simulation study we used a fixed number of overlapping stocks (20%), however the hidden slippage is a function of the overlapping stocks. Therefore, we run a Monte-Carlo (MC) simulation but this time we vary the number of overlapping stocks to quantify the absolute increase in cost per order. We use a MC simulation to make sure that on average the average order sizes for the two accounts are equal. We compare the cost with the case that there are zero overlapping stocks which also corresponds to the case in which we separate the rebalances. We refer to this case as the base case. In this simulation we compare the VWAP strategy with the impact optimal strategy. The results are shown in Figure 5.6.

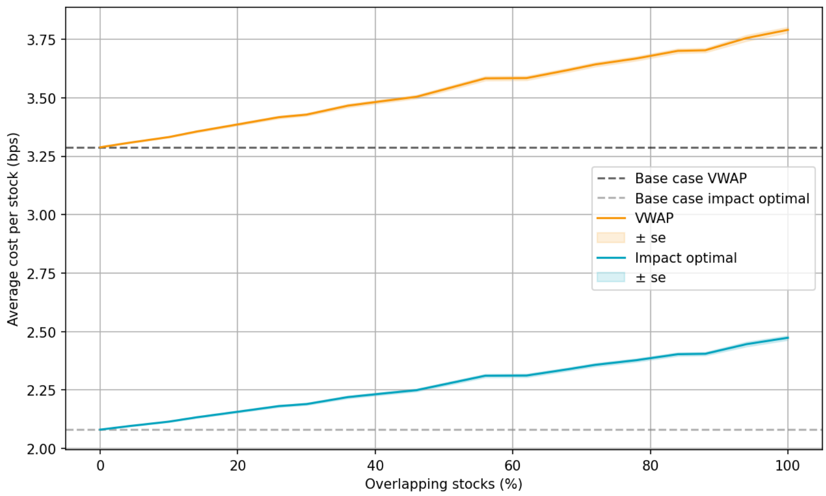

Figure 5.6: The average expected cost in bps per order for different percentages of overlapping stocks. The average order size is around 1% ADV and we use 50 MC samples.

In Figure 5.6 we find that for both strategies the cost increases linearly when the percentage of overlapping stocks increases. The increase is fastest in case of the VWAP strategy. Therefore, we conclude from this figure that the impact optimal strategy performs best on two fronts. In absolute terms the average cost is lowest but also the increase in cost when the number of overlapping stocks increases is slowest.

An interesting observation is that when we trade 100% of the same stocks, the cost increases by only approximately 20% for the impact optimal case. This aligns with the results shown in the 'Q_2 vs Q_1' column of Table 5.1. The primary reason for this is the sharp and rapid reversion of the impact on the second day, along with the absence of any permanent impact.

The simulation assumes that the rankings does not change due to new information. When the raking is only based on long-term signals this is a valid assumption and this simulation study shows that the hidden slippage can lead to significant under performance. However, when the raking does change a lot over night, it is not realistic to benchmark the second metaorder against S_0. In this scenario, we should use the TTSC framework because there exist no sequential trade cost anymore. This also allows us to use the decay of the previous impact as alpha signal in the optimization.

In conclusion, our simulation study demonstrates that rebalancing an account the day after a similar account trading the same stocks results in a significant hidden cost for the second account. This cost can be substantially reduced by employing an optimal trading strategy. However, even with the 'impact optimal' strategy, there can still be an additional cost of up to 20%, depending on the overlap of stocks traded. In practice, this hidden slippage is often not observed in transaction cost analysis because costs are typically benchmarked against the stock price just before execution begins. Consequently, when trading the same stock on consecutive days, the second order may even appear to have a lower expected cost due to the reversion of the first order's impact. Therefore, it's crucial to consider this factor in transaction cost analysis, especially when dealing with sequential metaorders, and when the stock ranking remains unchanged overnight due to reliance on long-term signals.

Returning to the three proposed approaches for rebalancing multiple similar accounts, namely, stock selection, optimizing the rebalance schedule, and execution strategies, we believe the most significant gains can be achieved at the stock selection level. By calculating the final impact state of orders executed on a given day and incorporating this as a penalty in the stock selection algorithm's objective function, we can avoid selecting stocks with a high impact state.

Moreover, if consecutive orders do arrive at the trading desk, traders can handle these by considering the impact reversion as a short-term alpha signal in their strategy. This approach allows for more effective management of the hidden costs associated with trading the same stocks on consecutive days. However, addressing the issue at its root, through careful stock selection, remains the most effective solution.

6 Conclusion

This thesis provides a comprehensive examination of the propagator market impact model and its associated optimal execution problem in both intraday and multiday settings. By exploring the theoretical foundations and practical implementations of these models, and leveraging Robeco's proprietary intraday order database, this thesis bridges the gap between academia and practice. The extensive review of existing literature and thorough examination of the order data led to the consideration of the propagator model as a suitable market impact model. It captures the significant autocorrelation in order flow and the transient nature of market impact observed in our data. The propagator model serves as the foundation for all applications discussed in this thesis.

In this thesis we introduced a new variant of the linear propagator model, which combines a general decay kernel with a time-varying liquidity process. We use a linear model because a non-linear model in the trading rate allows for price manipulation and makes the optimal execution problem intractable. Additionally, we introduce a time-varying liquidity process because the instantaneous impact scales less concave and more linearly with the participation rate, then it does with the usual normalization by the average daily volume. This time-varying liquidity process better reflects the variations in intraday liquidity, leading also to a better model fit. Within this framework, we derive sufficient conditions on the decay kernel and the liquidity process to ensure that the model is free of price manipulation in the sense of dynamic arbitrage.

To bridge the gap between theory and practice, we calibrated the model using Robeco's proprietary intraday order database. This is achieved through a two-step calibration approach. This method is particularly efficient because the first and most computationally intensive step has an analytical solution, which significantly speeds up the process and provides valuable insights into the underlying structure of the data. Additionally, access to Robeco's proprietary intraday order database enables detailed analyses of the model's behavior on proprietary order data. These results are rarely available in the literature, because it is not possible to identify the origin of an order in publicly available datasets.

From the detailed analysis of the calibrated model, we made several important observations. Firstly, regarding the estimates of the decay kernel, we found out that the non-projected decay kernel exhibits a two-stage decay: an initial rapid decay followed by a prolonged, slower decay. This pattern results in a short half-life of the impact. Additionally, we found no evidence of permanent market impact in the estimates. We also conclude that there is no single best admissible projection of the kernel estimate, as all projections have different characteristics and perform almost equally. As for the model performance, we found a R-squared value that is consistent with our expectations based on the low signal-to-noise ratio of our data.

