

Price Impact

Correlation Between Buy-/Sell-Pressure in the Stock Market and Subsequent Price Changes

Fredrik Lidefelt & Eric Saidac

Spring 2017

Abstract

In this thesis, the correlation between aggregated buy-/sell-pressure and price change is analyzed, something often referred to as price impact. The purpose is not only to find a model that can explain how prices change given a certain traded volume, but also to see how much of the variance in price changes results from the actual buy- or sell-pressure. Analysis is conducted on monthly trading data from stocks comprising OMXS30 in 2016. An odd and increasing function of net bought relative to average daily turnover is suggested as an appropriate way to model price impact. In line with previous research the function is also concave in order size. An exponentially decaying part dependent on previous trades is then added to the model in order to keep price changes uncorrelated. Further it is analyzed how the dependency on average daily turnover affects price impact, and how the rise of smaller market places might change the sensitivity to price impact at larger ones. Ultimately, it is concluded that the hypothesis stating that a fragmented market increases sensitivity to price impact cannot be rejected.

Preface

This is the Master's Thesis of Engineering Mathematics students Fredrik Lidelfelt and Eric Saidac, both pursuing their Master of Science in Engineering degree with the specialization Financial Modelling at Lund University, Faculty of Engineering.

Nasdaq Stockholm AB's Economic & Statistical Research department has conducted some small scale work into price impact a few years ago, but wanted a more thorough and impartial report of the phenomenon. The study fell into our hands through the department's Associate Vice President Björn Hertzberg; our external supervisor throughout this process. Björn Hertzberg and his team has provided large quantities of data as well as guidance through experienced insights, while remaining careful not to influence our work in any direction.

The process officially started January 30th 2017 in Stockholm, Sweden, but some preparatory work was conducted during the previous fall. The first two months were spent at the National Library of Sweden while awaiting security clearance for Nasdaq's Stockholm offices, after which we were provided with computers, data access and work space in their main offices in Frihamnen, Stockholm. A presentation of our work is scheduled June 1st 2017.

We would like to express our sincerest gratitude to Björn Hertzberg for his tireless aid through discussions that has proven indispensable for the completion of this work. Furthermore, we would like to thank Hossein Asgharian, Professor at Department of Economics - Lund University, for his assistance as internal supervisor from the university. Professor Asgharian has provided many important notes to improve the comprehensibility of the thesis as well as inspiration when new directions needed to be explored.

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 6 |
| 2 | Theory: Stock Market | 7 |
| 2.1 | Modelling the Stock Market | 7 |
| 2.2 | Competition Among Markets | 8 |
| 3 | Theory: Previous Research | 9 |
| 3.1 | Uncorrelated Stock Returns | 9 |
| 3.1.1 | How Can Order Signs Be Correlated? | 10 |
| 3.1.2 | Two Models | 10 |
| 3.2 | No Dynamic Arbitrage | 11 |
| 3.2.1 | Cost of Trading | 12 |
| 3.2.2 | Linear Permanent Impact | 12 |
| 3.2.3 | A Round Trip Strategy | 13 |
| 3.2.4 | Power-Law Decay | 14 |
| 3.2.5 | Other Decay Functions | 15 |
| 3.2.6 | Applicability of the Theory | 15 |
| 3.3 | Summary | 16 |
| 4 | Method & Data | 17 |
| 4.1 | Data | 17 |
| 4.1.1 | Broad Explanation of Data | 17 |
| 4.1.2 | Data Specifics | 18 |
| 4.1.3 | Distribution of Data | 19 |
| 4.2 | Method: Instantaneous Model | 19 |
| 4.3 | Method: Added Time Dependency | 21 |
| 5 | Results | 24 |
| 5.1 | Instantaneous Part of the Model | 24 |
| 5.2 | Temporal Part of the Model | 26 |
| 5.3 | Decay Function | 30 |
| 5.3.1 | Shape of the Decay Function | 31 |
| 5.3.2 | Performance on Individual Stocks | 33 |
| 5.4 | Summary | 35 |
| 6 | Discussion | 37 |
| 6.1 | Model Evaluation | 37 |
| 6.1.1 | Characteristics of Stock Returns | 37 |
| 6.1.2 | How Much is Explained? | 39 |
| 6.1.3 | Real vs. Predicted Prices | 39 |
| 6.1.4 | Is the Model Accurate? | 41 |
| 6.1.5 | This Model vs. Previous Models | 42 |
| 6.2 | The Effects of Price Impact | 44 |
| 6.2.1 | First Observations | 45 |
| 6.2.2 | Low Turnover Implies High Impact | 46 |
| 6.2.3 | Fragmentation of the Market | 48 |
| 7 | Conclusions | 52 |

| | |
|--|-----------|
| References | 54 |
| Appendix A Variables | 56 |
| Appendix B Tables | 57 |
| B.1 List of Stocks in OMXS30 | 57 |
| Appendix C Plots | 58 |
| C.1 Market Share of Different Venues | 58 |
| C.2 Visualization of all Stocks | 59 |
| C.3 Visualization of Quantile Means for all Stocks | 60 |
| C.4 Estimated β :s Using One and Six Months of Data for all Stocks . | 61 |
| C.5 Fitted $G(\cdot)$ over $\hat{\beta}$:s for all Stocks | 63 |
| C.6 Plots of Price Impact for Chi-X, Nasdaq and Turquoise | 64 |

List of Figures

| | | |
|------|--|----|
| 3.1 | Autocorrelation of order signs | 10 |
| 4.1 | First visualization | 19 |
| 4.2 | Histogram and QQ-plot to determine distribution | 20 |
| 4.3 | Visualization of means | 21 |
| 5.1 | Parameter estimates | 26 |
| 5.2 | Decay coefficients | 27 |
| 5.3 | Mean value of decay coefficients | 28 |
| 5.4 | Collectively estimated decay coefficients | 29 |
| 5.5 | Decay coefficients estimated for one and six months of data | 30 |
| 5.6 | Estimated $\hat{\beta}_i$:s for six stocks | 31 |
| 5.7 | $\hat{\beta}_i$:s estimated on all 29 stocks. | 32 |
| 5.8 | Fitted decay functions | 32 |
| 5.9 | Estimated values of the a- and b-parameters for all stocks | 34 |
| 5.10 | Exponential function fitted on decay coefficients | 34 |
| 5.11 | Exponential decay function with fixed a and b | 35 |
| 6.1 | Autocorrelation of observed vs simulated price changes | 38 |
| 6.2 | Autocorrelation of simulated price changes without decay | 38 |
| 6.3 | Observed vs simulated price changes | 39 |
| 6.4 | Real returns vs predicted price impact-returns in January 2016. | 40 |
| 6.5 | Real returns vs predicted price impact-returns in April 2016. | 41 |
| 6.6 | Data cloud of real returns vs Data cloud of η_t | 42 |
| 6.7 | The effect of turnover visualized | 44 |
| 6.8 | The effects of turnover for SEB A | 45 |
| 6.9 | The effects of turnover for Sandvik | 46 |
| 6.10 | Turnover's effect on incilnation of the price impact curve | 47 |
| 6.11 | Comparisons across venues | 50 |
| C.1 | Market share of the three largest venues | 58 |
| C.2 | First visualization: all stocks | 59 |
| C.3 | Visualization of means: all stocks | 60 |
| C.4 | Estimated β :s for all 29 stocks using only one month of data. | 61 |
| C.5 | Estimated β :s for all 29 stocks using six months of data | 62 |
| C.6 | Fitted $G(\cdot)$ over $\hat{\beta}$:s for all Stocks. | 63 |
| C.7 | Price impact curve for three venues: first quarter | 64 |
| C.8 | Price impact curve for three venues: second quarter | 65 |

1 Introduction

Considerable research has been made to explain the correlation between an incoming order and the price change that follows. This effect on stock prices is commonly referred to as Market Impact, which considers only one trade at a time. Empirically one can show that a buy order will result in a price increase whereas a decrease will follow a sell order. To accurately model this phenomenon and thus being able to forecast the cost of trading is of great importance for modern trading algorithms. In this thesis however we will consider what has been called Price Impact, where all trades during a certain time period on the market are considered as a simultaneous aggregate of order volume. The purpose is to understand how stock movements are affected by buy-/sell-pressure and also if and why this phenomenon behaves differently for different stocks, at different points in time and across various market places.

We will in this thesis first give a brief introduction to the properties of the stock market, and then present two theories that has been developed on the topic of market impact. The first one is used in the literature of Gerig [7] and is based on the discovery that there is a correlation in order signs (i.e. whether we see a buy- or a sell-order in the order book), which leads to certain constraints on the price impact function. The second theory starts from a no-arbitrage principle and is discussed in the paper of Gatheral [6]. The claim is that the way we model price impact should not give rise to any arbitrage opportunities. We then, based on the theories presented, move on to observing our data and fitting a price impact model which, in accordance with the theories of Gerig, turned out to consist of two parts; an instantaneous impact function and a decaying part depending on previous trades. What then follows is a discussion section which is divided into two parts. The first discusses the specifics of our model, and how it contradicts some of the theories presented by Gatheral. This part also discusses the implications of our model, and how it is used to quantify how much of the variance in price movements is a consequence of buy-/sell-pressure. The second part discusses the broader properties of price impact and why it is observed that different stocks and different market places are more or less sensitive to price impact.

2 Theory: Stock Market

In the stock market of today a buy order is a request to purchase a stock at a certain price decided by the buyer, and similarly, a sell order is a request to sell. These orders are collectively called an order book, and the number of price levels available in said order book constitutes the order book depth. A simple way of viewing the market is that a transaction will occur if there is a sell order that matches a buy order. We specify the nature of the order by its sign, a sell order has negative sign and a buy order has a positive sign, i.e. negative/positive volume of the trade. Furthermore, there are two basic categories of orders, *limit* and *market* orders. Limit orders specify a requested quantity and a price limit that the trader is not willing to surpass, either by selling their asset a price they consider too low or by buying at a price that is too high. A market order also specifies a requested quantity, but is not limited in its price. Instead it is automatically matched with the best sell/buy offer available. A limit buy order is further specified by its bid-price and bid-size, collectively called a *bid*. Analogously a limit sell order is specified by ask-price and ask-size, the sell order itself is called an *ask*. If there are no asks to match an incoming bid, i.e. if the bid-price is too low or the bid-size supersedes the available ask-size at said price, then the (remaining) bid will be saved in the order book and wait for a matching ask. The highest bid-price currently available in the order book is called the best bid and the lowest ask-price the best ask, these two prices are of great importance for stock market modelling as will be noticed throughout this thesis. The price in between the best ask and bid is often spoken of as the current stock price, more technically correct it is the *mid-point price*. The difference of the two prices is called the *spread*, which can also be used as a measure of implicit volatility of the stock. Another way to refer to limit orders and their prices is by the term *quotes*, which is used if the limit order exists at best price. When these properties are considered one realizes that continuous pricing of assets (i.e. of bid-/ask-prices) is not an option, since matches between ask and bid would become improbable, therefore the market regulates the traders' offers into discrete pricing increments where the minimum increment is called *tick size*. [7]

2.1 Modelling the Stock Market

Traditionally stock prices have been modelled as a stochastic process S_t that follows a geometric Brownian motion. This has been done since it was determined that prices move in percentage increments that were assumed to be drawn independently and identically from a Gaussian distribution [7]. This process satisfies the following stochastic differential equation (SDE) used famously by Black and Scholes, with the following dynamics [13],

$$dS_t = \mu S_t dt + \sigma S_t dW_t, \quad S_0 = s, \quad (1)$$

where s , μ ("percentage drift") and σ ("percentage volatility") are constants, and W_t is a Brownian motion or Wiener process.

The fact that there is a model for stock prices can seem contradictory since the prices themselves are varying with expectations and trading. This is however cared for since stock prices approximate a martingale, i.e. the best approximation of tomorrow's price conditioned on historical prices is the price today [12].

If it worked differently a prediction would be exploited until the predictability was nonexistent, which is the meaning of acting on an arbitrage opportunity. The martingale property of stock prices is very much related to the efficient market hypothesis (EMH), or more precisely the weak form of it. EMH was developed by Eugene Fama in three types of efficiency [5]:

1. The weak form of the EMH claims that all historic information publically available is reflected on traded assets.
2. The semi-strong form of the EMH claims in addition to (1) that prices instantly change to reflect new public information.
3. The strong form of the EMH claims in addition to (1) and (2) that prices also reflect hidden insider information.

2.2 Competition Among Markets

By November 1st 2007 a European Union law known as the Markets in Financial Instruments Directive (MiFID) was implemented, increasing competition in investment services [11]. Under MiFID there are two types of market places, regulated markets and Multilateral Trading Facilities (MTF). The regulations for an MTF are less strict than those of a regulated market. Furthermore, MiFID rendered the term *börs* (en: *exchange*, fr: *bourse*) obsolete, a term that prior to MiFID marked a trading venue as a primary exchange rather than a list [9]. The new dynamics in the trading business gave rise to rivalry spanning across the European Union. Shortly after MiFID was implemented, trading began at London based Chi-X (the company BATS acquired Chi-X in 2013, eventually earning the venue a Recognized Investment Exchange (RIE) status), and the effects were immediate. Since Chi-X holds RIE status, trading through their service does not affect the primary markets' order books, and they are indeed a stock exchange in their own right [14]. There are other alternative venues with market share less than Chi-X as well, the biggest of which is Turquoise, another London based MTF calling themselves a "pan-European trading platform" that was founded in 2008 [2].

By February 2009, approximately 10% of total traded volume of OMXS30 (an index describing the 30 most traded shares on the Stockholm Stock Exchange) was traded on Chi-X; by February 2016 that share had risen to 26% [16]. The trend is clearly an increasing one, but it seems to have stagnated somewhat the last years, and it thus appears as if the cannibalization of Nasdaq's market share by Chi-X has been terminally saturated (see Appendix C.1 for graphics).

3 Theory: Previous Research

When we are trying to explain the correlation between an incoming order and the subsequent price change there are a lot of aspects that should be considered. Firstly, there are questions regarding the actual input parameters to the price impact function. Is the impact only dependent on the traded volume, or does it also depend on characteristics that vary from stock to stock, such as volatility and turnover? Is the impact dependent on who is buying/selling? There is also a time aspect to the modelling. Does the impact at time t only depend on what happens at that specific time point, or does previous events affect the price impact? Finally, there is of course the more statistical/mathematical challenge to answer what is the shape and form of the impact function. Is it linear or nonlinear? If current trades impact future trades, for how long?

In this section we will seek answers to these questions. To do this we will try and see what certain characteristics of the impact function, purely mathematically, implies for the process of stock returns. To determine if a price impact function is reasonable we will present two different approaches which has been used for this purpose in many articles before this. The first one is that our way to model price impact cannot induce stock movements such that it leads to correlated stock returns. This approach is discussed by Gerig [7], and with this constraint alone one is able to exclude certain shapes of the impact function. The second approach is the one first used by Huberman and Stanzl [8] and later in numerous articles by Gatheral, which starts from a no-arbitrage principle. This principle merely states that a certain impact function cannot give rise to any arbitrage opportunities. That is, it should not be possible to have a strategy where we buy and sell the same amount of stocks but still, as a consequence of how we affect the price described by the price impact function, realize a profit.

3.1 Uncorrelated Stock Returns

As mentioned above we wish to find an impact function such that it does not give rise to correlated stock returns. To begin the process of finding a suitable shape for the price impact function we simply model the price change of a stock as

$$\Delta p = \epsilon_t f(v_t) + \eta_t, \quad (2)$$

where Δp is the price change of the stock at time/index t . The index variable t is commonly updated by one increment whenever a transaction occurs (i.e. when two market participants agree on a price for a stock). The function $f(\cdot)$ is the price impact function, which is a function of the transaction size v_t , and ϵ_t is the transaction sign (negative for sell and positive for buy). Finally, η_t is an uncorrelated noise term which could model specific changes in the return of a stock that are not a function of the traded volume. In the model above we assume that the price change caused by the volume traded at time t is independent of what has happened before this time. So one can model stock price changes between time 1 and T as the sum

$$\Delta p = p_0 + \sum_{t=1}^T \epsilon_t f(v_t) + \eta_t$$

As seen in the expression above price changes are only uncorrelated if neither trade sign nor volume exhibits autocorrelation. However, we have seen from empirical research that the sign of transactions, ϵ_t , are highly autocorrelated. A buy transaction tends to follow a buy, while a sell tends to follow a sell. This long-term memory in order sign seems to be universal and has been verified for stocks traded on the London Stock Exchange, Paris Bourse and the New York Stock Exchange [7]. This autocorrelation of order signs is seen in the example illustrated by Figure 3.1 on Nasdaq Stockholm.

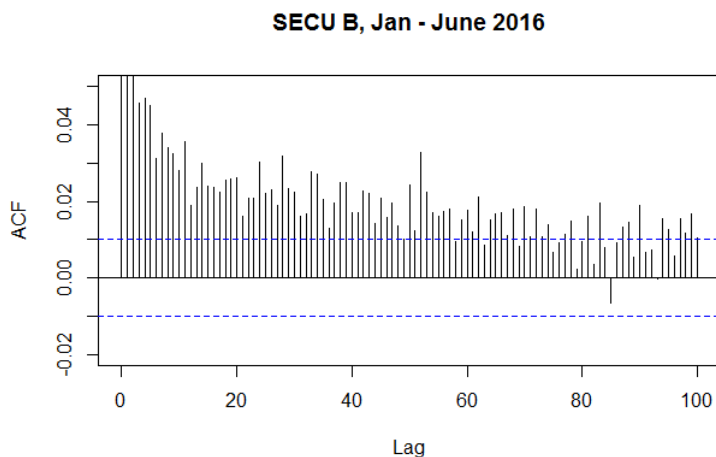


Figure 3.1: Autocorrelation of order signs, as observed on Securitas (Nasdaq Stockholm)

3.1.1 How Can Order Signs Be Correlated?

The problem of dealing with the autocorrelation of order flow signs has been the topic of discussion in many papers. If a buy or sell order tend to push the price up or down, and we can predict whether a buy or sell should appear next in the order book, wouldn't that mean we can predict the stock price? Why can't we? Gerig calls this the *efficiency puzzle*, as he refers to the uncorrelated and unpredictable nature of assets, also known as *market efficiency* [7]. A possible explanation to the autocorrelation of order flow could be that, even highly liquid markets such as the ones we analyze, only offer small volumes at the best quotes. This means that big trades will be fragmented and it could take up to several hours or days before they are completed, creating a long-term memory in the sign of the order flow. If this is true, and if we assume that stock prices follow a random walk, then, as Bouchaud describes it, a surprise in the order flow must impact prices more. That is, a buy following a buy should have less impact than a buy following a sell, otherwise trends would appear [4].