Building on the new variant of the linear propagator model, we addressed its associated optimal execution problem. Starting from the self-financing equation, we demonstrated that maximizing the trader's expected profit and loss is equivalent to minimizing the expected cost of trading. Furthermore, we showed that the solution of the OW optimal execution problem coincides with the solution of the optimal execution problem under the discrete linear propagator with an exponential kernel and constant liquidity process. To make the optimal execution problem applicable in more realistic settings, we extended the problem by formulating it as a mean-variance optimization problem and incorporate short-term alpha signals modeled as a Ornstein–Uhlenbeck process. The analytical solutions derived for the discrete analogs of these optimization problems are supported with realistic examples. In these examples, we additionally used a solver to restrict the solution space to prevent two-way trading, which is typically prohibited for asset managers. Although this approach leads to sub-optimal solutions, it still significantly outperforms the VWAP strategy in mean-variance framework.

We extended the intraday optimal execution framework to a multiday framework, motivated primarily by the significant autocorrelation in the order flow of metaorders. This extension is particularly relevant for asset managers with a large number of investment accounts following the same strategy, making consecutive-day trading of the same stock sometimes unavoidable. To address this, we developed a framework to calculate the expected cost of trading adjacent metaorders, which we used to determine optimal execution strategies in a multiday setup. Additionally, we quantified the hidden cost of trading the same stock on consecutive days. We extended this analysis to the portfolio level, where we examine the hidden cost of rebalancing similar accounts as a function of overlapping stocks.

The final simulation study highlights that rebalancing an account the day after a similar account trading the same stocks results, in a significant hidden costs for the second account. Even with an "impact optimal" strategy, additional costs of up to 20% may still occur due to the overlap in stocks. These costs are often overlooked in transaction cost analysis as they are benchmarked against the stock price before execution. To address the issue of rebalancing similar accounts on consecutive days, we proposed three approaches. Among these, we argued that the most significant gains can be achieved at the stock selection level. By calculating the final impact state of executed orders and incorporating it as a penalty in the stock selection algorithm, we can avoid high impact stocks and select those with similar rankings but lower hidden costs. Additionally, traders can manage consecutive orders more efficiently by considering impact reversion as a short-term alpha signal.

7

Discussion and future research

The discussion section of this thesis delves into some of the implications and limitations of the results presented. Additionally, we provide some suggestions for further research. We start by highlighting some limitations of our dataset. The data is constructed such that we only have intraday returns without any relation to other days. This limitation means we cannot measure the decay of trades executed in the last trading bin into the next day. Consequently, while there appears to be some evidence of permanent impact (see Figure 2.7), we cannot measure it because many metaorders have child orders close to the end of the day (see Figure C.1). As a result, our estimates of decay after the end of the metaorder and into a new day are extrapolations of the intraday decay observed after a child order. Additionally, this extrapolation might not accurately reflect reality, because our kernel estimate is for the continuous session, excluding the opening and closing auctions. The market behaves differently during these auction sessions, and thus the kernel might not be the same. Further research is needed to explore how a kernel can be simultaneously calibrated for both the continuous session and the auctions and whether there are notable differences between the estimates.

We continue with the limitations of the linear propagator model with a time-varying liquidity process introduced in this thesis. A key modeling choice was the use of linear scaling for the participation rate of every child order. While we extensively discussed its local-level implications in Appendix B.2, we did not explore its global-level impact. Empirical observations indicate that the impact or cost of a metaorder scales concavely with respect to order size. However, in our linear model based on the participation rate, this scaling is linear (see Figure 3.8). As a result, the model tends to underestimate the total impact for small orders and overestimate it for large orders. This discrepancy highlights a key limitation of the linear approach and suggests incorporating a global concave scaling, similar to the AFS model (see, Alfonsi et al. (2012)). In particular, the paper by Hey, Mastromatteo, et al. (2023) provides some promising results as it combines global connectivity with multiple exponential decay parameters and time dependent liquidity parameter, while the optimal execution problem remains tractable.

Another important assumption in the propagator model is that the decay kernel is time-invariant, meaning the decay kernel estimate remains constant throughout the day. Testing the validity of this assumption for a general decay kernel is challenging due to identification issues. Therefore, further research is needed to find out whether this assumptions holds. When this assumption is violated, the model might under of over estimate the impact. Consequently, affecting the solutions to the corresponding optimal execution problems. Possible solutions to mitigate this include: the introduction of a correction factor in the liquidity process or using volume time instead of regular clock time.

Moving on to the optimal execution problem, we solved all problems analytically using their discrete time analogs. We demonstrated that the solution to the OW optimal execution problem coincides with the solution of the discretized optimal execution problem for the linear model with an exponential kernel and constant liquidity process. However, we did not explicitly show that the optimal execution problem for the continuous time model with a time-varying liquidity process is the continuous time limit of its discrete version. Therefore, deriving analytical solutions in continuous time for the linear propagator model with a general decay kernel and time-varying liquidity process would be interesting for future research, despite its complexity. Additionally, it would be interesting to extend the problem to include a stochastic trading strategies, which includes the uncertainty of immediate execution or finding a block trade in a crossing network. Or one that dynamically adapts to a stochastic alpha signal. This could potentially be achieved using machine learning methods, such as reinforcement learning.

A

Appendix A: Additional proofs

A.1. Proof Lemma 3.4.2

Lemma. *Consider the discrete time version of the linear propagator model:*

$$S_n^N = S_0^N + \sum_{m=1}^{n} \Theta_m^N G(n-m)\Delta_m Q^N + Z_n^N,$$

where S_n^N is the mid-price just before time t_n^N. Then for a discrete admissible trading strategy $\Pi^N = Q^N$, the expected cost of trading in discrete time equals:

$$\mathbb{E}[C(\Pi^N)] = \mathbb{E}\left[\mathbf{q}^T\left(\frac{1}{2}\widetilde{\mathbf{\Theta}} \odot \mathbf{G}\right)\mathbf{q}\right] = \mathbb{E}\left[\mathbf{q}^T\mathbf{\Phi}\mathbf{q}\right],$$

where $q_n = \Delta_n Q^N$, $\widetilde{\Theta}_{n,m} = \Theta_{\min(n,m)}$, $G_{n,m} = G(|n-m|)$ and the operator $\odot$ means the Hadamard product between the matrices $\widetilde{\mathbf{\Theta}}$ and $\mathbf{G}$. Notice that $\widetilde{\mathbf{\Theta}}$ and $\mathbf{G}$ are both symmetric matrices. For simplicity we define $\mathbf{\Phi} = \frac{1}{2}\mathbf{\Theta} \odot \mathbf{G}$, which is a symmetric matrix as well.