3.1.2 Two Models

Gerig presents two different explanations to this *efficiency puzzle*. The first one is the one introduced by Lilo and Farmer (LF) [10], who argues that prices stay

uncorrelated with the help of a liquidity term, λ_t , which is dependent on the order sign.

$$\Delta p = \frac{\epsilon_t f(v_t)}{\lambda_t} + \eta_t,$$

They mean that if buying or selling is probable, λ_t increases or decreases keeping the price movements unpredictable. One can thus say that this model is dependent on the current state of the market.

The second model Gerig presents was introduced by Bouchaud, Gefen, Poters, and Wyart (BGPW) [3] who instead argues that the efficiency puzzle is explained, not by a state dependent model, but by introducing a decay function to act on the price impact. The total price change at time t is,

$$\Delta p = \epsilon_t f(v_t) - \sum_{k>0} \frac{G(k+1) + G(k)}{G(1)} \epsilon_{t-k} f(v_{t-k}) + \eta_t$$

where the price change at each time is dependent on transactions in the past. The function $G(t)$ is called the decay function and is according to BGPW a power-law with exponent < 1 . It is tuned in a way such that the autocorrelation of order signs is cancelled, leading to uncorrelated price changes.

Gerig shows in his report that the LF model is merely a generalized version of the BGPW model. And that the two models are equivalent under certain assumptions.

3.2 No Dynamic Arbitrage

So far we have argued that there should be some historic dependence to the price impact model based on the uncorrelated nature of stock returns. In a lot of papers the stock return is thus instead modelled as a modified version of (1)

$$S_t = S_0 + \int_0^t f(\dot{x}_s) G(t-s) ds + \int_0^t \sigma dZ_s \quad (3)$$

This is used in Gatheral's article [6] where \dot{x}_s is the rate of trading at time $s < t$ (positive for buying, negative for selling), $f(\dot{x}_s)$ is the impact of trading at time s , the *impact function*, and $G(t-s)$ is a decay factor sometimes referred to as the *decay kernel*. We can see from the equation above that S_t follows a random walk with drift term that depends on the accumulated impact of previous trades. The equation above can be seen as the limit of the discrete time process

$$S_t = S_0 + \sum_{i<t} f(\delta x_i) G(t-i) + \text{noise},$$

where δx_i is the quantity traded at some small interval δt . This discrete time process, and therefore time continuous process in (3), corresponds with the BGPW-picture [6].

The reason why we choose to model the stock price process as in (3) is to in a more stringent way understand which combinations of impact function and decay kernel are realistic. To do this we will use a well-known approach when it comes to financial modelling, namely the no-arbitrage principle. The no-arbitrage principle merely states that the expected cost of any trading strategy

which result in a zero net bought should be non-negative. This can be motivated as it is fair to assume that if there would exist a trading strategy that accumulated trades at a certain rate and then liquidated at a certain rate and resulted in an arbitrage profit, this arbitrage opportunity would be exploited away.

The no-arbitrage theory was first used for modelling price impact in the literature of Huberman and Stanzl [8]. Starting from a no-arbitrage principle they showed that price impact must be symmetric between buys and sells, which we will assume throughout this thesis. When Huberman and Stanzl first applied this principle to price impact they called it *no dynamic-arbitrage*, as it is in fact the *expected value* of trading costs that is being looked at.

The principle of no dynamic arbitrage allows for, in a more analytical way, computation of how the impact function and the decay kernel are related. However, before this can be done, a formal definition of the cost of trading is needed.

3.2.1 Cost of Trading

If we, in analogy with Gatheral [6], denote the number of shares traded at time t as x_t , then the expected cost $C[\Pi]$ for a specific strategy $\Pi = \{x_t\}_{0 \leq t \leq T}$ is given by

$$\begin{aligned} C[\Pi] &= \mathbb{E} \left[\int_0^T \dot{x}_t (S_t - S_0) dt \right] \\ &= \int_0^T \dot{x}_t \int_0^t f(\dot{x}_s) G(t-s) ds dt \end{aligned}$$

where $\dot{x}_t dt$ is the shares traded at time t at the expected price given by (3). If, under the constraint $\int_0^t \dot{x}_t dt = 0$, this cost was ever to become negative, it means that a certain amount of stocks can be bought and then sold at a negative cost (i.e. a profit), which means that there exists an arbitrage opportunity. As mentioned earlier, this is in fact not a "real" arbitrage opportunity as it is not certain what other traders will do during this time. Hence it is called *dynamic* arbitrage. By now the theory of no-dynamic arbitrage can be stated as a proposition. For future references a trading strategy that satisfies the condition $\int_0^t \dot{x}_t dt = 0$ will be called a round trip trade, i.e. when the strategy involves selling and buying the same number of shares.

Proposition 1 (*No Dynamic-Arbitrage*). For any strategy $\{x_t\}_{0 \leq t \leq T}$ such that $\int_0^t \dot{x}_t dt = 0$ we must have that

$$\int_0^T \dot{x}_t \int_0^t f(\dot{x}_s) G(t-s) ds dt \geq 0. \quad (4)$$

The equation above restricts $f(\cdot)$ and $G(\cdot)$ and we will in what follows, based on the methodology Gatheral [6] uses, learn more about these restrictions.

3.2.2 Linear Permanent Impact

As a first example let us assume that $f(\dot{x}) = k\dot{x}$ for some $k > 0$ and that the impact is permanent, i.e. $G(t-s) = 1$. For any arbitrary trading strategy

$\Pi = \{x_t\}_{0 \leq t \leq T}$ where we go from x_0 shares at time $t = 0$ to x_T shares at time $t = T$ we get

$$\begin{aligned} C(\Pi) &= \int_0^T \dot{x}_t \int_0^t f(\dot{x}_s) ds dt = k \int_0^T \dot{x}_t \int_0^t \dot{x}_s ds dt = k \int_0^T \dot{x}_t [x_s]_0^t dt \\ &= k \int_0^T \frac{dx_t}{dt} (x_t - x_0) dt = k \int_{x_0}^{x_T} (x_t - x_0) dx_t = \frac{k}{2} (x_T - x_0)^2. \end{aligned}$$

In the calculations above we have used that $\dot{x}_t = \frac{dx_t}{dt}$. Looking at Proposition 1 we can see that the condition for no dynamic-arbitrage is satisfied independently of the trading strategy Π since the cost only depends on the initial quantity, x_0 , and the final quantity, x_T , which are equal for a round trip trade. Thus, we can conclude that a linear permanent impact function does not violate the condition of no-dynamic arbitrage.

3.2.3 A Round Trip Strategy

Consider a strategy where we buy trades at the constant rate r_1 and then sell to the constant rate r_2 ($r_i > 0$, $i = 1, 2$), i.e.

$$\dot{x}_t = \begin{cases} r_1, & \text{if } 0 < t \leq \theta T, \\ -r_2, & \text{if } \theta T < t \leq T. \end{cases} \quad (5)$$

We choose θ so that we buy and sell the same amount of stocks, that is $r_1 \theta T - r_2 (T - \theta T) = 0$, which means

$$\theta = \frac{r_2}{r_1 + r_2}.$$

If we assume that $f(\cdot)$ is odd (i.e. $f(r) = -f(-r)$) we get

$$\begin{aligned} C(\Pi) &= \int_0^T \dot{x}_t \int_0^t f(\dot{x}_s) G(t-s) ds dt \\ &= \int_0^{\theta T} r_1 \int_0^t f(r_1) G(t-s) ds dt - \int_{\theta T}^T r_2 \int_{\theta T}^t f(-r_2) G(t-s) ds dt \\ &\quad - \int_{\theta T}^T r_2 \int_0^{\theta T} f(r_1) G(t-s) ds dt \\ &= r_1 f(r_1) \underbrace{\int_0^{\theta T} \int_0^t G(t-s) ds dt}_{= C_{11}} + r_2 f(r_2) \underbrace{\int_{\theta T}^T \int_{\theta T}^t G(t-s) ds dt}_{= C_{22}} \\ &\quad - r_2 f(r_1) \underbrace{\int_{\theta T}^T \int_0^{\theta T} G(t-s) ds dt}_{= C_{21}}. \end{aligned} \quad (6)$$

We see that the cost of a round trip strategy where trades are accumulated and then liquidated at constant rates is $C_{11} + C_{22} - C_{21}$. So the condition for no dynamic-arbitrage states that

$$C_{11} + C_{22} - C_{21} \geq 0.$$

The first and second term, C_{11} and C_{22} can be regarded as the impact from stock purchases (sales) and the impact of previous purchases (sales). The last term, or the cross term as Gatheral calls it, is the impact of the stock sales together with the impact from previous stock purchases.

3.2.4 Power-Law Decay

In the BGPW picture the impact is assumed to decay as a power-law, i.e.

$$G(t-s) = \frac{1}{(t-s)^\gamma}, \quad 0 < \gamma < 1.$$

Inserting this into the expression in (6) yields

$$\begin{aligned} C_{11} &= r_1 f(r_1) \int_0^{\theta T} \int_0^t \frac{1}{(t-s)^\gamma} ds dt = r_1 f(r_1) \int_0^{\theta T} \left[\frac{-1}{1-\gamma} (t-s)^{1-\gamma} \right]_0^t dt \\ &= r_1 f(r_1) \int_0^{\theta T} \frac{1}{1-\gamma} t^{1-\gamma} dt = r_1 f(r_1) \frac{T^{2-\gamma}}{(1-\gamma)(2-\gamma)} \theta^{2-\gamma} \\ C_{22} &= r_2 f(r_2) \int_{\theta T}^T \int_{\theta T}^t \frac{1}{(t-s)^\gamma} ds dt = r_2 f(r_2) \int_{\theta T}^T \left[\frac{-1}{1-\gamma} (t-s)^{1-\gamma} \right]_{\theta T}^t dt \\ &= r_2 f(r_2) \int_0^{\theta T} \frac{1}{1-\gamma} (t-\theta T)^{1-\gamma} dt = r_2 f(r_2) \frac{T^{2-\gamma}}{(1-\gamma)(2-\gamma)} (1-\theta)^{2-\gamma} \\ C_{21} &= r_2 f(r_1) \int_{\theta T}^T \int_0^{\theta T} \frac{1}{(t-s)^\gamma} ds dt = r_2 f(r_1) \int_{\theta T}^T \left[\frac{-1}{1-\gamma} (t-s)^{1-\gamma} \right]_0^{\theta T} dt \\ &= r_2 f(r_1) \int_{\theta T}^T \frac{1}{1-\gamma} [t^{1-\gamma} - (t-\theta T)^{1-\gamma}] dt \\ &= r_2 f(r_1) \frac{1}{(1-\gamma)(2-\gamma)} [t^{2-\gamma} - (t-\theta T)^{2-\gamma}]_{\theta T}^T \\ &= r_2 f(r_1) \frac{T^{2-\gamma}}{(1-\gamma)(2-\gamma)} [1 - \theta^{2-\gamma} - (1-\theta)^{2-\gamma}] \end{aligned}$$

If we assume that our trading strategy is a round trip trade the impact function $f(\cdot)$ must be of a shape such that Proposition 1 is satisfied. If we insert $\theta = r_2/(r_1 + r_2)$ into the expressions derived above we get after some simplifying algebra

$$\begin{aligned} C_{11} + C_{22} - C_{21} &\geq 0 \quad \Leftrightarrow \\ f(r_1) [r_1 r_2^{1-\gamma} - (r_1 + r_2)^{2-\gamma} + r_2^{2-\gamma} + r_1^{2-\gamma}] + f(r_2) r_1^{2-\gamma} &\geq 0 \end{aligned} \quad (7)$$

If we assume a permanent impact, i.e. we let $\gamma = 0$ in (7), we get

$$\begin{aligned} f(r_1) [r_1 r_2 - (r_1 + r_2)^2 + r_2^2 + r_1^2] + f(r_2) v_1^2 &\geq 0 \\ \Rightarrow r_1 (f(r_2) v_1 - f(r_1) r_2) &\geq 0 \\ \Rightarrow f(r_2) r_1 - f(r_1) r_2 &\geq 0. \end{aligned}$$

Since the inequality above should hold for any positive r_1, r_2 we must have that $f \propto r$. This discovery was also made by Huberman and Stanzl [8], namely that

if we have permanent impact then f must be linear otherwise there exists a dynamic arbitrage opportunity. Moreover, one notes that different values on γ gives different restrictions on $f(\cdot)$ and we can conclude that including a power-law decay function generates a model compatible with more realistic impact functions than if permanent impact is assumed. We can also see that if we let $r = r_1 = r_2$ in (7), which means we are trading in and out at the same rate we get

$$2f(r)r^{2-\gamma} \geq 0 \Leftrightarrow f(r) \geq 0$$

which is true since f is assumed to be positive for positive r . This means that trading in and out at the same rate could never give rise to dynamic arbitrage.

3.2.5 Other Decay Functions

In the language of Gatheral [6] more decay functions are being considered. For example, he shows that an exponential decay implies that the impact function has to be linear in order to satisfy the no-dynamic arbitrage condition, and is thus able to exclude this as a reasonable option since the impact function has been shown empirically to be concave. Despite this, he does however name two publications where exponential decay is assumed.

3.2.6 Applicability of the Theory

The theory above from which the inequalities and conditions are derived was included in order to give us a sense of what a realistic shape of the impact and decay function might be. It is however important to realize that if this theory contradicts with empirical estimations, the error may lay on either side. On a first note it should be underlined that the no-dynamic arbitrage condition (7) refers to the expected cost of *hidden orders*. That is, it is assumed that other traders in the market are unaware of how much and how fast we want to sell and will thus not act as a response to this. As a result one can say that the conditions above were mainly derived with the purpose of describing the *cost of trading*, i.e. how much does single trades affect stock prices. In this thesis, we are more interested in how aggregated buy-/sell-pressure affects the price. Since aggregated trades over some time interval could not really be modelled as a hidden order, the theory above might not explain the data we will be looking at as good as if we were actually trying to model the impact from single trades. For instance, if ten consecutive sell orders were registered at one minute and ten sell orders the next this would most likely affect the correlation of order signs, and thus the decay function. However, if we instead looked at the minutely aggregated data over these two minutes the correlation is not caught. That being said, we still find it important to present the theory in order to show what research has been done within the field of market/price impact, and we also expect our model to behave somewhat similar, especially the instantaneous part.

We can also note that far from all combinations of impact and decay functions have been considered above. And for the combinations we have looked at we have only considered trading strategies that involves trading at constant rates. Dropping this assumption might give rise to different constraints.

3.3 Summary

In the first part of this section it was described how the uncorrelated nature of stock prices together with the correlation of order signs implies that the impact function must have some historic dependence. From the theory of no-dynamic arbitrage Gatheral [6] rules out certain combinations of impact function and decay kernel. For example, if price impact is assumed to be permanent, then f must be linear in traded volume. This is however at odds with empirical data which leads us to believe that this is not a realistic model. If we instead, in analogy with BGPW [3], assume a power-law decay, the theory of Gatheral shows that the impact function f can take much more realistic shapes.

4 Method & Data

This section explicitly explains the modelling process used in this work. Initially a general introduction of the data is provided, which is thereafter specified more thoroughly. The exploration of the data clouds exposes some trends that will eventually allow model fitting and parameter estimation in the instantaneous case, i.e. where temporal effects are ignored. Following the description of the instantaneous model is an introduction of the previously discussed temporal effect, and ultimately, a decay function $G(\cdot)$.

4.1 Data

The Stockholm Stock Exchange, operating under the name Nasdaq Stockholm, uses the stock market index OMX Stockholm 30 (*OMXS30*) as a yardstick to how the market is doing. In simple terms, OMXS30 consists of the 30 most traded stocks on the market, and it is information on these stocks that has been analyzed in this thesis (see Appendix B.1 for named list). In addition to the Stockholm Stock Exchange, the OMXS30 stocks can also be traded on the pan-European equity exchanges *BATS Chi-X Europe* and *Turquoise*. Both are London based low cost alternatives, and thus also rivals, to the primary European exchanges such as the London and Stockholm Stock Exchanges (recall Section 2.2 for more details).

4.1.1 Broad Explanation of Data

Trading data normally consists of some information per individual incoming order, making the data sets extremely large. If the analyzed orders are limit orders, these data sets typically specify variables like time and date of order placement, number of stocks requested to buy or sell (bid-size or ask-size respectively), the limit price at which the purchase or sale should be made (bid-price or ask-price respectively), and naturally some stock identifier. The computational cost that follows handling data of this magnitude becomes enormous, and considering that the objective of this thesis is to measure price impact over all aggregated trading data, a contraction of the data points into minute by minute intervals is easily justified. A simplified example of how the data looks at this point for one asset (Stock A) during one minute follows. Note that during this minute, only four orders come in and the order book is blank at the beginning of the example.