Proof. We derive the expected cost of trading for a discrete admissible trading strategy $\Pi^N = (Q_n^N)_{n=0}^N$ in a similar way as in Lemma 3.4.1. The total cost of trading is given by:

$$C(\Pi^N) = \sum_{n=1}^{N}(S_n^N - S_0^N)\Delta_n Q^N.$$

Substituting the discrete model, we find:

$$C(\Pi^N) = \sum_{n=1}^{N}\sum_{m=1}^{n} \Theta_m^N G(n-m)\Delta_m Q^N \Delta_n Q^N + \sum_{n=1}^{N} Z_n^N \Delta_n Q^N.$$

Then for the first part we observe that:

$$\sum_{n=1}^{N}\sum_{m=1}^{n} \Theta_m^N G(n-m)\Delta_m Q^N \Delta_n Q^N = \sum_{n=1}^{N}\sum_{m=1}^{n} \Theta_m^N G(|n-m|)\Delta_m Q^N \Delta_n Q^N,$$

which is equal to:

$$= \frac{1}{2}\sum_{n=1}^{N}\sum_{m=1}^{n-1} \Theta_m^N G(|n-m|)\Delta_m Q^N \Delta_n Q^N + \frac{1}{2}\sum_{n=1}^{N}\sum_{m=n+1}^{N} \Theta_n^N G(|n-m|)\Delta_m Q^N \Delta_n Q^N$$
$$+ \sum_{n=1}^{N} \Theta_n^N G(0)\left(\Delta_n Q^N\right)^2$$

To be consistent with the continuous case, we assume that we trade uniformly in a time interval such we pay the average instantaneous impact. This results in:

$$= \frac{1}{2}\sum_{n=1}^{N}\sum_{m=1}^{n-1}\Theta_m^N G(|n-m|)\Delta_m Q^N \Delta_n Q^N + \frac{1}{2}\sum_{n=1}^{N}\sum_{m=n+1}^{N}\Theta_n^N G(|n-m|)\Delta_m Q^N \Delta_n Q^N$$
$$+ \frac{1}{2}\sum_{n=1}^{N}\Theta_n^N G(0)\left(\Delta_n Q^N\right)^2.$$

Ensuring causality and simplifying gives:

$$= \frac{1}{2}\sum_{n=1}^{N}\sum_{m=1}^{N}\Theta_{\min(n,m)}^N G(|n-m|)\Delta_m Q^N \Delta_n Q^N,$$

such that the cost of trading is equal to:

$$C(\Pi^N) = \frac{1}{2}\sum_{n=1}^{N}\sum_{m=1}^{N}\widetilde{\Theta}_{n,m}^N G(|n-m|)\Delta_m Q^N \Delta_n Q^N + \sum_{n=1}^{N} Z_n^N \Delta_n Q^N,$$

where $\widetilde{\Theta}_{n,m}^N = \Theta_{\min(n,m)}^N$. Then using that Q is adapted and Z a martingale, we find that the expected cost of trading equals:

$$\mathbb{E}[C(\Pi^N)] = \mathbb{E}\left[\frac{1}{2}\sum_{n=1}^{N}\sum_{m=1}^{N}\widetilde{\Theta}_{n,m}^N G(|n-m|)\Delta_m Q^N \Delta_m Q^N\right],$$

which in matrix-vector notation is equal to:

$$\mathbb{E}[C(\Pi^N)] = \mathbb{E}\left[\mathbf{q}^T\left(\frac{1}{2}\widetilde{\mathbf{\Theta}} \odot \mathbf{G}\right)\mathbf{q}\right],$$

where $q_n = \Delta_n Q^N$, $\widetilde{\Theta}_{n,m} = \Theta_{\min(n,m)}$, $G_{n,m} = G(|n-m|)$ and the operator $\odot$ means the Hadamard product between the matrices $\widetilde{\mathbf{\Theta}}$ and $\mathbf{G}$. Notice that $\widetilde{\mathbf{\Theta}}$ and $\mathbf{G}$ are both symmetric matrices. For simplicity we define $\mathbf{\Phi} = \frac{1}{2}\mathbf{\Theta} \odot \mathbf{G}$, which is symmetric as well. □

A.2. Proof Lemma A.2.1

Lemma A.2.1. *Let $\mathbf{g} = [g_0, g_1, \dots, g_{M-1}]^T$ and $\mathbf{u} = [u_0, u_1, \dots, u_{M-1}]^T$ be vectors, and let $\mathbf{G}$ and $\mathbf{U}$ be their corresponding lower triangular Toeplitz matrices. Then $\mathbf{Ug} = \mathbf{Gu}$.*

Proof. The matrices $\mathbf{G}$ and $\mathbf{U}$ can be represented as:

$$\mathbf{G} = \begin{bmatrix} g_0 & 0 & \cdots & 0 \\ g_1 & g_0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ g_{M-1} & g_{M-2} & \cdots & g_0 \end{bmatrix},$$
$$\mathbf{U} = \begin{bmatrix} u_0 & 0 & \cdots & 0 \\ u_1 & u_0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ u_{M-1} & u_{M-2} & \cdots & u_0 \end{bmatrix}.$$

For the matrix-vector multiplication $\mathbf{Ug}$, the i-th element of the resulting vector is:

$$(\mathbf{Ug})_i = \sum_{j=0}^{i} u_j g_{i-j}, \quad \text{for } i = 0, 1, \dots, M-1.$$

Similarly, for $\mathbf{Gu}$, the i-th element of the resulting vector is:

$$(\mathbf{Gu})_i = \sum_{j=0}^{i} g_j u_{i-j}, \quad \text{for } i = 0, 1, \dots, M-1.$$

To show that these two expressions are equivalent, observe the structure of the lower triangular Toeplitz matrices. In both matrices, the i-th row contains elements that are a reverse sequence of the first $i+1$ elements of the vector. Specifically, in $\mathbf{Ug}$, the i-th element is a dot product of the i-th row of $\mathbf{U}$ with $\mathbf{g}$, and in $\mathbf{Gu}$, it is the dot product of the i-th row of $\mathbf{G}$ with $\mathbf{u}$.

Since the order of multiplication in the dot product is commutative, we have:

$$\begin{aligned}(\mathbf{Ug})_i &= \sum_{j=0}^{i} u_j g_{i-j} = \sum_{k=0}^{i} u_k g_{i-k} = \sum_{j=0}^{i} g_j u_{i-j} \\ &= (\mathbf{Gu})_i.\end{aligned}$$

Since this equality holds for each $i = 0, 1, \dots, M-1$, we conclude that:

$$\mathbf{Ug} = \mathbf{Gu}.$$

This completes the proof. □

A.3. Proof Corollary A.3.1

Corollary A.3.1. *Consider the general continuous time optimal execution problem from Equation (4.7) including the risk aversion term w.r.t the OU alpha signal:*

$$\begin{aligned}\sup_{Q\in\mathcal{D}} \mathbb{E}\Bigg[&\int_0^T \alpha_t dQ_t - \frac{1}{2}\int_0^T\int_0^T \widetilde{\Theta}(t,s)G(|t-s|)dQ_s dQ_t - \psi_1\left(\int_0^T Z_t dQ_t\right)^2 \\ &-\psi_2\left(\varsigma\int_0^T\int_0^t e^{-\kappa(t-s)}dW_s' dQ_t\right)^2\Bigg].\end{aligned}$$

Consider the discrete time linear propagator model from Equation 3.4 for which we derived the expected cost of trading in Lemma 3.4.2 and the variance of the cost in Lemma 4.4. Then the discrete version of the optimal execution problem above can be formalized in matrix-vector notation as follows:

$$\begin{aligned}&\min_{\mathbf{q}\in\mathcal{D}} \mathbf{q}^T(\mathbf{\Phi} + \psi_1\mathbf{\Sigma} + \psi_2\mathbf{A})\mathbf{q} - \boldsymbol{\alpha}^T\mathbf{q} \\ &\textit{s.t.} \quad \mathbf{1}^T\mathbf{q} = Q_0^N,\end{aligned}$$

where the vector $\boldsymbol{\alpha}$ contains the Ornstein–Uhlenbeck alpha signal, i.e. $\mathbb{E}[\alpha_n] = \alpha_0 e^{-\kappa n} + \mu(1-e^{-\kappa n})$ and the symmetric matrix $\mathbf{A}$ represents the risk aversion w.r.t the alpha signal. The upper-triangular part of this matrix is equal to:

$$\mathbf{A}_{upper} = \begin{cases} A_{n,n} = \varsigma^2\sum_{i=1}^{n} e^{\kappa(i-1)} & \forall n = 1,\cdots T \\ A_{n,n+j} = \varsigma^2\sum_{i=1}^{n} e^{\kappa(2i-1)} & j \textit{ odd and}, \forall n = 1,\cdots T \textit{ and } j = 1,\cdots T-1 \\ A_{n,n+j} = \varsigma^2\sum_{i=1}^{n} e^{\kappa(2i)} & j \textit{ even and}, \forall n = 1,\cdots T \textit{ and } j = 2,\cdots T-1 \end{cases}$$

This is a quadratic optimization problem with optimal solution equal to:

$$\mathbf{q}^* = \frac{1}{2}(\mathbf{\Phi} + \psi_1\mathbf{\Sigma} + \psi_2\mathbf{A})^{-1}\left(\boldsymbol{\alpha} - \frac{\mathbf{1}}{\mathbf{1}^T(\mathbf{\Phi} + \psi_1\mathbf{\Sigma} + \psi_2\mathbf{A})^{-1}\mathbf{1}}\left(\mathbf{1}^T(\mathbf{\Phi} + \psi_1\mathbf{\Sigma} + \psi_2\mathbf{A})^{-1}\boldsymbol{\alpha} - 2Q_0^N\right)\right)$$

Proof. Consider the optimal execution problem with a deterministic admissible strategy $Q \in \mathcal{D}$:

$$\sup_{Q\in\mathcal{D}} \mathbb{E}\Bigg[\int_0^T \alpha_t dQ_t - \frac{1}{2}\int_0^T\int_0^T \widetilde{\Theta}(t,s)G(|t-s|)dQ_s dQ_t - \psi_1\left(\int_0^T Z_t dQ_t\right)^2 \\ -\psi_2\left(\varsigma\int_0^T\int_0^t e^{-\kappa(t-s)}dW_s' dQ_t\right)^2\Bigg].$$

Then we can rewrite the risk aversion term w.r.t the alpha signal using results from the proof of Lemma 4.4 in the following way:

$$\begin{aligned}
&\mathbb{E}\left[\left(\varsigma\int_0^T\int_0^t e^{-\kappa(t-s)}dW_s' dQ_t\right)^2\right] \\
&= \mathbb{E}\left[\left(\varsigma\int_0^T\left(\int_t^T e^{-\kappa(t-s)}dQ_s\right)dW_t'\right)^2\right] \\
&= \varsigma\int_0^T\left(\int_t^T e^{-\kappa(t-s)}dQ_s\right)^2 d[W']_t \\
&= \varsigma\int_0^T\left(\int_t^T e^{-\kappa(t-s)}dQ_s\right)^2 dt,
\end{aligned}$$

where in the second to last line we use Ito's isometry. Substituting this in the control problem we get:

$$\sup_{Q\in\mathcal{D}}\int_0^T \mathbb{E}[\alpha_t]dQ_t - \frac{1}{2}\int_0^T\int_0^T \widetilde{\Theta}(t,s)G(|t-s|)dQ_s dQ_t - \psi_1\int_0^T\left(\int_t^T \sigma_s dQ_s\right)^2 dt \\ -\psi_2\varsigma\int_0^T\left(\int_t^T e^{-\kappa(t-s)}dQ_s\right)^2 dt.$$

Using Theorem 4.5.1, we write this in matrix-vector notation as follows:

$$\begin{aligned}
&\min_{\mathbf{q}\in\mathcal{D}} \mathbf{q}^T(\mathbf{\Phi}+\psi_1\mathbf{\Sigma}+\psi_2\mathbf{A})\mathbf{q} - \boldsymbol{\alpha}^T\mathbf{q} \\
&\text{s.t.}\quad \mathbf{1}^T\mathbf{q} = Q_0,
\end{aligned}$$

where the vector $\boldsymbol{\alpha}$ contains the Ornstein–Uhlenbeck alpha signal, i.e. $\mathbb{E}[\alpha_n] = \alpha_0 e^{-\kappa n} + \mu(1-e^{-\kappa n})$ and the symmetric matrix $\mathbf{A}$ represents the risk aversion w.r.t the alpha signal.

The proof is a direct consequence of Theorem 4.3.1 and Theorem 4.5.1 since $\mathbf{\Phi}+\psi_1\mathbf{\Sigma}+\psi_2\mathbf{A}$ is a symmetric matrix. We only need to show that we can discretize the risk aversion w.r.t alpha signal as the symmetric matrix $\mathbf{A}$. Using the the notation in Definition 3.1.4, we discretize the risk-aversion w.r.t the alpha signal as follows:

$$\sum_{n=1}^T\left(\sum_{m=n}^T e^{-\kappa(n-m)}\Delta_m Q^N\right) = \sum_{n=1}^T\sum_{m=1}^T A_{n,m}\Delta_m Q^N \Delta_n Q^N.$$

If one writes out the left hand side for different T and inspects it structure we find:

$$\mathbf{A}_{upper} = \begin{cases} A_{n,n} = \sum_{i=1}^n e^{\kappa(i-1)} & \forall n = 1,\cdots T \\ A_{n,n+j} = \sum_{i=1}^n e^{\kappa(2i-1)} & j \text{ odd and },\forall n = 1,\cdots T \text{ and } j = 1,\cdots \\ A_{n,n+j} = \sum_{i=1}^n e^{\kappa(2i)} & j \text{ even and },\forall n = 1,\cdots T \text{ and } j = 2,\cdots \end{cases}$$

So for example we take $T = 5$, we find:

$$\mathbf{A} = \varsigma^2 \cdot \begin{bmatrix} 1 & e^{\kappa} & e^{2\kappa} & e^{3\kappa} & e^{4\kappa} \\ e^{\kappa} & e^{2\kappa}+1 & e^{3\kappa}+e^{\kappa} & e^{4\kappa}+e^{2\kappa} & e^{5\kappa}+e^{3\kappa} \\ e^{2\kappa} & e^{3\kappa}+e^{\kappa} & e^{4\kappa}+e^{2\kappa}+1 & e^{5\kappa}+e^{3\kappa}+e^{\kappa} & e^{6\kappa}+e^{4\kappa}+e^{2\kappa} \\ e^{3\kappa} & e^{4\kappa}+e^{2\kappa} & e^{5\kappa}+e^{3\kappa}+e^{\kappa} & e^{6\kappa}+e^{4\kappa}+e^{2\kappa}+1 & e^{7\kappa}+e^{5\kappa}+e^{3\kappa}+e^{\kappa} \\ e^{4\kappa} & e^{5\kappa}+e^{3\kappa} & e^{6\kappa}+e^{4\kappa}+e^{2\kappa} & e^{7\kappa}+e^{5\kappa}+e^{3\kappa}+e^{\kappa} & e^{8\kappa}+e^{6\kappa}+e^{4\kappa}+e^{2\kappa}+1 \end{bmatrix}$$