Example 4.1. A sell order comes in for 100 shares of Stock A at a price of 10 SEK; it is placed in the order book since no matching bid exists. A few seconds later a bid order of 90 shares comes in at bid-price 9 SEK, since this bid-price is lower than the available ask-price no match is made and it is saved in the order book. Later during that minute, another bid comes in at 10 SEK for the volume 80 shares, match! The bidder takes their 80 shares from the available 100 at the price 10 SEK, and thus a buy has occurred. Remaining in the order book is one ask order of 20 shares at ask-price 10 SEK and a buy order of 90 shares at bid-price 9 SEK. The net bought volume up to this moment is calculated as, $80 \cdot 10 = 800$ SEK (positive). Now another sell order comes in at ask-price 9 SEK for 50

shares. This matches a bid-price in the order book and a sell occurs, leaving the order book with one ask of 20 shares at 10 SEK and one bid of 40 shares at 9 SEK. The total net bought volume for this minute becomes,

$$80 \cdot 10 - 50 \cdot 9 = 350 \text{ SEK.}$$

It is clear from the example that a *sell* occurs when an incoming ask matches an existing bid, which generates a negative volume, whereas a *buy* occurs when an incoming bid matches an existing ask generating a positive volume.

4.1.2 Data Specifics

The trading data provided by Nasdaq AB has been aggregated per the above scheme into minute by minute observations. Each of the 30 stocks constituting OMXS30 has one row of data per minute. In particular, the data set consists of roughly half a million observations in 23 variables that all in all make up for the month of November¹ 2016 for both Nasdaq Stockholm and Chi-X. Until further notice however, only Nasdaq Stockholm data will be considered. As for the 23 variables, some have been used more extensively than others, and those are explained in Appendix A.

Of the 30 stocks in OMXS30, one has been excluded due to irregular behavior. Fingerprint Cards has had a troublesome year with news coverage that has impacted the stock in a very unpredictable manner, and has thus become a stock where traded volume is not the key driver in price changes. Therefore, only 29 stocks will be considered henceforth.

The current aim is to get a first view of the data in two variables; price impact as a function of aggregated volume. When it comes to observed price impact, one out of two measurements is usually preferred. The *mid-price change* is the change in mid-point price during the minute in question, recall that mid-point price is the mid-point between the best bid-price and ask-price. The option favored in this thesis is *micro-price change*, which is an order book weighted mid-point price, to account for possible differences in ask-size and bid-size. As previously mentioned, the input variable has been chosen as an aggregated volume, or *net bought*, which has been divided by the average turnover per minute for each particular stock in order to clarify comparisons. This input variable will be called v throughout the rest of this section (the font differs in the plots, but $v = \mathfrak{v}$). Note that the division by average turnover constitutes a significant alteration of the data in a way that becomes fruitful for the analysis in this section. This alteration will be analyzed more thoroughly in Section 6.2. For a more detailed explanation of variable calculations, see Appendix A (MICROPRICE_CHANGE and NET_BOUGHT_Vpct).

In Figure 4.1, the data clouds of the first nine stocks are displayed. Note that the values on the y-axis are multiplied by a factor of 10^4 in order to have price impact in BPS². For similar plots of all 29 stocks, see Figure C.2 in Appendix C.2.

¹In later sections, similar data sets for January-June will also be used.

²Basis Points = one hundredth of a percent




Figure 4.1: A visualization of the first nine stocks (alphabetically) in OMXS30.
 Axes: observed price impact = micro-price change
 v = aggregated volume \div individual minutely average turnover

4.1.3 Distribution of Data

As described in the Theory section, stock returns are assumed to follow a Gaussian distribution. Before we move on to estimating parameters, and ultimately drawing conclusions based on these, it is essential to know if normality is a good approximation for the data we will analyze. In Figure 4.2 a histogram of 30,000 randomly chosen points are drawn from the stock LUPE (Lundin Petroleum; which was randomly chosen). In the histogram, the shaded area represents the density function of a normal distribution. It is evident that more points are gathered around zero compared to the normal distribution. From the QQ-plot it is also evident that the actual distribution of the data is much more heavy tailed than a normal distribution. The data is actually rather poorly approximated by a normal distribution. This must be taken into consideration when looking at confidence intervals and standard deviations of estimates, as the R-functions from which these are calculated often assume normally distributed data.

4.2 Method: Instantaneous Model

The data in Figure 4.1 appears to be very noisy, although some trend can be spotted. Can the noise be reduced by adding a time dependent variable? In other words, does price impact contain some form of time dependency? This line of reasoning is recognized from previous sections and it results in adding a decay function. But before we account for any possible time dependency, it




Figure 4.2: Histogram and QQ-plot of data for the stock LUPE (Lundin Petroleum). The shaded area over the histogram represents the density function of a normal distribution. It is evident that the data is not very well approximated by a normal distribution. The same plots for different stocks showed the same patterns.

might prove fruitful to try some simplifications on the data.

The mentioned trend of the plots in Figure 4.1 is a slightly increasing one that some data transformations hopefully will make clearer. If we place the sorted v -values in N groups of more or less equal size; that is, equal in number of elements and not the length of intervals, the mean value of v within each quantile can be calculated. The new quantile mean of v will be denoted by \tilde{v} . Finally, the averages of the values on the y-axis are calculated within each quantile, i.e. for each of the N relevant \tilde{v} :s. The new, altered clouds of data for the first nine stocks can be seen in Figure 4.3. For similar plots of all 29 stocks, see Figure C.3 in Appendix C.3.

It is evident that the patterns in Figure 4.3 are following increasing and symmetric, or more specifically odd, functions. Moreover the patterns seem to be concave for positive x -values and convex for negative; we will refer to this as *signed concavity*. These convenient shapes provide a foundation for this stage of the model-fitting, which will be performed on the original data rather than this simplified version. To approximate the returns Y_t described above, it is reasonable to assume a function $f(\cdot)$ with the structure,

$$f(v_t) = \alpha_0 + \alpha_1 \cdot \text{sign}(v_t) \cdot |v_t|^\psi, \quad \psi \in (0, 1), \alpha_0 \in \mathbf{R}, \alpha_1 > 0, \quad (8)$$

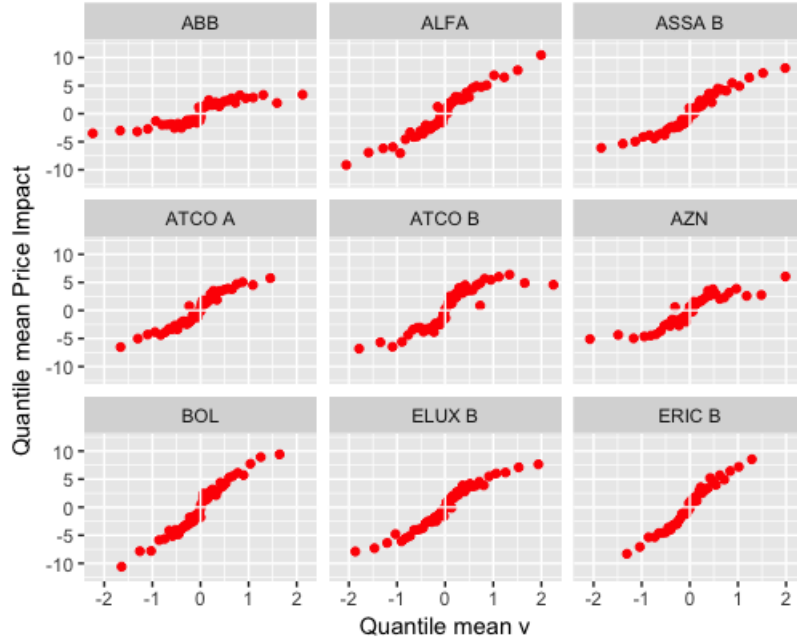


Figure 4.3: A visualization of the quantile means for the first nine stocks in OMXS30.

where α_0 , α_1 and ψ are to be estimated with a nonlinear least squared method, i.e. through,

$$Y_t = f(v_t; \alpha_0, \alpha_1, \psi) + \eta_t, \quad (9)$$

where η_t is the residual term. In terms of notation, the arguments after the semicolon in $f(\cdot)$ specifies what parameters are to be used and estimated, whereas v_t of course is the variable.

Before any statistical tests are performed however, it is in order to express a hypothesis about the estimated functions' behaviors. In order to satisfy the observed signed concavity, ψ should be strictly larger than zero and strictly smaller than one. We note that when ψ approaches one, the model approaches linearity, and when ψ approaches zero, the model approaches a step function with values $+1$ or -1 depending on the sign of v . Clearly ψ differs in size from stock to stock, as is seen in Figure 4.3, but values around $\psi = 1/2$ are to be expected. The variable α_0 should prove to be insignificant, i.e. $\alpha_0 = 0$ since all plots cross the origin. In qualitative terms that hypothesis is derived from the assumption that zero aggregated volume generates zero impact. As for α_1 , it explains the inclination of the function, and it is expected to differ in size between different stocks. Hopefully statistical tests will show that these instantaneous models can be concatenated into one collective function without loss of too much accuracy.

4.3 Method: Added Time Dependency

If one again looks at the plots in Figure 4.1 where we see the micro-price change in basis points as a function of the aggregated net bought over a minute divided

by turnover, one can see that this somewhat noisy relationship cannot easily be described by a function of aggregated net bought and turnover. Similar, or even equal, x-values seem to give rise to a variety of different y-values, which contradicts the very definition of a *function*.

As an example, we can take a closer look at the data we have from the stock ATCO A (Atlas Copco). At 2016-01-01 09:24:00-09:24:59 the aggregated net bought was -1563743 SEK, i.e. there were 1563743 more SEK worth of stock sold than bought during this minute. The price impact was at this minute -9.4507 BPS. Based on the theory we have presented this is somewhat expected; a big pressure on the sell side should lead to a drop in price. However, the same day but between 13:40:00 and 13:40:59 a similar net bought was recorded (-1515100 SEK) but during this minute the price of the stock instead increased with 4.6064 BPS! This example is maybe extreme and could possibly have been caused by some important news. There are however more examples like this, and they show that there is more than just the aggregated net bought that decides what the impact should be.

As we already have argued, the price impact model should have some dependence on historic data, described by the decay function. In order to examine this, we form the following time lagging model for the price impact R_t

$$R_t = f(v_t) + \sum_{s=1}^N \beta_s f(v_{t-s}), \quad (10)$$

where $f(\cdot)$ is the impact function presented in the previous section and β_s are the coefficients that should describe how the impact depends on previous data, we call them *decay coefficients*. Before estimating these values, it is interesting to analyze what kind of behavior we should expect. In Section 3 we explained how Gerig argues that, as a consequence of the highly correlated order signs, a surprise in order flow has to impact prices more in order for price changes to remain uncorrelated. This means that we should expect negative decay coefficients. We can demonstrate this by a simple example.

Example 4.2. We wish to calculate the price impact at $t = 3$ using the impacts from $t = 0, 1, 2$. Assume the net bought, v_t , has been positive for all previous time points, which implies $f(v_t) > 0 \quad \forall \quad t \leq 2$. Now consider two scenarios; we buy at $t = 3$ ($f(v_3) > 0$) or we sell at $t = 3$ ($f(v_3) < 0$). The impact at this time becomes

$$R_3 = f(v_3) + \beta_1 f(v_2) + \beta_2 f(v_1) + \beta_3 f(v_0)$$

If all the β_s are negative the absolute value of the impact will be greater for the scenario where $f(v_3) < 0$, i.e. where we sell at $t = 3$, than for the scenario where we buy. This is in accordance with the theory that a surprise in the order flow should impact prices more.

If we instead choose to compare our model setup with (3) in Section 3.2 (the one Gatheral uses), we can see that the β_s -values should form the decay function $G(\cdot)$. However, Gatheral, and a lot of other researchers, argue that this decay function should be positive and decay as a power-law. This seems to contradict the "surprise-in-order-flow"-theory.

We choose to model the price impact function $f(\cdot)$ in the same way as before. So (10) can be written as

$$R_t = \alpha_0 + \alpha_1 \cdot \text{sign}(v_t)|v_t|^\psi + \sum_{s=1}^N \beta_s (\alpha_0 + \alpha_1 \cdot \text{sign}(v_{t-s})|v_{t-s}|^\psi),$$

where the parameters α_0 , α_1 , and ψ are to be estimated with the instantaneous model. The decay coefficients will then be fitted using a linear regression model according to

$$Y_t = f(v_t) + \sum_{s=1}^N \beta_s f(v_{t-s}) + \eta_t, \quad (11)$$

where η_t is the residual term.

Once the decay coefficients β_s have been determined, they are to be approximated by a function, the decay function $G(\cdot)$. The shape of this function will be determined by the values obtained from (11) through another non-linear least square fitting.

5 Results

By now the general model has been set, and it is to be adapted into its final form by estimating parameters and analyzing the generated results. Through various tests and approximations the instantaneous model is justifiably simplified into a final model of square root shape with only one individual parameter per stock. Interestingly enough the decay function, which is an aggregate of the estimated coefficients over time lagged data, proves to be approximated best by an exponentially decreasing function at odds with our previous hypothesis.

The analysis in this chapter is conducted on both individual stocks and the entire set of stocks at once in order to make comparisons and draw conclusions about the nature of the market. When it came to extracting the temporal effects of price impact on one stock at a time, it proved necessary to introduce more data than originally needed, which will be explained in detail as the time arrives. Keep in mind that one of the objectives of this thesis is to make a comparative analysis of price impact on different trading markets, although the model is derived from the primary trading ground in Sweden; the Stockholm Stock Exchange (holds about 60%³ of trades involving OMXS30 when comparing all venues as of end 2016).

5.1 Instantaneous Part of the Model

As mentioned in Section 4.2 parameters are to be fitted using a non-linear least square method with the following model,

$$Y_t = \alpha_0 + \alpha_1 \cdot \text{sign}(v_t) \cdot |v_t|^\psi + \eta_t, \quad \psi \in (0, 1), \alpha_0 \in \mathbf{R}, \alpha_1 > 0, \quad (12)$$

where Y_t are the observed price changes, i.e. micro-price changes, and η_t is the residual term.

Table 1 presents the resulting parameters of each stock when the model in (12) is used for the month of November 2016.

From an initial look at Table 1 it is evident that all α_1 s and ψ s are significant on a very strict level, and that the hypothesis of $\psi \approx 1/2$ was reasonable. That the parameters α_0 were to be insignificant turns out to be partially true. The statistically significant α_0 s might exhibit their relatively large magnitude due to the method of division into quantiles, or simply due to noisy data. In fact, if the same nonlinear least square estimate scheme is conducted for the model without the parameter α_0 , the residual sum of squares increases only marginally (0.036% at most), which is why model (12) is updated to

$$Y_t = \alpha_1 \cdot \text{sign}(v_t) \cdot |v_t|^\psi + \eta_t, \quad \psi \in (0, 1), \alpha_1 > 0. \quad (13)$$

The parameters, α_1 and ψ , in the fitted models obtain new but very similar values when version (13) is fitted. The new values, that are not presented here, fall in the ranges $\alpha_1 \in [2.533, 9.671]$ and $\psi \in [0.304, 0.702]$.

It would be convenient if the 29 individual models could be contracted into one, but is it justifiable to assume pairwise equality between the ψ s and α_1 s?

³As seen in Figure C.1 in Appendix C.1.

³Residual Sum of Squares

⁴Simultaneous fit for model including α_0 , see (12)

⁵Simultaneous fit for model excluding α_0 , see (13)

| Stock Code | α_0 ($\times 10^{-4}$) | α_1 ($\times 10^{-4}$) | ψ | RSS ³ ($\times 10^{-3}$) |
|------------------------------|---------------------------------|---------------------------------|---------|---------------------------------------|
| ABB | 0.091 | 2.532 | • | 1.72 |
| ALFA | 0.322 • | 5.825 • | • | 3.15 |
| ASSA B | 0.111 | 4.749 • | • | 2.61 |
| ATCO A | 0.349 • | 4.454 • | • | 3.21 |
| ATCO B | -0.076 | 4.823 • | • | 3.96 |
| AZN | -0.122 | 3.670 • | • | 5.82 |
| BOL | 0.148 | 7.202 • | • | 6.37 |
| ELUX B | -0.064 | 5.467 • | • | 3.06 |
| ERIC B | 0.266 • | 6.854 • | • | 4.39 |
| GETI B | 0.160 ◦ | 5.089 • | • | 2.29 |
| HM B | 0.057 | 4.599 • | • | 2.41 |
| INVE B | 0.031 | 3.260 • | • | 1.76 |
| KINV B | -0.169 ◦ | 4.300 • | • | 1.84 |
| LUPE | -0.048 | 6.664 • | • | 4.54 |
| NDA SEK | 0.166 • | 4.096 • | • | 2.41 |
| NOKIA SEK | -0.159 | 4.562 • | • | 6.19 |
| SAND | 0.265 • | 6.092 • | • | 3.28 |
| SCA B | -0.132 ◦ | 4.283 • | • | 2.20 |
| SEB A | 0.137 ◦ | 4.027 • | • | 2.31 |
| SECU B | 0.014 | 5.948 • | • | 4.11 |
| SHB A | 0.114 | 4.793 • | • | 2.49 |
| SKA B | 0.053 | 4.251 • | • | 2.44 |
| SKF B | 0.133 | 5.807 • | • | 4.00 |
| SSAB A | 0.020 | 9.670 • | • | 8.96 |
| SWED A | 0.053 | 3.760 • | • | 1.96 |
| SWMA | -0.029 | 3.874 • | • | 2.02 |
| TEL2 B | -0.220 • | 3.958 • | • | 2.13 |
| TELIA | -0.136 • | 3.411 • | • | 1.32 |
| VOLV B | 0.048 | 4.499 • | • | 2.79 |
| Mean | 0.0429 | 4.92 | 0.535 | 3.3 |
| Standard Dev. | 0.0214 | 2.07 | 0.093 | 1.72 |
| Sim. (12)⁴ | 0.0491 ◦ | 4.87 • | 0.527 • | 99.2 |
| Sim. (13)⁵ | | 4.87 • | 0.527 • | 99.3 |

Table 1: Level of significance indicated by • = 0.1%, ◦ = 5%; thus • implies ◦.

A 5%-confidence interval ($\pm 2.5\%$) is provided with the nonlinear least square function in R. In Figure 5.1, estimates of the parameters α_1 and ψ are plotted as dots along with their confidence intervals as vertical blue lines. The red line represents the estimates found in the final row of Table 1.