□

B

Appendix B: Additional numerical results

B.1. Residual analysis least squares problem

To estimate the non-projected kernel we use a least-squares approximation. To make sure that the least-square approximation gives the best unbiased solution to the problem, some assumptions need to be satisfied. However, we are dealing with a rather unique version of the least-squares approximation. In our problem, $\mathbf{U}_k$ is a lower-triangular Toeplitz matrix instead of a normal design matrix and $\mathbf{U}_k\mathbf{g}$ represents a convolution operation. Therefore, our assumptions deviate from the classical ones.

Below we list the assumptions and whether they apply to our model or not:

1. *Linearity:* In the linear propagator model we assume that the participation rate and the instantaneous impact have a linear relationship. In the Section B.2, we have looked into this.
2. *No endogeneity:* holds by construction of the model. In the linear propagator model we assume that only our participation in the market influences the impact.
3. *No perfect multicollinearity:* does not apply by construction of the model. The linear propagator model is based on the assumption that the impact of our trade is an 'accumulation' of the impacts of all our previous trades by means of a convolution. Or in different words, the predictors are correlated by construction. Therefore, this assumption does not apply in our case.
4. *Homoscedasticity:* the residuals have constant variance. This is something we check.
5. *Normality of the residual distribution:* The residuals are normally distributed. This is something we check.
6. *No autocorrelation in the residuals*: the residuals are not correlated with each other and the correlation does not change over time. This is something we check.
7. *Stationary of the residuals*: The distribution of the residuals does not change to much over time. This is something we check.

From the revision of the least-squares assumptions we conclude that we only need to check assumptions $4-7$. To check assumption 4 we make a scatter plot of the standardized square-root residuals versus the fitted values and use a spline regression (see Hastie (1986)) to check for a pattern. This is also called a scale-location plot of the standardized vs fitted values in the literature. We use this version of the residuals versus fitted plot because it is more robust for outliers. The scale-location plot is displayed in left hand side of Figure B.1. To check assumption 5, we make a Q-Q plot of the residuals. This plot is present in the right hand side of Figure B.1.

From the spline regression and the way how the observations are scattered around the grid we deduce that there is not a clear pattern between the standardized residuals and the fitted values. This means that the variance of the residuals can be assumed to be constant across the observations. Therefore, the homoscedasticity assumption holds. Furthermore, from the right hand side of Figure B.1 we find that

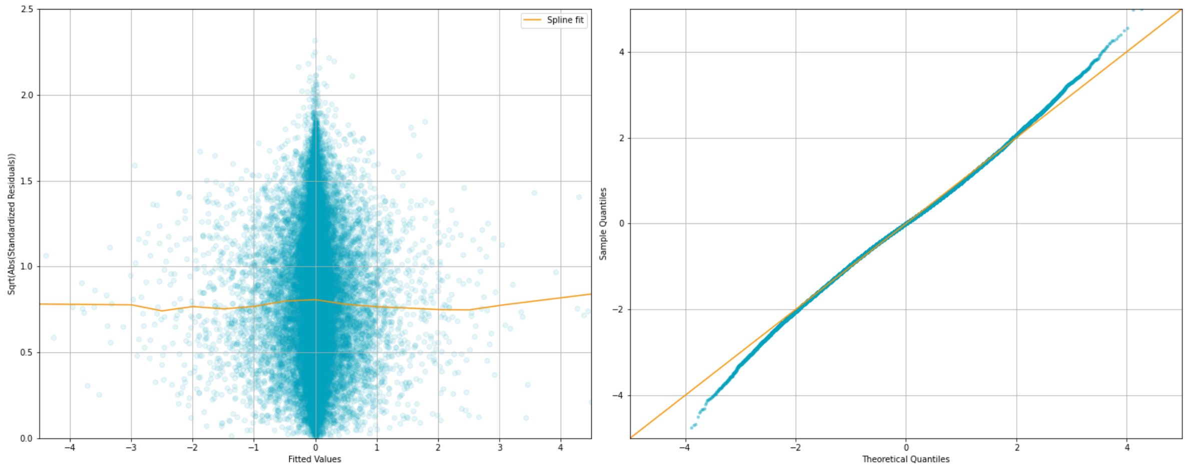

Figure B.1: (Left) scatterplot of the fitted values versus the standardized absolute square-root residuals. The orange line represents a spline regression trough the data. (Right) Quantile-Quantile (Q-Q) plot of the standardized residuals.

the residuals are by approximation normally distributed because the sample and theoretical quantiles align. However, we see a bit of a fat tail behaviour.

We continue by checking assumption 6, i.e. we investigate whether there is a significant autocorrelation between the residuals. To do this we use the Durbin–Watson (DW) statistic (see, Durbin and Watson (1950)). The DW statistic is a value between 0 and 4. If the value of the test statistic is close to 0 there is statistical evidence of a positive autocorrelation, if it is close to 2 there is no statistical evidence of autocorrelation and if it is close to 4 than there is statistical evidence for negative autocorrelation between the residuals.

Calculating the DW statistic for the total prediction error gives a value of 0.032. This indicates that the residuals are strongly positively lag 1 autocorrelated. This means that if the prediction is wrong in a certain direction, the next prediction will again be wrong in the same direction. This suggest that there has been a level shift which can not be explained by the model. For example, this could imply that there exist a non-zero short-term alpha signal.

To overcome this, we check for autocorrelation in the difference between the return per bin and the difference in predicted impact per bin, i.e. $e_n^N = \Delta_n S^N - \Delta_n I^N$. This approach makes use less vulnerable for a large level shift by a short-term alpha signal. In this new approach we calculate the DW statistic for the change in the residuals per metaorder and plot the empirical distribution. To compare the empirical distribution we formulate a base distribution in which we calculate the DW statistic of ϵ_n^N; the increments of the martingale, which are i.i.d standard normal distributed. In addition, we calculate the empirical distribution for different subsets of the dataset, where these subsets are all of the same length and represent different time windows. Both plots are present in Figure B.2. We see in Figure B.2 that the empirical distribution of the DW statistic is centered around 1.5. This means that the there is some evidence for positive autocorrelation in the change of the residuals but it is not critical. The same we have observed in Figure 2.3, in which we calculate the autocorrelation of the return. In this analysis we found a significant first lag, even if we remove our own impact. Therefore, we can conclude that the autocorrelation in the residuals is due to the data instead of the model. In the right hand side of the figure above we find that the empirical distribution of the of the DW-statistic does not change to much over time.

Lastly, we check of the stationary of the residuals. Since we are dealing with time series data it is important that the residuals are stationary over time. We check this by calculating the empirical distribution of the residuals for different time windows. All subsets are of equal length and based on time. The empirical distribution of the residuals for different time windows are displayed in Figure B.3. We see in Figure B.3 that the empirical distribution for different time widows are approximately the same. This is enough to conclude that the residuals are stationary.

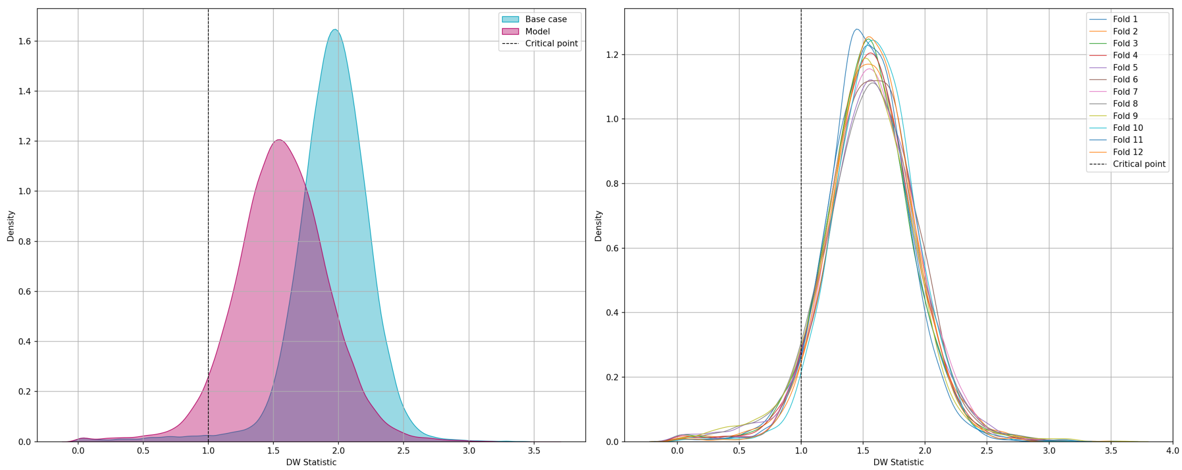

Figure B.2: (Left) the empirical distribution of the DW statistic of the rate of change in the residuals per metaorder in pink. The base distribution is the DW statistic of the increments of the martingale Z in blue. (Right) the empirical distribution of the DW statistic for 12 different time windows .

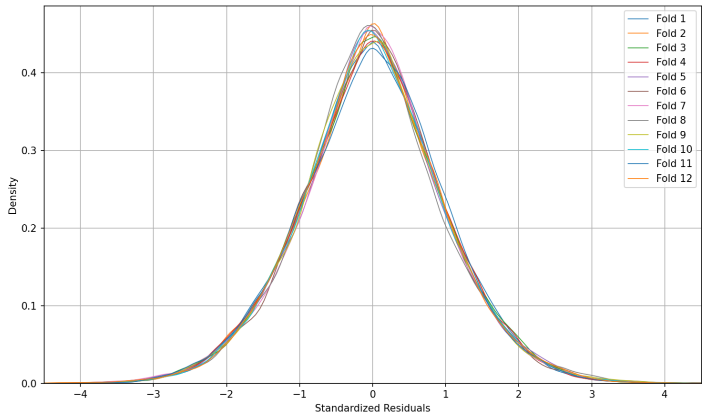

Figure B.3: Empirical distribution of the residuals for 12 different time windows.

B.2. Locally concave model performance

The main model use throughout this thesis is the linear propagator model on the trader's participation rate. In Chapter 3, we have seen that there are alternative modeling choices possible, including variations in the instantaneous impact function (linear vs. non-linear) and normalization methods (participation rate vs. normalized by ADV). The key reason for using the linear model is that a non-linear model permits price manipulation and lacks tractable solutions for the optimal execution problem.

The decision to use participation rate as the normalization method is twofold. Firstly, as demonstrated in Section 2.3.1, child orders normalized by the intraday volume curve exhibit a more linear relationship compared to those normalized by ADV. Secondly, normalizing by the intraday volume curve enables the model to account for the varying liquidity throughout the trading day, a critical factor influencing market impact.

In this section of the appendix, we demonstrate that this modeling choice also results in a better model fit in terms of R-squared.

Consider the discrete propagator model from Equation (3.4), in which we scale every child order by a concave function $f(\cdot)$:

$$S_n^N = S_0^N + \sum_{m=1}^{n} G(n-m) f(\Theta_m^N \Delta_m Q^N) + Z_n^N,$$

For demonstration purposes we consider an exponential kernel such that we obtain:

$$S_n^N = S_0^N + \sum_{m=1}^{n} \lambda e^{-\beta(n-m)} f(\Theta_m^N \Delta_m Q^N) + Z_n^N,$$

To investigate the concavity of the model for different normalizations of the traded quantity, we consider two functional forms for the instantaneous impact function $f(\cdot)$. First, we consider a power-law function normalized by ADV:

$$f_1(\Theta_n^N \Delta_n Q^N) = \sigma \text{ sign } (\Delta_n Q^N) \left(\frac{|\Delta_n Q^n|}{ADV} \right)^{c_1},$$

where c_1 is the concavity parameter. The second functional form is a power-law on the participation rate:

$$f_2(\Theta_n^N \Delta_n Q^N) = \sigma \text{ sign } (\Delta_n Q^N) \left(\frac{|\Delta_n Q^N|}{V_n^N} \right)^{c_2},$$

where V_n^N is the total volume traded in a time bucket.

By comparing these two functional forms, we aim to understand how different normalizations of the traded quantity affect the concavity and performance of the model. To analyze this, we employ a similar optimization approach to that outlined in Section 3.5, but we find the parameters of the exponential kernel directly. For each concavity parameter $c \in [0.3, 1]$, we determine the parameters using a variant of Equation 3.12:

$$\beta(c), \lambda(c) = \arg\min_{\beta,\lambda} \sum_{k=1}^{K} \left\| \mathbf{y}_k - \mathbf{U}_k \mathbf{g} \right\|^2,$$

where $\mathbf{g}$ is a vector representing the exponential decay kernel. Each entry U_n of the matrix $\mathbf{U}_k$ is scaled using either f_1 or f_2. To solve the optimization problem, we use the solver *"SLSQP"* from the *scipy.optimize.minimize* package in Python. Using the calibrated parameters, we evaluate the model performance on all orders using the R-squared.

Performing the analysis for different values of the concavity parameters results in Figure B.4. In this model we fit the concavity of the instantaneous impact function versus the model performance measured in R-squared for the two normalizations.