According to the confidence intervals neither parameter is equal for all stocks. However, ψ does seem to be more forgiving in this respect. It should be mentioned again that the confidence intervals are calculated under the assumption of normally distributed data; an assumption that we showed is a rather poor approximation of our data. However, what is seen in Figure 5.1 is what inspired the idea to test whether or not it is reasonable to fix one or both of the parameters. It is in our interest to minimize the number of variables in this early stage of the model building since an effort will be made to explain the remaining variables through external factors.

If parameters α_1 and ψ are both fixed (either as the mean of all 29 estimated individual ones or as the estimated values generated through collective fitting), the accuracy lessens significantly. The same worsening effect is observed when ψ is allowed to vary but α_1 is fixed. As it turns out, a fixed ψ and varying

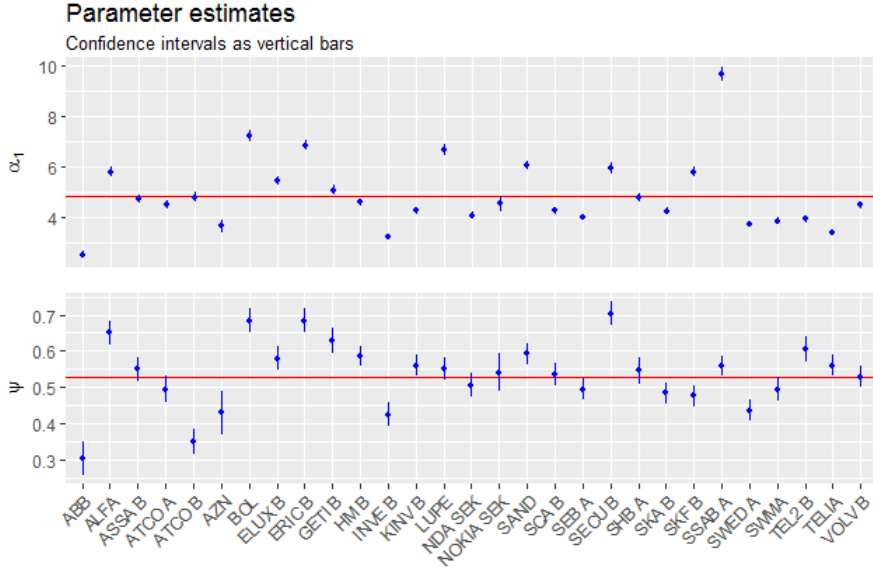


Figure 5.1: Estimates of the parameters in model (13) along with their confidence intervals.

α_1 gives the best trade-of between convenience in number of parameters and accuracy. Therefore, the parameter ψ will hereafter be set as,

$$\psi = 0.5,$$

hence the non-temporal model obtains a square root appearance. Very much in analogy with previous research. Note that also the ψ estimated overall ($\psi = 0.527$) was tested, but the result turned out nearly identical in terms of RSS as well as outcome of the parameters α_1 .

In more detail, the chosen simplifications discussed above are derived from a comparison between (13) and the model

$$Y_t = \alpha_1 \cdot \text{sign}(v_t) \cdot \sqrt{|v_t|} + \eta_t, \quad \alpha_1 > 0. \quad (14)$$

We find a very slight increase in error (maximum 1.3% in terms of RSS), and new but yet again very similar values of α_1 . The parameters α_1 exhibit the exact same pattern as that seen in Figure 5.1 when plotted, however the confidence intervals become somewhat tighter.

Given the presented results and approximations, we leave the instantaneous model as seen in (14), hence generating an updated version of the function $f(\cdot)$ seen in (8):

$$f(v_t) = \alpha \cdot \text{sign}(v_t) \cdot \sqrt{|v_t|}, \quad \alpha > 0, \quad (15)$$

where the index of α_1 has been dropped, and is therefore renamed α .

5.2 Temporal Part of the Model

By the end of Section 4.3, it was determined that a linear regression method was to be used to estimate the parameters β_s in model (11). Explicitly, the

following model will be fitted,

$$Y_t = f(v_t) + \sum_{s=1}^N \beta_s f(v_{t-s}) + \eta_t, \quad (16)$$

where $f(\cdot)$ is taken from (15) with its parameters set as constants, and η_t is the residual term.

In the plots of Figure 5.2 decay coefficients have been fitted for 100 lags on six of the 29 analyzed stocks. We have also included the 95% confidence intervals for the estimates, represented by the vertical lines crossing each data point. Since the confidence intervals are a bit hard to see we have market significant estimates as red, and insignificant as blue. Again, the confidence intervals, and thus the significance of estimates, are based on the rather inaccurate approximation of normally distributed data. Conclusions are thus not solely based on these.

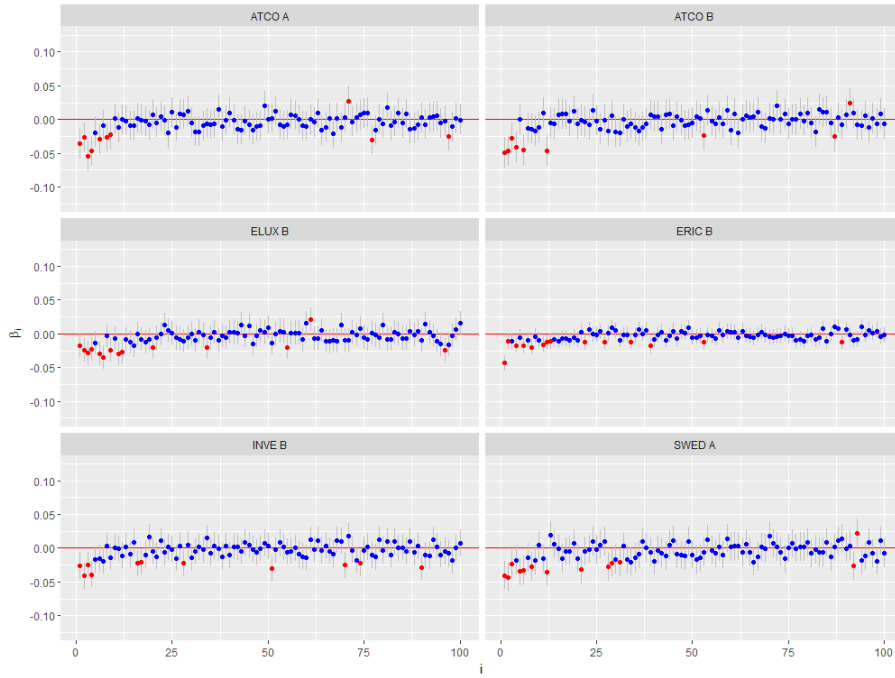


Figure 5.2: Decay coefficients for six stocks. Significant estimates are marked red, while insignificant ones are marked blue.

One can start by looking at the decay coefficients after lag 30 and note that there is very little structure. The decay coefficients seem to be close to uniformly distributed, not around zero but around some small negative number as the majority of the points are negative. If we shift focus to the first 30 data points we can see that they seem to follow an increasing trend. For the stocks plotted in Figure 5.2, this is probably most evident for ATCO A (Atlas Copco), but for the other stocks one can at least see that the majority of the data points are negative. This is in line with what we expect when arguing from the Gerig

perspective. However, far from all of the 30 first data points, only around 5-10 depending on which stock we look at, are classified as significant when using a 95% confidence interval. When looking at these plots for the other 23 stocks we see the same pattern; mostly negative values, an increasing trend for the first ~ 30 lags, but a great part of them classified as insignificant. The trend indicates that recent trades affect the price more than earlier ones; a behavior that seems reasonable and is thus in our interest to verify. However, the fact that a lot of the decay coefficients are very close to zero and also classified as insignificant makes it hard to do this when looking at stocks individually. The small values (and the insignificance) of the coefficients could be interpreted as price impact not being dependent on past trading history. This interpretation does however contradict the theory that has been presented in earlier sections, which is why we will investigate this further.

In order to find out whether the insignificance of the decay coefficients is reasonable or not we can first compare the estimates for all stocks. When doing this we discover that for the first ten lags separately, the average number of stocks that has a positive coefficient are 5.44. To clarify, for each β_i when $1 \leq i \leq 10$, roughly 5-6 stocks out of 29 have $\beta_i > 0$, the rest have negative β_i . Moreover, the positive coefficients seem to be randomly distributed among different stocks, i.e. we do not see that one stock have positive decay coefficients for all $i \in [1, 10]$. If the coefficients were truly insignificant, wouldn't we see roughly the same amount of positive and negative ones? From a probabilistic point of view it is highly unlikely that, if positive and negative coefficients were equally likely, we would for a certain lag, say β_k , see 6 or less coefficients with positive sign. One can easily calculate this probability as

$$\sum_{n=0}^6 \binom{29}{n} \cdot 0.5^{29} \approx 0.001158.$$

And since this is the case for almost all β_i when $i \leq 10$ this becomes even more improbable. Furthermore, we can look at the mean of the decay coefficients for all stocks, which is presented in Figure 5.3.

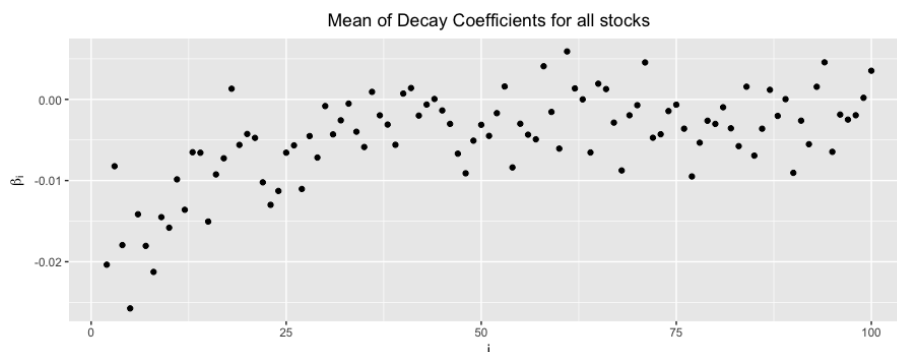


Figure 5.3: Mean value of decay coefficients for all 29 stocks

Here we can see that all the coefficients, but perhaps two or three, are negative (from the first 30 lags), and they seem to be increasing (or decreasing in magnitude). From this plot one can also note that the vast majority of the points are

negative even for lags up to 100, although the increasing trend seems to stop after around lag 30.

Finally, if we try to model the price impact on all the stocks at the same time, similar to what we did for the instantaneous impact function, the pattern discussed above becomes even more clear. The estimated decay coefficients for 100 lags is show in Figure 5.4.




Figure 5.4: Decay coefficients estimated for all 29 stocks collectively. The increasing structure is much more clear. It is also easier to see that the majority of the estimates are negative.

First we note that this plot is very similar to that of Figure 5.3. We see a very clear pattern of negative and increasing coefficients. A lot more coefficients are also classified as significant. This could be an indicator that the coefficients are indeed less than zero, but when estimating these for individual stocks we do not have enough data to get a small enough standard deviation to verify this. We can further see that 81/100 coefficients are negative and that none of the positive coefficients are significant. We did also look at lags greater than 100 but after this there seemed to be a close to even split between negative and positive coefficients. We can conclude that when estimating the decay coefficients for all stocks at the same time they satisfy the following:

1. The majority is less than zero up to lag ~ 100
2. Based on the amount of data we have used, the coefficients after lag ~ 100 can be approximated by zero.
3. The magnitude decreases until lag ~ 30 , after which there seems to be no structure except for the majority of the coefficients being negative.

We cannot statistically say that the conditions above hold for individual stocks, even though we have reasons to believe they should. What is important to note is that when the decay coefficients for all stocks combined were estimated, 29 times more data could be used compared to when each stock where estimated individually. A hypothesis is that the characteristics mentioned in the three points above hold for each stock individually, it is merely a question of having

enough data to support it. In order to, maybe not prove, but at least strengthen this hypothesis we estimate the decay coefficients for each stock on a bigger data set. Instead of just using one month of data we now use six, from January 2016 to June 2016. When doing this, the structure of the decay coefficients is much more similar to the one shown in Figure 5.4. In the Figure 5.5, one can see the estimates of the decay coefficients when one month of data is used compared to when six months of data is used. We have deliberately plotted three stocks where the change is easily spotted. In Appendix C.4 one can see all estimates both when one and six months of data are used.

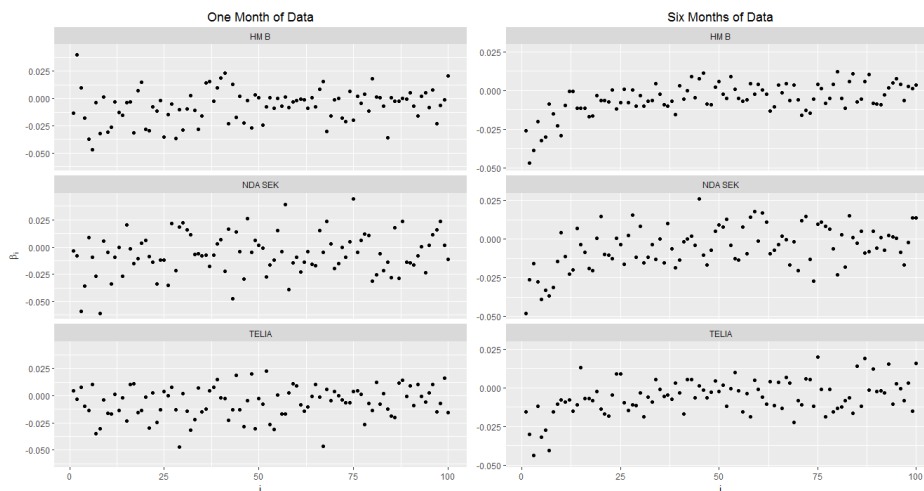


Figure 5.5: Decay coefficients estimated for three stocks with one month and six months of data

With the result from the six months estimates we can conclude that for all of the stocks, a great majority of the decay coefficients are negative, and for almost all stocks one can detect an increasing pattern for the first 15-25 lags. Since this increasing pattern is important and was not as clear when only one month of data was used, all estimates on individual stocks will from this point be made on the six-month data set.

5.3 Decay Function

As mentioned above the decay coefficients seem to be negative and increasing with the number of lags. This is not really in line with the theory of a positive power-law decay function. As a matter of fact, negative and increasing betas is the opposite of positive decaying ones. Since this section is dedicated to finding a functional form of the decay coefficients comparable to previous research the stock returns can be expressed as

$$R_t = f(v_t) - \sum_{i=1}^N G(i)f(v_{t-i}), \quad (17)$$

where $G(i)$ will be estimated using $-\beta_i$. For notational convenience, $-\beta_i$ will be denoted by $\hat{\beta}_i$. In Figure 5.6, $\hat{\beta}_i$ is plotted for the same stocks as in Figure 5.2.

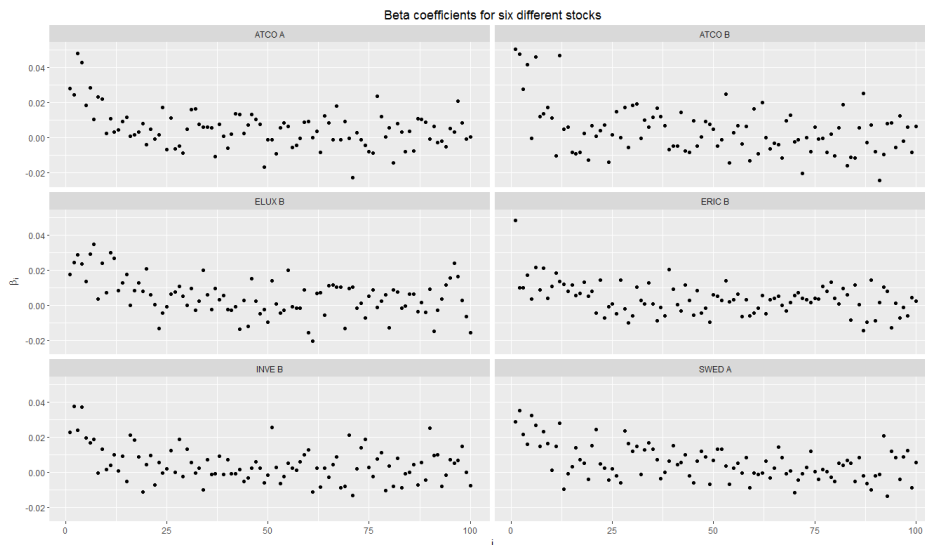


Figure 5.6: Estimated $\hat{\beta}_i$:s for six stocks

Even though the data points are somewhat noisy one can see that there is a non-linear decay that flattens out after around lag 30. The intercept and the steepness of the decay seem to vary a bit from stock to stock, but overall the structure seems to be very similar for all stocks.

5.3.1 Shape of the Decay Function

Since the decay seems to be similar for each stock we will use the data seen in Figure 5.7 to find the most suitable shape of the decay function $G(\cdot)$ and then estimate the parameters of this model for each stock. The lag coefficients seen in the plot are estimated in the same way as the ones seen in Figure 5.4, however these ones are estimated on a bigger data set (Jan 2016 - June 2016). This is because, as mentioned above, we will later on use this data set to estimate the parameters on each stock. As one can see from the plot in Figure 5.7, the points follow a very similar structure to what one can see in Figure 5.4, but as a result of the bigger data set the confidence intervals (that are not included in the plot) are smaller.

It is obvious that $G(\cdot)$ should not be linear, which is why we try to fit three different non-linear decreasing functions to the data seen in Figure 5.7:

1. Exponential function: $G_{exp}(\tau) = a \cdot e^{b\tau}$
2. Logarithmic function: $G_{log}(\tau) = a + b \cdot \log(\tau)$
3. Power-Law function: $G_{pow}(\tau) = a \cdot \tau^b$

The results of the non-linear least square estimation are presented in the following table and visualized in Figure 5.8.