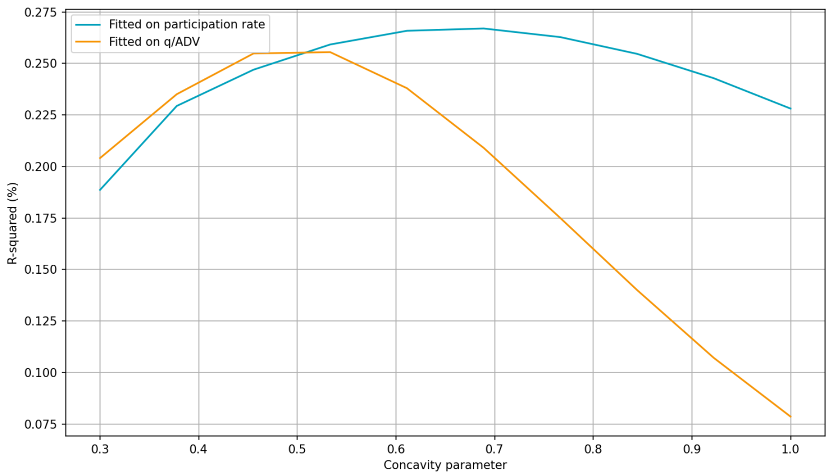

Figure B.4: Locally concave model performance for concavity parameter $c \in [0.3, 1]$ for the model fitted on participation rate and normalized by ADV

Let's start with the model fitted on the normalized child orders by ADV. Figure B.4 shows that the model performs best with a concavity parameter around $c_1 = 0.5$. This aligns with results from the literature (see Bouchaud et al. (2009)). However, as the model becomes less concave and approaches a linear form, its performance decreases significantly, by approximately 70%.

However, for the model fitted on the participation rate, the optimal concavity parameter is $c_2 = 0.65$. This indicates a less concave relationship compared to the ADV normalization, which is consistent with our observations in Figure 2.6. Additionally, the performance drop when transitioning towards a linear model is much smaller, only around 30%. Interestingly, the linear model based on the participation rate performs almost as well as the model normalized by ADV with a concavity of $c_1 = 0.5$.

In conclusion, the main reason for fitting a linear model is to ensure its viability. To compensate the performance loss associated with ADV normalization, fitting on the participation rate offers a strong alternative. The participation rate normalization maintains better performance even as the model becomes linear, making it a better choice for practical applications.

C

Appendix C: Additional figures

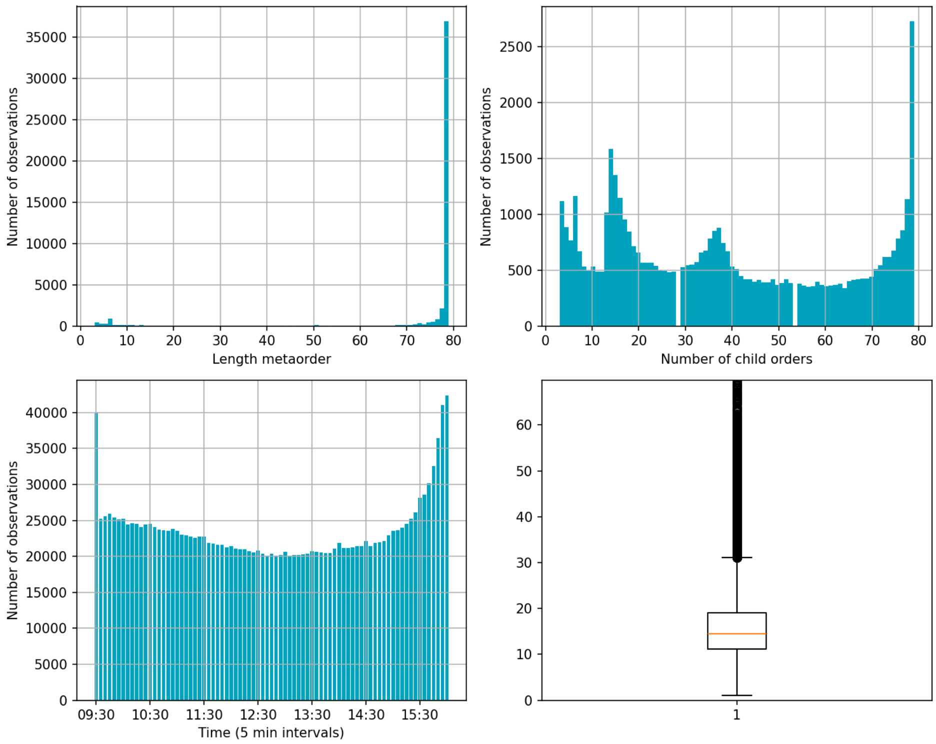

Figure C.1: After applying the filters in Table 2.1. (Upper left) histogram of the length of the metaorders. (Upper right) histogram of the number of child orders per metaorder. (Lower left) plot which displays the amount of trades in each intraday time bin. (Lower right) boxplot of the of the intraday volatility's.

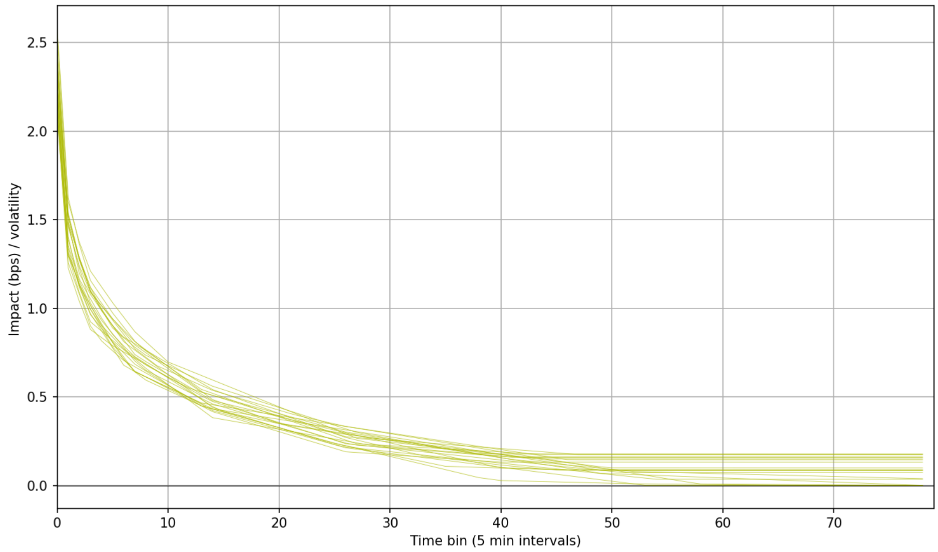

Figure C.2: Bootstrapped nonparametric decay kernel with 20 bootstraps.

Bibliography

Abi Jaber, E., & Neuman, E. (2022). Optimal liquidation with signals: The general propagator case. *SSRN*.

Alfonsi, A., Fruth, A., & Schied, A. (2008). Constrained portfolio liquidation in a limit order book model. *Banach Center Publ*, *83*, 9–25.

Alfonsi, A., Fruth, A., & Schied, A. (2010). Optimal execution strategies in limit order books with general shape functions. *Quantitative finance*, *10*(2), 143–157.

Alfonsi, A., & Schied, A. (2013). Capacitary measures for completely monotone kernels via singular control. *SIAM Journal on Control and Optimization*, *51*(2), 1758–1780.

Alfonsi, A., Schied, A., & Slynko, A. (2012). Order book resilience, price manipulation, and the positive portfolio problem. *SIAM Journal on Financial Mathematics*, *3*(1), 511–533.

Almgren, R., & Chriss, N. (1999). Value under liquidation. *Risk*, *12*(12), 61–63.

Almgren, R., & Chriss, N. (2001). Optimal execution of portfolio transactions. *Journal of Risk*, *3*, 5–40.

Almgren, R., Thum, C., Hauptmann, E., & Li, H. (2005). Direct estimation of equity market impact. *Risk*, *18*(7), 58–62.

Bacry, E., Iuga, A., Lasnier, M., & Lehalle, C.-A. (2015). Market impacts and the life cycle of investors orders. *Market Microstructure and Liquidity*, *1*(02), 1550009.

Bertsekas, D. (2012). *Dynamic programming and optimal control: Volume i* (Vol. 4). Athena scientific.

Bertsimas, D., & Lo, A. W. (1998). Optimal control of execution costs. *Journal of financial markets*, *1*(1), 1–50.

Bochner, S. (1932). *Vorlesungen uber fouriersche integrale*. Akademische Verlagsgesellschaft.

Bordigoni, G., Figalli, A., Ledford, A., & Ustinov, P. (2022). Strategic execution trajectories. *Applied Mathematical Finance*, *29*(4), 288–330.

Bouchaud, J.-P., Bonart, J., Donier, J., & Gould, M. (2018). *Trades, quotes and prices: Financial markets under the microscope*. Cambridge University Press.

Bouchaud, J.-P., Farmer, J. D., & Lillo, F. (2009). How markets slowly digest changes in supply and demand. In *Handbook of financial markets: Dynamics and evolution* (pp. 57–160). Elsevier.

Bouchaud, J.-P., Gefen, Y., Potters, M., & Wyart, M. (2003). Fluctuations and response in financial markets: The subtle nature ofrandom'price changes. *Quantitative finance*, *4*(2), 176.

Bouchaud, J.-P., Kockelkoren, J., & Potters, M. (2006). Random walks, liquidity molasses and critical response in financial markets. *Quantitative finance*, *6*(02), 115–123.

Brokmann, X., Serie, E., Kockelkoren, J., & Bouchaud, J.-P. (2015). Slow decay of impact in equity markets. *Market Microstructure and Liquidity*, *1*(02), 1550007.

Bucci, F., Benzaquen, M., Lillo, F., & Bouchaud, J.-P. (2018). Slow decay of impact in equity markets: Insights from the ancerno database. *Market Microstructure and Liquidity*, *4*(03–04), 1950006.

Busseti, E., & Lillo, F. (2012). Calibration of optimal execution of financial transactions in the presence of transient market impact. *Journal of Statistical Mechanics: Theory and Experiment*, *2012*(09), 09010.

Cartea, Á., Jaimungal, S., & Penalva, J. (2015). *Algorithmic and high-frequency trading*. Cambridge University Press.

Cont, R., Kukanov, A., & Stoikov, S. (2014). The price impact of order book events. *Journal of financial econometrics*, *12*(1), 47–88.

Curato, G., Gatheral, J., & Lillo, F. (2017). Optimal execution with non-linear transient market impact. *Quantitative Finance*, *17*(1), 41–54.

Durbin, J., & Watson, G. S. (1950). Testing for serial correlation in least squares regression: I. *Biometrika*, *37*(3–4), 409–428.

Efron, B., & Tibshirani, R. J. (1994). *An introduction to the bootstrap*. CRC press.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The journal of Finance*, *25*(2), 383–417.

Fruth, A., Schöneborn, T., & Urusov, M. (2014). Optimal trade execution and price manipulation in order books with time-varying liquidity. *Mathematical Finance*, *24*(4), 651–695.

Gatheral, J. (2010). No-dynamic-arbitrage and market impact. *Quantitative finance*, *10*(7), 749–759.

Gatheral, J., Schied, A., & Slynko, A. (2012). Transient linear price impact and fredholm integral equations. *Mathematical Finance*, *22*(3), 445–474.

Harvey, C. R., Ledford, A., Sciulli, E., Ustinov, P., & Zohren, S. (2022). Quantifying long-term market impact. *Journal of Portfolio Management*, *48*(3).

Hastie, T. J. (1986). Generalized additive models. *Statistical Science*, *1*(3), 297–318.

Hey, N., Bouchaud, J.-P., Mastromatteo, I., Muhle-Karbe, J., & Webster, K. (2023). The cost of misspecifying price impact. *arXiv preprint*.

Hey, N., Mastromatteo, I., Muhle-Karbe, J., & Webster, K. (2023). Trading with concave price impact and impact decay-theory and evidence. *SSRN*.

Huberman, G., & Stanzl, W. (2004). Price manipulation and quasi-arbitrage. *Econometrica*, *72*(4), 1247–1275.

Hull, J. C. (2016). *Options, futures, and other derivatives*. Pearson Education.

Isichenko, M. (2021). *Quantitative portfolio management: The art and science of statistical arbitrage*. John Wiley & Sons.

Lehalle, C.-A., & Neuman, E. (2019). Incorporating signals into optimal trading. *Finance and Stochastics*, *23*, 275–311.

Merton, R. C. (1975). Optimum consumption and portfolio rules in a continuous-time model. In *Stochastic optimization models in finance* (pp. 621–661). Elsevier.

MSCI Inc. (2024). MSCI USA Index (USD) - Index Factsheet. https://www.msci.com/documents/10199/67a768a1-71d0-4bd0-8d7e-f7b53e8d0d9f

Neuman, E., Stockinger, W., & Zhang, Y. (2023). An offline learning approach to propagator models. *arXiv preprint*.

Obizhaeva, A. A., & Wang, J. (2013). Optimal trading strategy and supply/demand dynamics. *Journal of Financial markets*, *16*(1), 1–32.

Perold, A. F. (1988). The implementation shortfall: Paper versus reality. *Journal of Portfolio Management*, *14*(3), 4.

Protter, P. (2005). *Stochastic integration and differential equations*. Springer.

Shreve, S. E., et al. (2004). *Stochastic calculus for finance ii: Continuous-time models* (Vol. 11). Springer.

Stock, J. H., & Watson, M. W. (2020). *Introduction to econometrics*. Pearson.

Tóth, B., Lemperiere, Y., Deremble, C., De Lataillade, J., Kockelkoren, J., & Bouchaud, J.-P. (2011). Anomalous price impact and the critical nature of liquidity in financial markets. *Physical Review X*, *1*(2), 021006.

Uhlenbeck, G. E., & Ornstein, L. S. (1930). On the theory of the brownian motion. *Physical review*, *36*(5), 823.

Wagner, W. H., & Edwards, M. (1993). Best execution. *Financial Analysts Journal*, *49*(1), 65–71.

Webster, K. T. (2023). *Handbook of price impact modeling*. CRC Press.

Yong, J., & Zhou, X. Y. (2012). *Stochastic controls: Hamiltonian systems and hjb equations* (Vol. 43). Springer Science & Business Media.