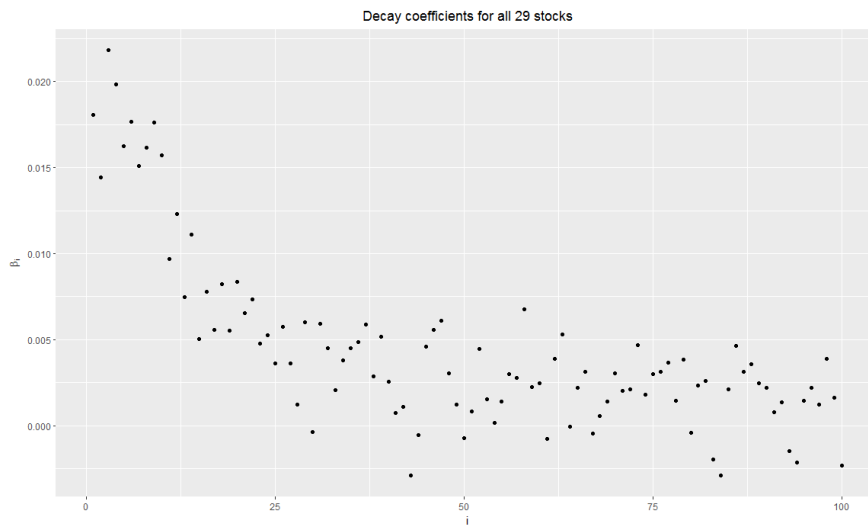


Figure 5.7: $\hat{\beta}_i$ s estimated on all 29 stocks.

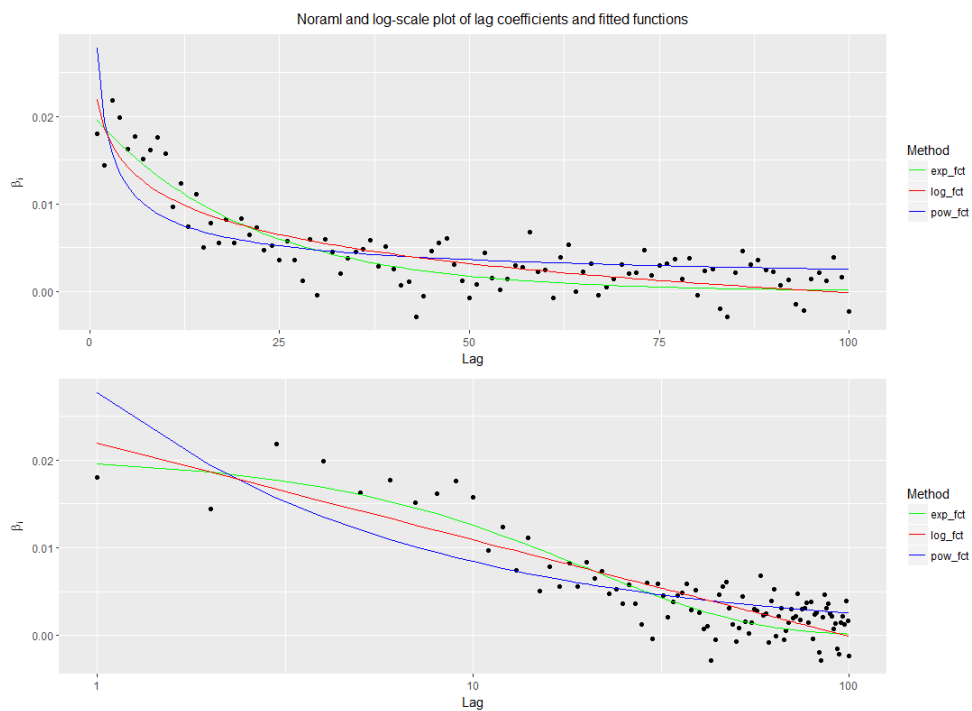


Figure 5.8: Fitted functions in normal and logarithmic scale

The number after the \pm -sign indicates the 95% confidence interval for the estimate. From Figure 5.8 it is not obvious, but possible to see that the power law-function fits the data worse than the exponential and logarithmic function. This becomes clearer when looking at the residual sum of squares. A surprising result since the most discussed decay function in previous literature is in fact

| Method | Estimates (a, b) | | RSS |
|-----------------|---------------------|----------------------|---------|
| $G_{pow}(\tau)$ | 0.0277 ± 0.0047 | -0.5191 ± 0.0701 | 0.00096 |
| $G_{log}(\tau)$ | 0.0219 ± 0.0021 | -0.0048 ± 0.0050 | 0.00067 |
| $G_{exp}(\tau)$ | 0.0205 ± 0.0022 | -0.0491 ± 0.0037 | 0.00058 |

Table 2: Estimated parameters of three different types of decay functions along with their respective confidence intervals and residual sum of squares (RSS).

the power-law function. The table above also tells us that, even though it is very hard to tell from the plots, the exponential function is, at least based on the RSS, a better fit than the logarithmic. In terms of the residuals it is hard to separate the accuracy of the models. A Shapiro-Wilk normality test indicates that we cannot reject normality of neither the log nor the exponential function, and QQ-plots further show that both the exponential and the logarithmic fit produce close to normally distributed residuals.

Furthermore, each of the different functional forms has a disadvantage. The exponential decay violates the no-dynamic arbitrage principle as discussed in the Theory section, while the logarithmic decay can give rise to negative $\hat{\beta}_i$:s. Since we expect the $\hat{\beta}_i$:s to be positive the logarithmic function should satisfy the following constraint:

$$a + b \cdot \log(\tau) > 0 \quad \forall \quad \tau \in [1, 100]$$

We have $a > 0$ and $b < 0$ and since $\log(x)$ is strictly increasing it suffices to say that

$$-\frac{a}{b} > \log(100) \approx 4.60517.$$

This condition is in fact not satisfied by the estimated parameters a and b in the above table, and it is not unreasonable to assume that when fitting these parameters on individual stocks the condition above will be violated as well.

All in all, the logarithmic and the exponential decay both follow the data well. Based on the theory presented however, both functional forms have their pros and cons. The exponential decay is strictly positive but violates the no-dynamic arbitrage condition. The logarithmic decay on the other hand gives rise to negative values, and it has not been calculated whether it violates the arbitrage conditions or not, making a comparison between the two either inconclusive or in favor of the exponential function. The final means of comparison available is the RSS, which is slightly higher in the logarithmic case, which is why we move on to using the exponential decay when fitting the parameters on individual stocks in the next section.

5.3.2 Performance on Individual Stocks

Having concluded that the decay coefficients follow an exponential structure we now move on to fitting the a - and b -parameters on all stocks separately. In the following plots one can see the estimates for the a - and b -parameters together with their 95% confidence intervals.

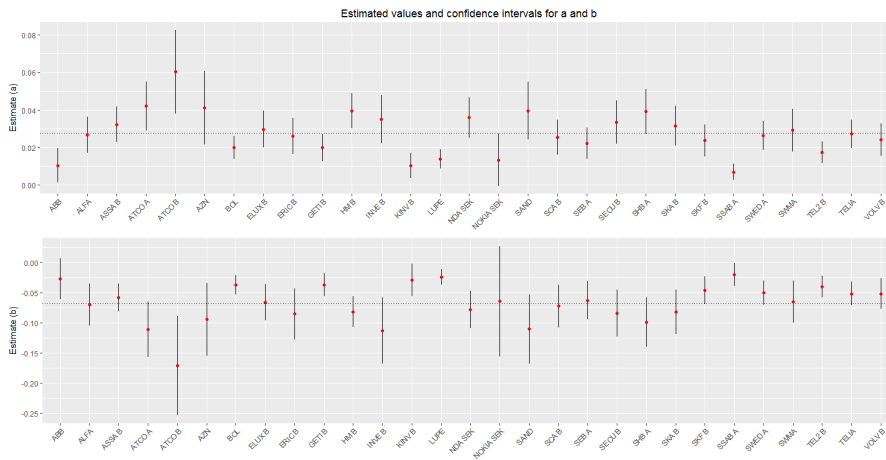


Figure 5.9: Estimated values of the a- and b-parameters for all stocks

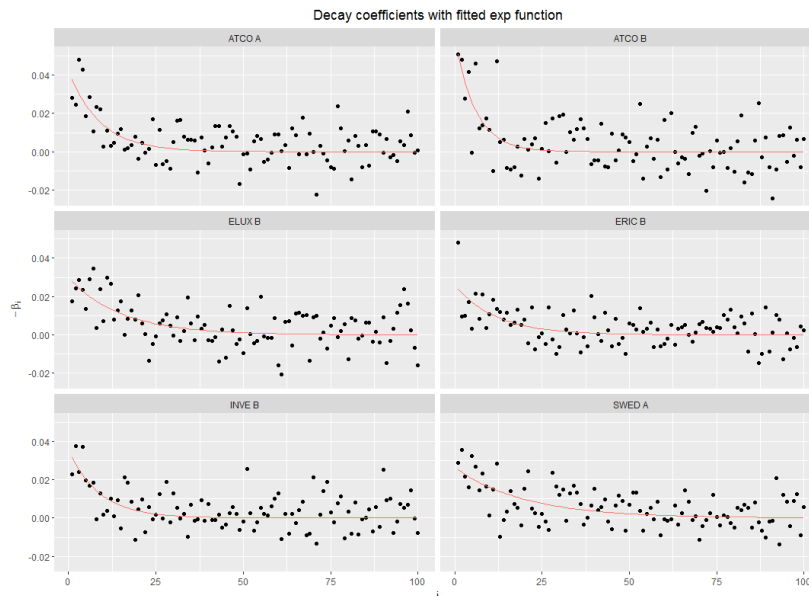


Figure 5.10: Exponential function fitted on decay coefficients of six different stocks

As a result of the noisy structure of the decay coefficients, the fitted function is a pretty rough estimate. The mean RSS of all estimates is 0.00885, which is roughly 15 times higher than the RSS for the function estimated on all stocks at the same time. However, one can see in Appendix C.5, where the fitted functions for all 29 stocks are shown, that some stocks have very high variance of the data, which obviously increases the mean RSS. In order to determine whether exponential decay is a reasonable fit for the individual stocks the mean RSS can be compared to some base line estimate, such as just estimating the data with the mean value of all data points. The RSS of this mean estimate is 0.085, i.e. 10 times higher than the mean RSS of the exponential fit.

In Figure 5.9 one can see that the confidence intervals of the estimates are relatively broad, especially for ATCO B, AZN and NOKIA (Atlas Copco, Astra Zeneca and Nokia). Many of the confidence intervals covers the mean of all estimates which is represented by the dotted line, the parameter b more so than a . This could indicate that using the mean values of a and b might not be that much worse of an estimate. In Figure 5.11, we have fitted the decay functions using $a = 0.0277$ and $b = -0.0681$ for the same stocks as in Figure 5.2. One can see that the fit is quite rough, but the RSS of 0.00938 reveals that it is actually not that much worse than when unique parameters for each stock were used. Thus, we have found an estimation of the decay function that is independent of what stock we try to model.




Figure 5.11: Exponential function with fixed a and b and decay coefficients of six different stocks

5.4 Summary

The obtained results of Section 5.1 allowed for a justifiable simplification of the instantaneous model. This part was developed to capture the trends of price impact as a function of aggregated volume, which in its own constitutes the foundation of a final model. The section started with having three stock-specific parameters that were later lessened to only one with minimal loss of accuracy. A result that improves the simplicity of the model, and in turn allows for more generality.

When analyzing the decay coefficients estimated on all stocks there was a clear decaying structure and the data was well estimated with both a logarithmic and an exponential decay. According to the RSS the exponential decay was a slightly better fit. Even though it violates the no dynamic arbitrage condition,

we used this model for our shape of the decay function over the logarithmic shape based on three key points:

1. The logarithmic function produced negative $\hat{\beta}_i$:s that are unwanted.
2. We still do not know if the logarithmic function satisfies the no dynamic arbitrage conditions.
3. The RSS of the exponential fit is lower.

When using an exponential decay on each stock separately the estimates became more noisy resulting in higher RSS. However, despite the noisy data and rough estimate, the decay coefficients are clearly better approximated by an exponential function than the mean value of all data points for all of the 29 stocks, indicating that the exponential shape is present. Further we saw that fixing both a and b to the mean of all estimated values gave a relatively good estimate which is why we choose to model the decay coefficients as:

$$G(\tau) = 0.02877e^{-0.0681\tau}. \quad (18)$$

6 Discussion

It is evident from the previous sections of this thesis that the way price impact affects stock movements cannot be ignored. Price impact has also been seen to be an odd and increasing function of net bought relative to a stock's turnover with a dependence on previous trades. Moreover it is concave in order size. These characteristics have been verified for 29 different stocks on data taken with minutely intervals for six months. We will in this section discuss what these characteristics might imply for the stock market in general. We will also more thoroughly discuss what was just mentioned, that price impact is a function net bought relative to a stocks *turnover*, and how this might affect the market structures we see today. However, before discussing the broader, and maybe more easily verified, characteristics of price impact we will analyze the specifics of the model. Since a lot of simplifications were made when deciding both the instantaneous impact function and the decay function it is interesting to see if our model is reasonable, if can reconstruct the characteristics of how stock prices move and how it compares to other models derived in previous literature on this topic.

6.1 Model Evaluation

At this point a model for price impact has been derived. First the immediate impact function was estimated, and then in order to account for the correlation in order signs as well as the noisy behavior of the data clouds, we estimated a decay effect to the impact function. The final model, after some simplifications, turned out to be:

$$Y_t = f(v_t) - \sum_{\tau=1}^{100} G(\tau)f(v_{t-\tau}) + \eta_t, \quad (19)$$
$$f(v_t) = \alpha \cdot \text{sign}(v_t)\sqrt{|v_t|}, \quad G(\tau) = 0.02877 \cdot e^{-0.0681\tau},$$

where v_t is the fraction between the number of stocks bought and the average minutely turnover for the specific stock analyzed, and α is a stock specific parameter. Moreover η_t is a noise term which includes price changes caused by other factors than the buy-/sell-pressure, such as budget reports and other news.

It is now time to see if this model is able reconstruct the characteristics of stock returns, i.e. are the returns uncorrelated and to what extent do they follow the actual stock movement? We will also see how much of the variance in price movements our model suggests to be a result of the buy-/sell-pressure.

6.1.1 Characteristics of Stock Returns

Theoretically, if we managed to split up the stock returns in changes caused by price impact and changes caused by other factors, and if we assume that these other factors are not autocorrelated, we would expect the price changes caused by price impact to be uncorrelated as well. As mentioned numerous times in this thesis the autocorrelation of trade sign would give rise to autocorrelated

price changes if no decay function was added to the model. Thus, we added this, and as one can see in Figure 6.1, the estimated change in micro-price caused by price impact seems to be uncorrelated, at least if we assume that the actual micro-price changes are not autocorrelated. In the plots in question, it looks as if there is a correlation at lag one. However it should be underlined that the micro-price changes are not normally distributed, one should therefore not pay too much attention to the confidence intervals, but instead compare the two autocorrelation functions.

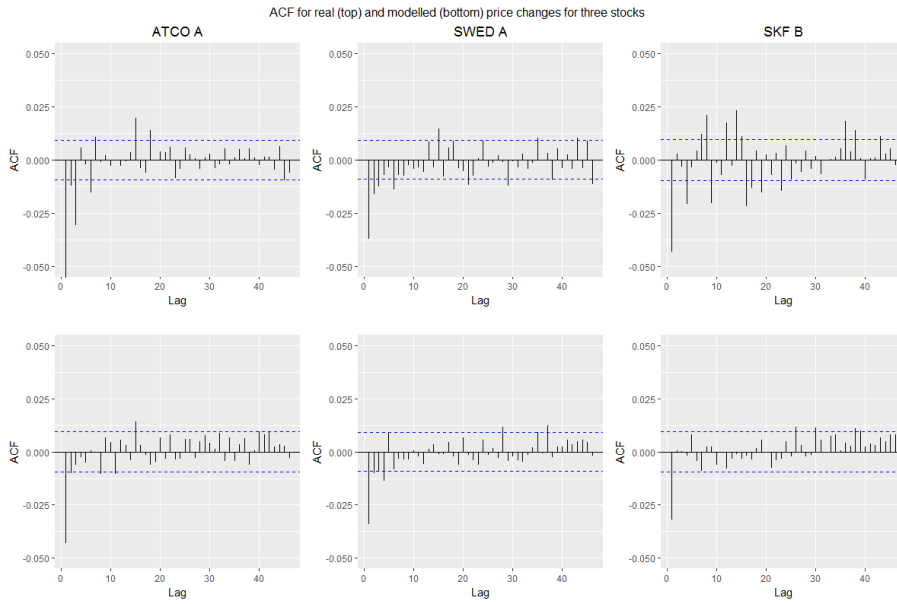


Figure 6.1: Autocorrelation function for real micro-price changes (top) vs. price changes caused by price impact (bottom).

We can compare the plots in Figure 6.1 to the autocorrelation function plot of the model without decay function. As one can see in Figure 6.2 the returns are now (more) autocorrelated, meaning that the decay function does serve its purpose by cancelling out the autocorrelation from the order signs.




Figure 6.2: Autocorrelation function for price impact when no decay function is added.

6.1.2 How Much is Explained?

As we saw above the decay function is necessary in order to produce uncorrelated stock returns, but we also added the decay function with hopes of explaining more of the variance in the micro-price changes. That being said, we did not expect the micro-price change explained by our price impact function to have the same variance as the real data cloud, since that would imply that the variance caused by factors such as news would not induce any variance in the stock prices (i.e. $\text{Var}(\eta_t) = 0$).

The real micro-price changes are plotted as a function of the aggregated net bought together with the micro-price changes estimated from our price impact function in Figure 6.3. One can clearly see that the variance is much higher for the real data. This is the case for all 29 stocks analyzed, and based on these plots it can be concluded that the decaying part of the impact function is necessary in order to keep prices uncorrelated, but in terms of predicting the price changes, the mean and variance of the residuals are not that much worse without the decay function.

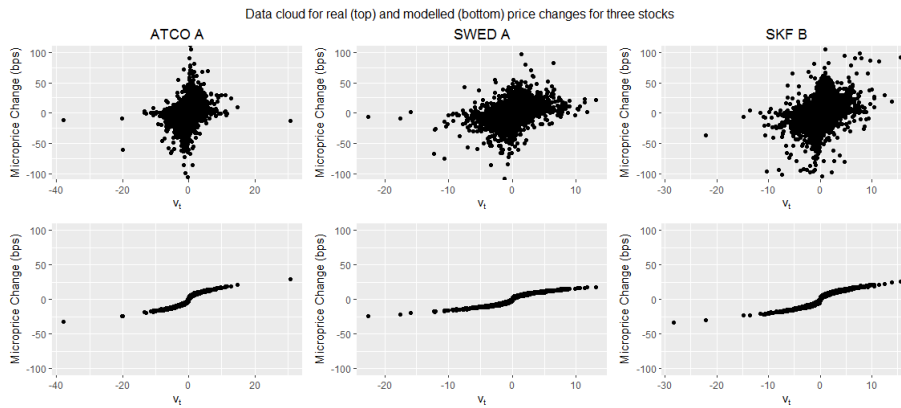


Figure 6.3: Real cloud of micro-price change as a function of buy-/sell-pressure vs changes in micro-price caused by price impact.

In order to quantify, the variance of the modelled price impact is only about 10-35% of the variance in the real micro-price change, depending on which stock we look at. It can also be mentioned that the mean variance of the micro-price change for all stocks between January 2016 and June 2016 is around 84 (BPS) while the mean variance of the estimated price changes caused by price impact is only 20 (BPS). Almost all of this variance comes from the instantaneous impact function (around 19 BPS on average).

6.1.3 Real vs. Predicted Prices

Now observe the plot of the returns caused by price impact and compare it to the actual returns. It should be mentioned that in the plots of Figure 6.4 & 6.5, the blue line named 'pred' might be misinterpreted as an actual prediction of the stock movements. This is not the case, the blue line rather represents how the stock would have moved if, according to our model, all the price movements of a stock were a result of aggregated buy-/sell-pressure. However, we will

sometimes refer to the blue line as 'predicted returns' due to lack of a better word.

When we looked at the data between January and June 2016 we saw that in January, the blue and the red lines differed a lot, and we thus chose to compare this plot to a month where the real and predicted returns differed less. That is why there are two plots of the same six stocks, one for returns in January (Figure 6.4) and one for April (Figure 6.5).




Figure 6.4: Real returns vs predicted price impact-returns in January 2016.

For the other months (Feb, Mar, May and June), the two lines followed each other pretty well, not as good as in April, but much better than the January plot. This indicates that returns are (often) increasing if there is a pressure on the buy side, and decreasing if there is a sell pressure. However, for all stocks one can find segments where the red and the blue line move in opposite directions, and in January this seems to happen at all instants. The first thought we had was that maybe it is not a coincidence that the price impact function followed the returns poorly this specific month, as usually there is a lot of news and people might be trading based on upcoming full year reports and such. However, would it not be reasonable to assume that if a lot of bad news hit a stock a specific month this would be reflected in the buy-/sell-pressure? For instance, would we not see a greater sell pressure if there is a lot negative news regarding a certain stock? Clearly, this does not seem to be the case, which might sound counter intuitive, but imagine a certain time period where there are no specific news regarding a certain stock. If there is a pressure on the buy side, i.e. sell orders are being filled more often than the buy orders, then of course market participants will notice




Figure 6.5: Real returns vs predicted price impact-returns in April 2016.

this and is able to add more expensive sell orders which will continue being filled and cause a price increase until the buy pressure stops. Imagine then, without loss of generality, that negative news about the stock is published. People will then expect the price to go down and add sell orders at a certain price into the order book. This price will however rather be conditioned on the importance of the news than the buy-/sell-pressure in the order book. And hence a lot of the sell orders could potentially be filled quickly if people value the importance of the news differently, so even though more sell orders than buy orders are being filled (i.e. there is a buy-pressure) the stock price will fall because the orders are being filled at a lower price. This could explain why the actual stock price goes down even though we experience a buy-pressure. Recall the Efficient Market Hypothesis which states that all available information regarding a stock should be reflected in its price. In times of no new public information about stocks the only information market participant receives is the one reflected in the buy- or sell-pressure and thus it seems reasonable that prices will move according to this. However as public news are released the trading is instead conditioned on these, and the buy-/sell-pressure is no longer as important.

6.1.4 Is the Model Accurate?

What is clear from the plots and numbers shown in the previous section is that the model in (19) reconstructs the data clouds poorly (as we see in Figure 6.3). The variance of the real micro-price changes is clearly higher and behave differently than our model. However, this fact does not have to imply that our impact function is not accurate. Remember that we are not trying to find a model for the total movements in stock prices, but only the movements caused

by buy-/sell-pressure. What Figure 6.3 tells us is that (under the assumption that the price impact function is accurate) a lot of the movements in stock prices are caused by other factors than the buy-/sell-pressure. But how do we know if the price impact function in (19) is accurate? Testing our model becomes a bit tricky since there is no data that says exactly why a stock price changes, just that it does. However, one way to try and find out how well the model is able to describe the price changes caused by price impact is to study the residual term η_t . If the first terms of (19) describe all price changes caused by the buy-/sell-pressure then the residual term η_t should theoretically describe the price changes caused by other factors, and one would thus expect this term to be uncorrelated with v_t . As one can see in Figure 6.6, the data points of η_t seems to be neither increasing nor decreasing with v_t . The pattern one can detect from the plots is that there are more extreme values around $v_t = 0$. This can however be explained by the stochastic nature of η_t . There are more data points around $v_t = 0$ and thus the number of extreme values would from a probabilistic point of view increase.

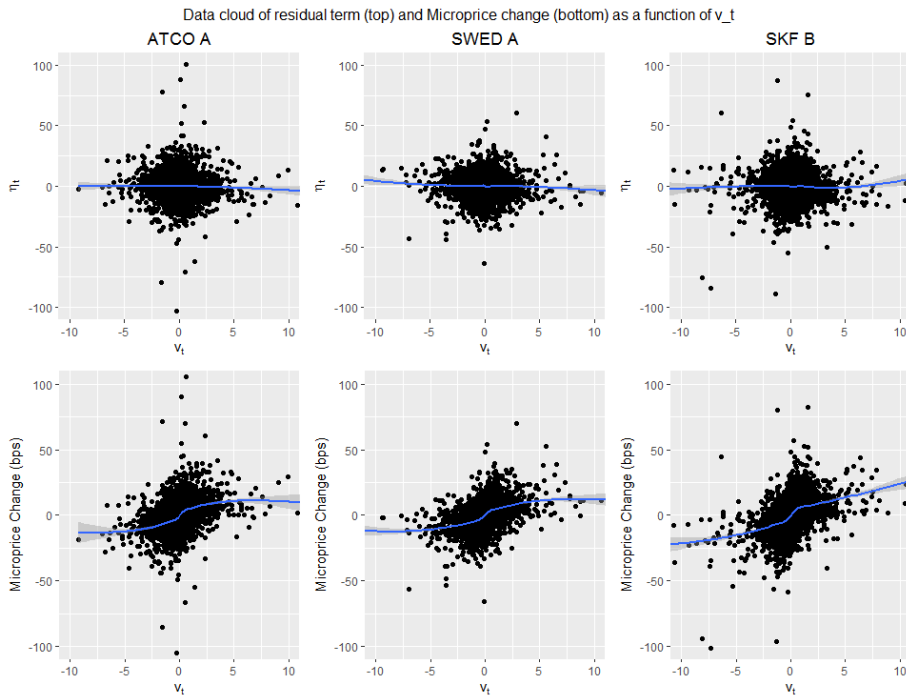


Figure 6.6: Data cloud of real returns vs Data cloud of η_t .

6.1.5 This Model vs. Previous Models

As mentioned in the beginning of this thesis, considerable research has been done on the subject of price impact. However, most of this research has been done from a trader's perspective and the objective has thus been to examine how much an individual trade affects the price of stocks. This research is important since traders, or more accurately trading algorithms, want to know how one

accumulates a certain amount of stocks at a certain time point to the lowest cost. When our model was estimated, it was done from a market perspective where we wanted to find to what extent, and how, aggregated buy-/sell-pressure on stocks affect their movement. These two ways of attacking the problem are indeed similar, since they both boil down to finding a model for price impact. However, since we look at aggregated sell-/buy-pressure we divided our data into minutely intervals, instead of looking at each individual trade. Therefore, the constraints mentioned in Section 3 about what impact functions are deemed reasonable and which decay function that are compatible might not apply for us. That being said, it is still interesting to make the comparison.

There are a lot of similarities between our model and previous ones, mainly on two key parts; the shape of the instantaneous impact function and the surprise in order flow-theory. The instantaneous impact function was in line with what previous researchers had discovered. As written in the Theory section, the instantaneous impact function was argued to be concave and odd, with the signed concavity best described by a power-law. These are all characteristics that we found to be true for our data. We did however see that, even though we fixed the power-law exponent to 0.5, this value did vary in the range 0.4-0.8 from stock to stock. We did also have a scale factor to the instantaneous impact function which had to be estimated for each separate stock. A lot of time was spent on understanding why this factor, the sensitivity to price impact, varied between different stocks. Some correlation was found between these parameters and the stock specific spread (averaged for the relevant time period), and in fact a rough approximation was obtained albeit not explanatory enough to give any further meaning to this thesis. All in all the instantaneous impact function we estimated was pretty much in line with what people had argued before us.

Secondly, we mentioned how the decay function was necessary in order to keep prices uncorrelated. Our estimated decay function was tuned in such way that it cancelled out the correlation in order sign, and it did also have negative decay coefficients which verified the "surprise in order book"-theory that was mentioned in the Gerig literature. That is, a surprise in the order book does in fact impact prices more.

When it came to the actual shape of the decay function there were more dissimilarities. In the Theory section, we mentioned that a non-linear instantaneous impact function should not be compatible with an exponential decay function since it would give rise to arbitrage opportunities. The decay function we found to fit our data the best was however exponential. Although, after evaluating the result and showing how little the decay function actually impacted the final model (from a residual point of view) we could probably have used a power law or a logarithmic decay function and achieved similar results. With respect to what has been mentioned, it is not surprising that the part of our model that depends on previous trades was the one that did not really agree with the theory presented. The time dependence is naturally weakened when the data is partitioned into minutely intervals rather than individual trades. For instance, it is argued that the autocorrelation of order signs comes from large orders being divided up in smaller pieces. This property is not captured as well when the data is partitioned into minutely intervals, which may both weaken the importance of the decay function as well as changing the shape of it.

6.2 The Effects of Price Impact

What was discovered when analyzing the instantaneous price impact function was that there was a concave relationship between aggregated volume and the micro-price changes. This was evident for all 29 stocks analyzed, as one can see in Appendix C.3. When modelling this concave shape, we chose to define our input variable v_t as the aggregated net bought divided by the minutely average turnover for each stock. This was done as we wanted to find a model as general as possible that varied as little as possible between different stocks. In Figure 6.7 we see data from all 29 stocks that has been divided into partitions depending on the size of the input (Net Bought) in order to calculate the mean Net Bought and mean Impact for those partitions. This method should be recognized from previous sections. One can see how the division of average minutely turnover (bottom plot) creates more similar trends between the different stocks. In the Method section we did not analyze what effect this (possible) dependency on turnover has, we only recognized its importance and moved on to fitting parameters. At this point in the analysis however, with a ready model for price impact, it is in order to see how said turnover affects the observed impact more thoroughly.



Figure 6.7: All 29 with data divided into partitions depending on the size of the input (Net Bought). In the top plot the input parameter is Net Bought while in the bottom plot it is Net Bought divided by average minutely turnover. Each color represent data from a certain stock. As one can see the differences between stocks are smaller in the bottom plot.

6.2.1 First Observations

An illustrative way to see the effects of an asset's turnover is to compare that asset on two different venues for the same time period, where turnover differs between the two. Starting with an example, the stock SEB A is traded on Nasdaq Stockholm as well as Chi-X (the pan-European venue explained in Section 4.1 & 2.2), specifically, 69.4% of the two venues' added turnover was traded at Nasdaq Stockholm in June 2016 (leaving 30.6% on Chi-X). As seen on the inclination of the bottom left plot in Figure 6.8, Chi-X appears to be considerably more sensitive to price impact than Nasdaq Stockholm. But how do we know if this effect is not just venue specific, and actually independent on the difference in turnover between the two venues? Would a 50/50 divide impose similar inclination in the two venues? Note that a steeper inclination implies more sensitivity in the sense that small differences in traded volume affect prices to a larger extent.

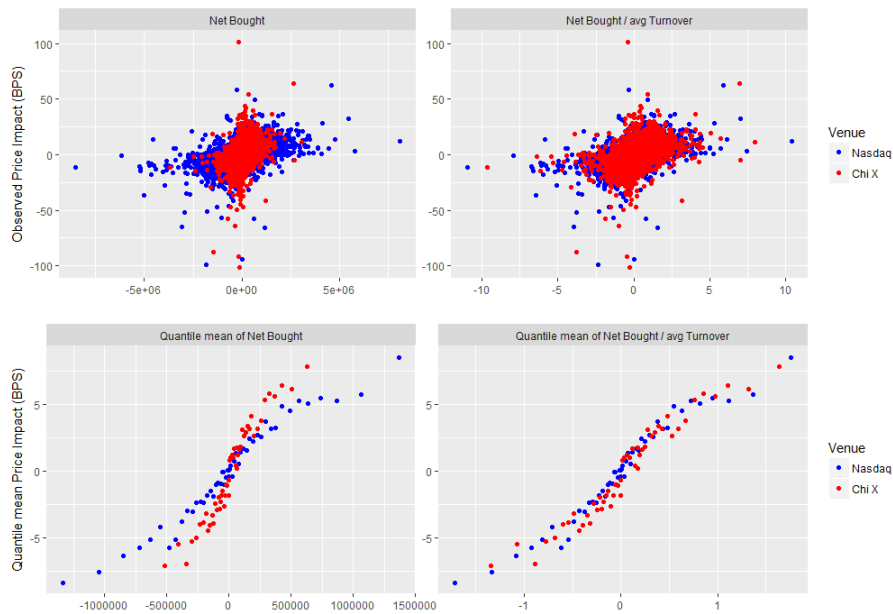


Figure 6.8: SEB A during June 2016; 69.4/30.6 in favor of Nasdaq Stockholm

Figure 6.8 depicts the original data cloud for one stock during one month in the top left corner, where Chi-X data (red) is plotted on top of Nasdaq Stockholm data (blue). Clearly, there is a difference between the two venues, although to say exactly what differs is difficult from this plot alone. In the bottom left corner, the same data has been partitioned as in Figure 6.7. In the plot in question, the difference between the two venues becomes much more visible; a steeper inclination is observed from the Chi-X data, hence leaving us with the conclusion that the stock prices on Chi-X in this case are affected more by trade than those on Nasdaq Stockholm. But as the input variable is divided by the average daily turnover (before partitioning and averaging the data) the difference diminishes, see bottom right plot in Figure 6.8. The conclusion is updated. Perhaps it is not Chi-X that is more susceptible to price change due

to trading than Nasdaq Stockholm per se, but instead it is an effect of how large the traded volumes are on the venue.

Since the alteration of input data is meant to simulate a scenario where equal turnover is observed on the two venues, we take a look at another example. In Figure 6.9, the stock SAND (Sandvik) is plotted for the month of January 2016, when the division of turnover happened to be near equal.

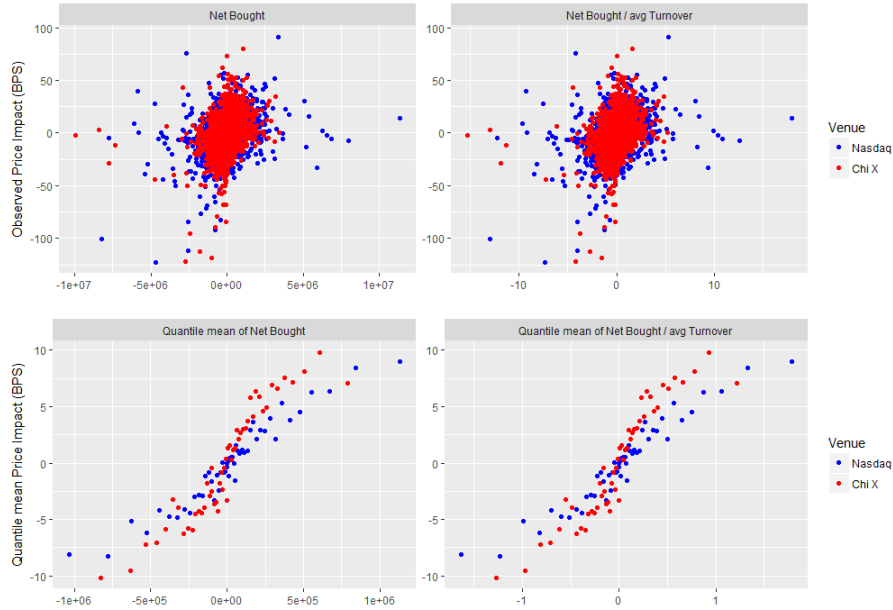


Figure 6.9: SAND during January 2016; 50.2/49.8 in favor of Nasdaq Stockholm

As expected, the dissimilarities between the left and right plots are negligible since we are dividing both the Nasdaq and Chi-X data with similar numbers. But what further strengthens our hypothesis that turnover is important in deciding the sensitivity of price impact is that the difference is small before the division of average minutely turnover.

6.2.2 Low Turnover Implies High Impact

We have now seen two examples that suggests a strong connection between magnitude of traded volume and price impact. However, before trying to analyze what effect this might have we include some more examples like the one we see in Figure 6.9. In Figure 6.10 one can see the mean data of six different stocks traded on Nasdaq and Chi-X before and after division by average minutely turnover. Again, the pattern is clear; in the original data the venue with smaller turnover (Chi-X) is the one with steeper inclination, and the difference is clearly decreased as the input variable (Net Bought) is divided by average minutely turnover. Before making a claim that this relationship holds we can recall the model used in previous sections to model the instantaneous impact function

$$f(v) = \alpha \cdot \text{sign}(v) \cdot \sqrt{|v|}, \quad \alpha > 0.$$

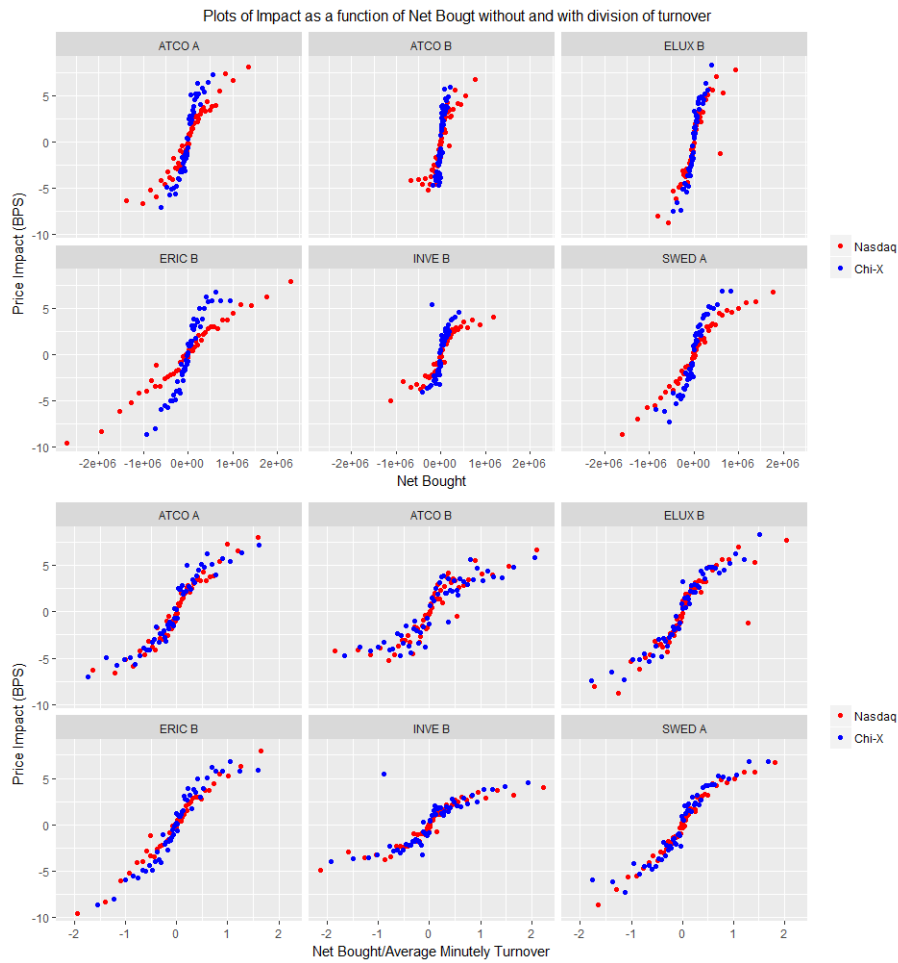


Figure 6.10: Impact in BPS as a function of Net Bought in the first six plots and Net Bought divided by average minutely turnover in the last six plots.

The value of α describes the inclination of the s-shaped curve seen in the previous plots. We can estimate two alphas for each stock; one for the Nasdaq and one for Chi-X. As the plots in Figure 6.10 indicate, all the stocks have a higher Chi-X alpha, say α_C , than Nasdaq alpha, α_N . More specifically, when looking at the data used to construct the plots in question (April 2016), the average ratio α_C/α_N is 1.58. That is, the value which describes the inclination of the instantaneous impact function for separate stocks is on average 58% higher on Chi-X compared to Nasdaq. However, when we instead estimate the α :s but with v defined as net bought divided by the stock's average turnover this ratio decreases to 1.05. In the table below we have done the same comparison for all months between January and June. The fraction $\alpha_C^{pct}/\alpha_N^{pct}$ is ratio of the α :s on Chi-X and Nasdaq when the input parameter is net bought divided by average turnover. As we can see from the values in Table 3, the difference in sensitivity to price impact is almost removed (except for January) when we model net bought relative to average turnover. We make the following

| Month | Jan | Feb | Mar | Apr | May | Jun |
|---------------------------------|------|------|------|------|------|------|
| α_C/α_N | 1.75 | 1.46 | 1.52 | 1.58 | 1.43 | 1.44 |
| $\alpha_C^{pct}/\alpha_N^{pct}$ | 1.37 | 1.05 | 1.11 | 1.05 | 0.96 | 0.98 |

Table 3: Fractions of the inclination parameters α_C and α_N , for Chi-X and Nasdaq respectively, compared to the same fraction where a division by average minutely turnover has been made (pct).

claim:

Trading venues with a higher turnover are less sensitive to price impact.

An implication of this claim is that on trading markets with higher turnover it is harder to manipulate prices by accumulating and liquidating stocks according to certain strategies.

6.2.3 Fragmentation of the Market

As described in Section 2.2, new market places for trading has emerged during the last decade. The shift in trade from a primary market to a secondary one is often referred to as fragmentation of the market, and naturally it comes at a cost for the primaries. In the case of Nasdaq Stockholm and Chi-X, continued cannibalization possibly lessens the total turnover of Nasdaq Stockholm, and increases that of Chi-X. As previously concluded, a large turnover has a stabilizing effect on price impact. We form the following hypothesis:

A fragmented market invokes sensitivity to price impact.

The hypothesis states that fragmentation of the stock market generates stock dynamics that are more sensitive to incoming trade. The previous section inspired this hypothesis as we saw how the inclination of the impact curve was affected by turnover. However, changes that occur over time might contribute to price impact sensitivity to the extent that decreasing turnover becomes irrelevant.

In order to analyze how fragmentation affects sensitivity for price impact we need to look at data where there has been a considerable shift in market share. As mentioned in the Theory section, the cannibalization of Nasdaq's market share by Chi-X was steadily increasing until around 2013 but has stagnated since. Unfortunately, no data from years prior to 2013 were available rendering the fragmenting dynamics of Chi-X and Nasdaq uninteresting for the purposes of this section. However, the smaller MTF Turquoise has been increasing in market share between 2013 and 2016. In Table 4 the average daily turnover (ADT) is presented for Nasdaq, Chi-X and Turquoise, along with the average market share for each venue. The data is taken from the first six months of 2013 and 2016. It should be added that the numbers and plots in this section come from the 26 stocks that the OMXS30 index has in common for 2013 and 2016.

As the market share of these three venues combined is almost 95% it is clear that Turquoise has taken market share from Nasdaq and Chi-X. The increase in market share has resulted in larger turnover, which according to our hypothesis, would result in less sensitivity to price impact, i.e. a flatter curve, compared to 2013. For Nasdaq and Chi-X the market share has decreased for all months,

| | | | Jan | Feb | Mar | Apr | May | Jun |
|-----------|------|-----------|---------------|---------------|---------------|---------------|---------------|---------------|
| Nasdaq | 2013 | ADT MS | 6.21 60% | 6.83 60% | 6.30 60% | 7.67 62% | 6.10 60% | 7.01 59% |
| | 2016 | ADT MS | 6.33 56% | 7.79 57% | 5.73 55% | 6.99 56% | 5.78 57% | 6.22 56% |
| Chi-X | 2013 | ADT MS | 2.68 29% | 2.99 29% | 2.63 29% | 2.78 28% | 2.41 29% | 3.01 30% |
| | 2016 | ADT MS | 4.00 26% | 3.91 25% | 2.68 24% | 2.60 24% | 2.14 24% | 2.54 25% |
| Turquoise | 2013 | ADT MS | 0.50 4.2% | 0.61 4.5% | 0.59 4.7% | 0.66 4.7% | 0.59 5.0% | 0.82 6.0% |
| | 2016 | ADT MS | 2.16 10.8% | 2.38 11.9% | 2.20 14.3% | 2.22 12.5% | 1.69 12.1% | 1.98 12.2% |

Table 4: Average Daily Turnover (ADT) and average Market Share (MS) of the three venues Nasdaq Stockholm, Chi-X and Turquoise for each of the six available months in 2013 and 2016. ADT values are in SEK billions. MS values are taken from bats.com [15][16][17].

while the turnover has decreased for some and increased for some, which should result in a steeper curve for some while a flatter for some. However, plots showed that the steepness of the price impact curve of Chi-X and Nasdaq had increased or remained stable for all months. We noticed that an increase in steepness always came with an increase in spread (i.e. the average of TWA_SPREAD explained in Appendix A), and for the months were the impact curve remained stable the difference in spread was negligible. In order to neutralize this potential spread-effect the y-axis in Figure 6.11, and in the following plots, is therefore divided by the average monthly spread.

In Figure 6.11 data from March and April is plotted and a linear fit is added in order to give a rough estimate of how the steepness of the curves has changed. From Table 4 one can see that between March 2013 and March 2016 there was a small decrease in turnover for Nasdaq while it increased very slightly for Chi-X. Despite this increase in turnover we see a small increase in the steepness of the impact curve for Chi-X. However, the increase is greater for Nasdaq, where there actually was a decrease in turnover. We further see that the impact curve for Turquoise has moved in the opposite way, maybe as a result of the significantly increased turnover. For April the inclination of the impact curve seems to have increased slightly for both Nasdaq and Chi-X, which is in line with our hypothesis since they both decreased in turnover slightly. Moreover, we can note that the decrease in steepness is obvious for Turquoise yet again, which strengthen the hypothesis that an increase (decrease) in turnover yields less (more) sensitivity to price impact also over time.

If we shift focus to the lines measuring the steepness we can in a more quantitative way study how change in turnover changes sensitivity to price impact over time. Of course the lines is not a very good fit to the data, but they do indicate the average inclination of the impact curve. We denote the inclination with λ . In Table 5 the change in inclination is presented and one can from these numbers together with the data in Table 4 conclude the following:

1. For Nasdaq turnover has increased for Jan and Feb and λ has increased

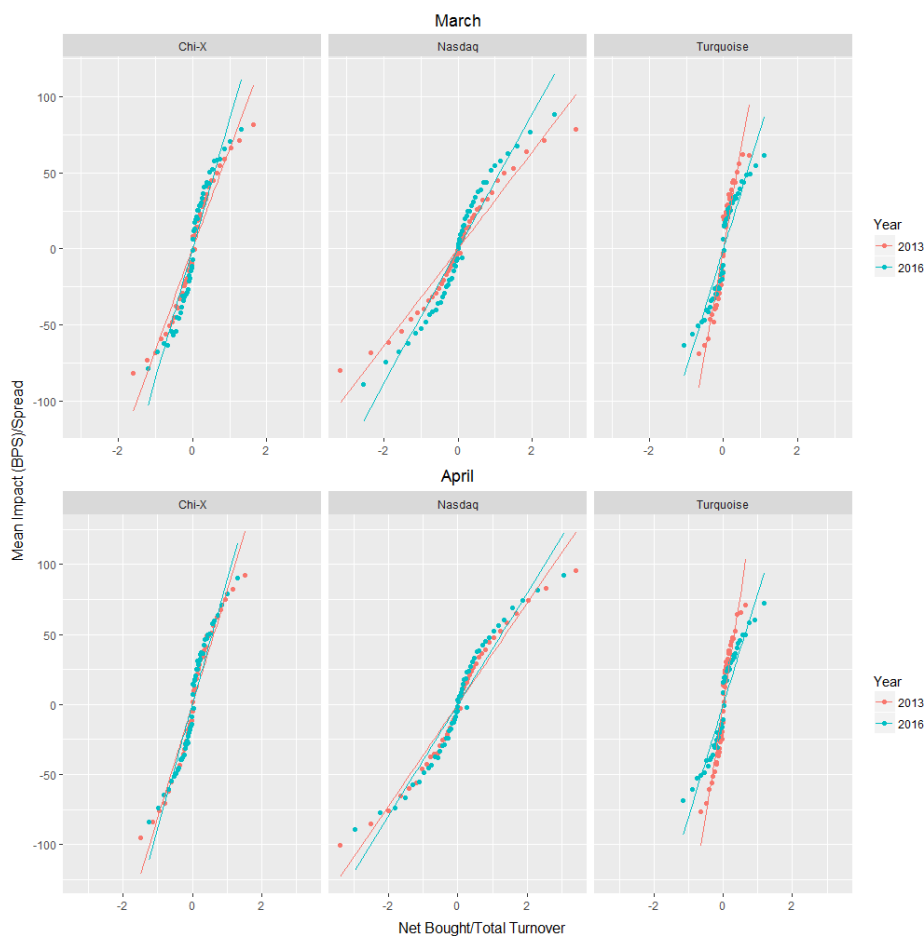


Figure 6.11: Plots of price impact for Chi-X, Nasdaq and Turquoise for March and April. The x-axis is divided by the total turnover over all venues in order to simplify comparisons, and in order to neutralize the effect of spread the y-axis is divided by spread. For all six months (January - June), see Figures C.7 & C.8 in Appendix C.6.

| | | Jan | Feb | Mar | Apr | May | Jun |
|------------------|--------|------|------|------|------|------|------|
| Nasdaq | 2013 | 28 | 33 | 32 | 36 | 37 | 39 |
| | 2016 | 44 | 48 | 44 | 40 | 46 | 46 |
| | Change | 57% | 43% | 39% | 10% | 24% | 19% |
| Chi-X | 2013 | 56 | 66 | 66 | 81 | 77 | 81 |
| | 2016 | 85 | 86 | 85 | 88 | 98 | 93 |
| | Change | 53% | 39% | 29% | 9% | 28% | 15% |
| Turquoise | 2013 | 118 | 144 | 138 | 158 | 145 | 145 |
| | 2016 | 99 | 91 | 79 | 80 | 94 | 92 |
| | Change | -17% | -37% | -43% | -49% | -37% | -37% |

Table 5: Average Daily Turnover (ADT) and average Market Share (MS) of the three venues Nasdaq Stockholm, Chi-X and Turquoise for each of the six available months in 2013 and 2016. ADT values are in SEK billions.

for all months.

2. For Chi-X turnover has increased for Jan, Feb and Mar and λ has increased for all months.
3. For Turquoise turnover has increased for all months and λ has decreased for all months.
4. All venues have the highest increase (least negative for TRQX) in λ in Jan and the lowest (most negative for TRQX) in Apr.

Based on these facts it seems as there is more than just change in turnover that affects sensitivity to price impact over time. More specifically, there seems to be factors pushing the sensitivity of price impact up from 2013 to 2016 based on the fact that inclination has increased for months where turnover has remained stable. How much the price impact sensitivity increases varies a lot from month to month and what causes this increase is outside the scope of this analysis. However, what is evident based on the figures shown is that the sensitivity has gone up for Nasdaq and Chi-X where the market share has remained fairly constant (or slightly decreasing), while for Turquoise the inclination of the price impact curve has decreased. One can thus argue that the naturally caused positive change to the inclination of the impact curve across all venues instead turned to a negative change for Turquoise as a result of the big increase in turnover. What is argued here is that if everything but turnover would have stayed constant, we would have seen a greater negative change in λ for Turquoise and a closer to zero change in λ for Chi-X and Nasdaq. A natural counter argument to this could be that the factors pushing the inclination up for Chi-X and Nasdaq could be the ones pushing it down for Turquoise, making the change in turnover irrelevant. However, based on the data for the months analyzed above there is a correlation of -0.814 between the change in λ and the change in turnover, which, combined with the plots and the data presented above, is a strong indicator that the hypothesis stated earlier should not be rejected. One could thus, based on the idea that a greater price impact sensitivity induces more price manipulation opportunities, question whether the rise of smaller trading venues and the fragmentation that comes with it really is a good thing for the stock market.

7 Conclusions

The correlation between an incoming order's nature and subsequent price change is no new field of financial studies, in fact considerable research has been made on the subject commonly known as market impact. A less exploited field is that of price impact, where all trades during a specific time period is considered. Throughout this thesis, various previously composed theories and models of market impact have been evaluated and refitted onto the financial data provided by Nasdaq Stockholm AB in order to explain the phenomenon of price impact thoroughly and hopefully in a comprehensible way.

A first look at the data suggested a signed concave model to explain price changes as a function of buy-/sell-pressure where, to begin with, no time dependency was present. The concave instantaneous model was equipped with the sign-function, leaving the sign of the output directly dependent of whether the bulk of the input consists of sell- or buy-orders. A.N. Gerig among others suggested that the order signs of historical trades are autocorrelated, and that autocorrelation must not be reflected in the model's output. Hence, after verifying the claim of the order signs' property, a time dependent part was added to account for the discrepancy. The added time dependency was supposed to serve two purposes, eliminate autocorrelation of output and account for the observed variability that was not covered by an instantaneous model. The effects of the time dependency did not cover as much of the variability as expected, but served its purpose well in the question of autocorrelation. Explicitly, the time dependence was added to the instantaneous model as a function of time lagged inputs, making the completed model dependent on several estimated parameters. Some variables could be generalized over time and assets, but the modelling part of this thesis is left with one stock-specific parameter that additional research might be able to explain in a more general way. What was concluded was that a square root shape best described the shape of price impact, adjusted by an exponentially decaying function of time lagged data.

When the performance of the model was tested onto observed data, under the assumption that the model covers what was intended, it became clear that a great part of price changes is independent of the actual buy-/sell-pressure. It was observed that, according to our model, price changes caused by price impact only explains around 10-35% of the variability in real price changes. Moreover, it was seen that the stock movements implied by price impact followed the real stock movements poorly in the month of January, and much better for other months. Suggesting that this could be explained by the amount of news released in January led to the conclusion that buy-/sell-pressure indicate stock movements poorly when a lot of news are released, and better in times of less news. The price impact model derived in the thesis was also compared to earlier models derived for market impact. The instantaneous part followed previous theory well while the decaying part did not. An exponential decay was found to fit the data the best despite this shape of decay function was argued against as it led to arbitrage opportunities. However, the reason the decay function differed from theory could be a consequence of the fact that we are comparing a model for aggregated trades to models for single trades, which most likely affects the time dependency of the model.

Moving on from the process of modelling, price impact was analyzed with respect to turnover. In order to keep as many stock market properties as possible

constant, the analysis was conducted over one month and stock at a time but observed on different trading venues where turnover differed. As it turned out, turnover has a clear effect on how sensitive an asset is to price impact, namely that less turnover implies more sensitivity. This "discovery" was further tested by another approach, by comparing the three major trading venues involved in OMXS30 trade.

Among the three largest venues that trade OMXS30 in Europe (Nasdaq Stockholm, Chi-X and Turquoise; in descending order in terms of turnover), it is evident that considerable market share has shifted from the two larger ones to Turquoise when looking at data from 2013 and 2016. This shift in market share comes hand in hand with a shift in turnover, thus leading to increased stability in terms of price impact at Turquoise, and a less stable environment at Nasdaq and Chi-X. The results invite the question of whether or not multiple trading platforms is in the best interest of the market or not.

References

- [1] R. Almgren and N. Chriss, *Optimal Execution of Portfolio Transactions*, Journal of Risk, 3:5-40 (2001).
- [2] B. Bland, *Analysts divided over trading 'threat' to LSE*, The Daily Telegraph (2007-05-29) [cited 2017-04-20]. Available at: <https://www.pressreader.com/uk/the-daily-telegraph-business.../20070529/textview>
- [3] J.P. Bouchaud, Y. Gefen, M. Potters and M. Wyart, *Fluctuations and response in financial markets: The subtle nature of random price changes*, Quantitative Finance, 4(2):176-190 (2004).
- [4] J.P. Bouchaud, *Price Impact*, Encyclopedia of Quantitative Finance (2010).
- [5] E. Fama, *Efficient Capital Markets: A Review of Theory and Empirical Work*, Journal of Finance, 25:383 (1970).
- [6] J. Gatheral, *No-dynamic-arbitrage and market impact*, Quantitative Finance, 10(7):749-759 (2010).
- [7] A.N. Gerig, *A Theory for Market Impact: How Order Flow Affects Stock Price*, PhD thesis, University of Illinois (2007). Available at arXiv:0804.3818.
- [8] G. Huberman and W. Stanzl, *Price manipulation and quasi-arbitrage*, Econometrica, 72(4):1247-1275 (2004).
- [9] A. Jansson, *MiFID – nya regler för värdepappersmarknaden* [Internet], Aktiespararna (2007-10-31) [cited 2017-05-03]. Available at: <http://www.aktiespararna.se/artikelarkiv/opinion/2007/oktober/mifid-nya-regler-for-vardepappersmarknaden/>
- [10] F. Lilo and J.D. Farmer, *The Long Memory of the Efficient Market*, Studies in Nonlinear Dynamics & Econometrics, 8 (2004).
- [11] C. Af Sandberg, R. Sevenius, *Börsrätt*, Studentlitteratur (2007-2008).
- [12] S. Åberg, *Valuation of Derivative Assets: Lecture Notes* (p.57), Mathematical Statistics, Lund Faculty of Engineering (2010).
- [13] S. Åberg, *Valuation of Derivative Assets: Lecture Notes* (p.98), Mathematical Statistics, Lund Faculty of Engineering (2010).
- [14] Bats website, *About Bats: History* [cited 2017-05-03]. Available at: http://www.bats.com/our_company/about/history/
- [15] Bats website, *Market Statistics* [downloaded 2017-05-10]. Available at: http://www.bats.com/europe/equities/market_statistics/venue.../nasdaqomx/index/OMXS/csv
- [16] Bats website, *Market Statistics* [downloaded 2017-05-10]. Available at: http://www.bats.com/europe/equities/market_statistics/venue.../batseurope/index/OMXS/csv

[17] Bats website, *Market Statistics* [downloaded 2017-05-10]. Available at:
[http://www.bats.com/europe/equities/market_statistics/venue...
.../turquoise/index/OMXS/csv](http://www.bats.com/europe/equities/market_statistics/venue.../turquoise/index/OMXS/csv)

A Variables

1. `CODE_CUR`. Code under which the stock is traded.
2. `PID`. Venue where data has been gathered. `BOOK` for Nasdaq Stockholm, `CHIX` for Chi-X, or `TRQX` for Turquoise.
3. `TIME_MINUTE`. Time of the day from 540 (9:00AM) until 1044 (5:24PM); market closes at 5:30PM.
4. `NET_BOUGHT`. Aggregated buying pressure on the market, i.e.

$$\text{bidprice} * \text{bidsize} - \text{askprice} * \text{asksize},$$

summed for all instants during that minute. The prices are the best offers at said instant and the sizes correspond to those prices. Measured in Swedish kronor (SEK).

5. `TURNOVER`. Total turnover, measured as `NET_BOUGHT` but each instant calculated as,

$$\text{bidprice} * \text{bidsize} + \text{askprice} * \text{asksize}.$$

6. `NET_BOUGHT_Vpct`. Calculated as `NET_BOUGHT` divided by the mean `TURNOVER` for this stock on this market. This is the variable used as the input of the models.
7. `MICROPRICE_CHANGE`. An order book weighted mid-point price, calculated according to,

$$\frac{\text{bidprice} * \text{asksize} + \text{askprice} * \text{bidsize}}{\text{bidsize} + \text{asksize}},$$

where the quantities have the same properties and relationship as in `NET_BOUGHT`. This is the output of the model, i.e. the observed price impact.

8. `TWA_SPREAD`. Time Weighted Average of the spread during relevant minute.

B Tables

B.1 List of Stocks in OMXS30

As of mid 2016 and end of 2016, the following stocks were included in OMXS30 (updated biannually).

| Company name | Stock code | GICS ¹ Sector | Weight ² |
|------------------------------------|------------|-----------------------------------|---------------------|
| ABB Ltd | ABB | Industrial Machinery | 2.36 |
| Alfa Laval | ALFA | Industrial Machinery | 1.65 |
| Assa Abloy B | ASSA B | Building Products | 3.55 |
| AstraZeneca | AZN | Pharmaceuticals | 3.13 |
| Atlas Copco A | ATCO A | Industrial Machinery | 5.24 |
| Atlas Copco B | ATCO B | Industrial Machinery | 1.67 |
| Boliden | BOL | Diversified Metals and Mining | 4.97 |
| Electrolux B | ELUX B | Household Appliances | 1.84 |
| Ericsson B | ERIC B | Communications Equipment | 7.04 |
| Fingerprint Cards B ^{3,4} | FING B | Diversified Support Services | 5.23 |
| Getinge B | GETI B | Health Care Equipment | 0.76 |
| Hennes & Mauritz B | HM B | Apparel Retail | 7.56 |
| Investor B | INVE B | Multi-Sector Holdings | 3.40 |
| Kinnevik B ⁴ | KINV B | Multi-Sector Holdings | 1.03 |
| Lundin Petroleum | LUPE | Oil and Gas Exploration and Prod. | 1.22 |
| Nokia Oyj ⁴ | NOKIA SEK | Communications Equipment | 0.94 |
| Nordea Bank | NDA SEK | Diversified Banks | 6.24 |
| Sandvik | SAND | Industrial Machinery | 4.36 |
| SCA B | SCA B | Paper Products | 3.69 |
| SEB A | SEB A | Diversified Banks | 4.72 |
| Securitas B | SECU B | Professional Services | 1.83 |
| Skanska B | SKA B | Construction and Engineering | 2.06 |
| SKF B | SKF B | Industrial Machinery | 3.22 |
| SSAB A | SSAB A | Steel | 2.16 |
| Svenska Handelsbanken A | SHB A | Diversified Banks | 3.38 |
| Swedbank A | SWED A | Diversified Banks | 4.75 |
| Swedish Match | SWMA | Tobacco | 1.47 |
| Tele2 B | TEL2 B | Integrated Telecom Services | 1.60 |
| Telia Company ⁴ | TELIA | Integrated Telecom Services | 3.05 |
| Volvo B | VOLV B | Heavy Trucks | 5.90 |

¹Global Industry Classification Standard.

²Contribution to OMXS30 index (%) in terms of turnover by end of November 2016.

³Excluded stock due to irregular behavior.

⁴Not part of OMXS30 during 2013, thus excluded from comparisons in Section 6.2.3.

C Plots

C.1 Market Share of Different Venues

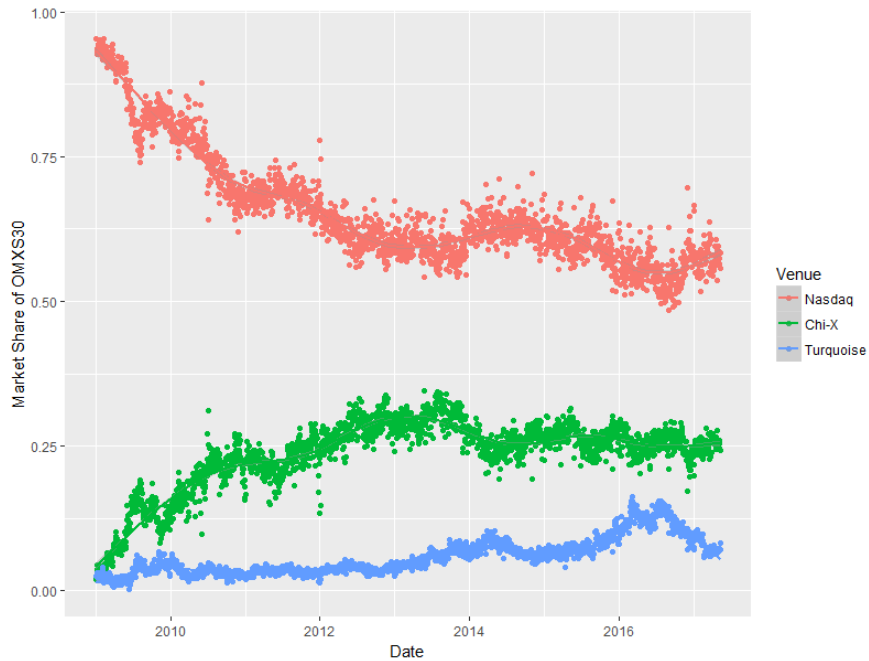


Figure C.1: Share of total traded volume of the stocks comprising OMXS30 for Nasdaq Stockholm, Chi-X and Turquoise. Clearly, the primary trading venue has lost a lot of market shares since the introduction of alternative venues. It does look like the secondary venues, and Chi-X especially, have stabilized around a market share percentage. The data was downloaded from `bats.com` [15][16][17].

C.2 Visualization of all Stocks




Figure C.2: A visualization of all stocks in OMXS30.

Axes: observed price impact = micro-price change
v = aggregated volume \div individual minutely average turnover

C.3 Visualization of Quantile Means for all Stocks




Figure C.3: A visualization of all stocks in OMXS30, altered according to the mean-scheme described in Section 4.2.

C.4 Estimated β :s Using One and Six Months of Data for all Stocks

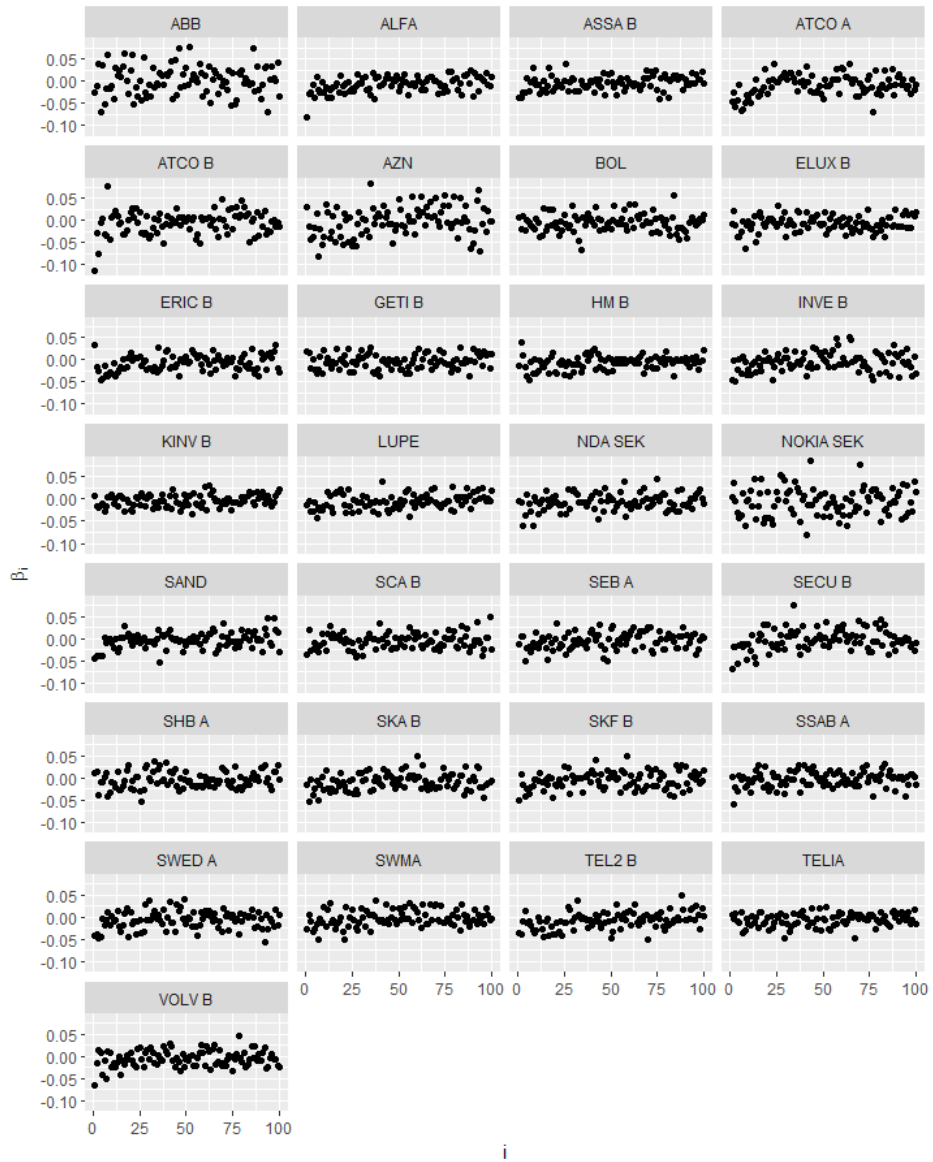


Figure C.4: Estimated β :s for all 29 stocks using only one month of data.

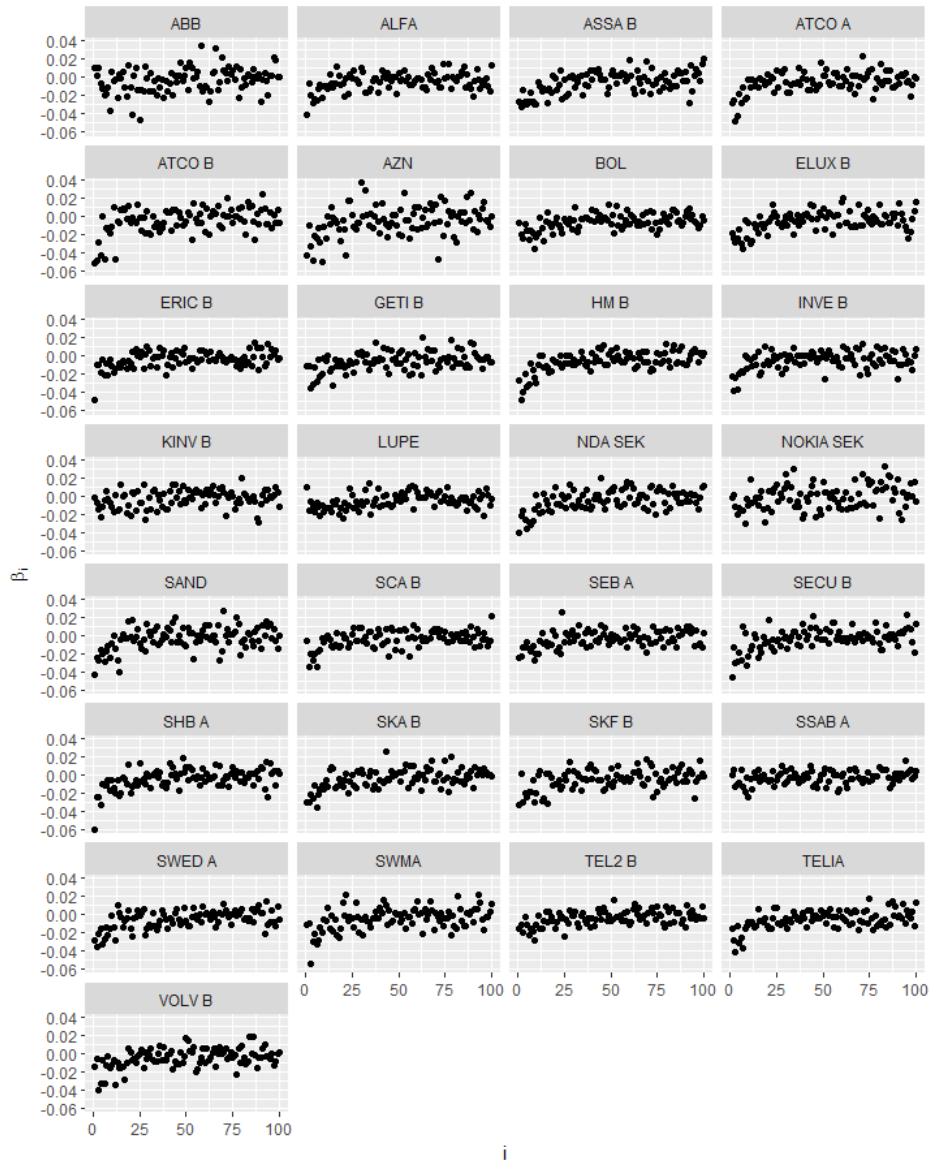


Figure C.5: Estimated β :s for all 29 stocks using six months of data. Note that the increasing structure is much more clear than in Figure C.4.

C.5 Fitted $G(\cdot)$ over $\hat{\beta}$:s for all Stocks

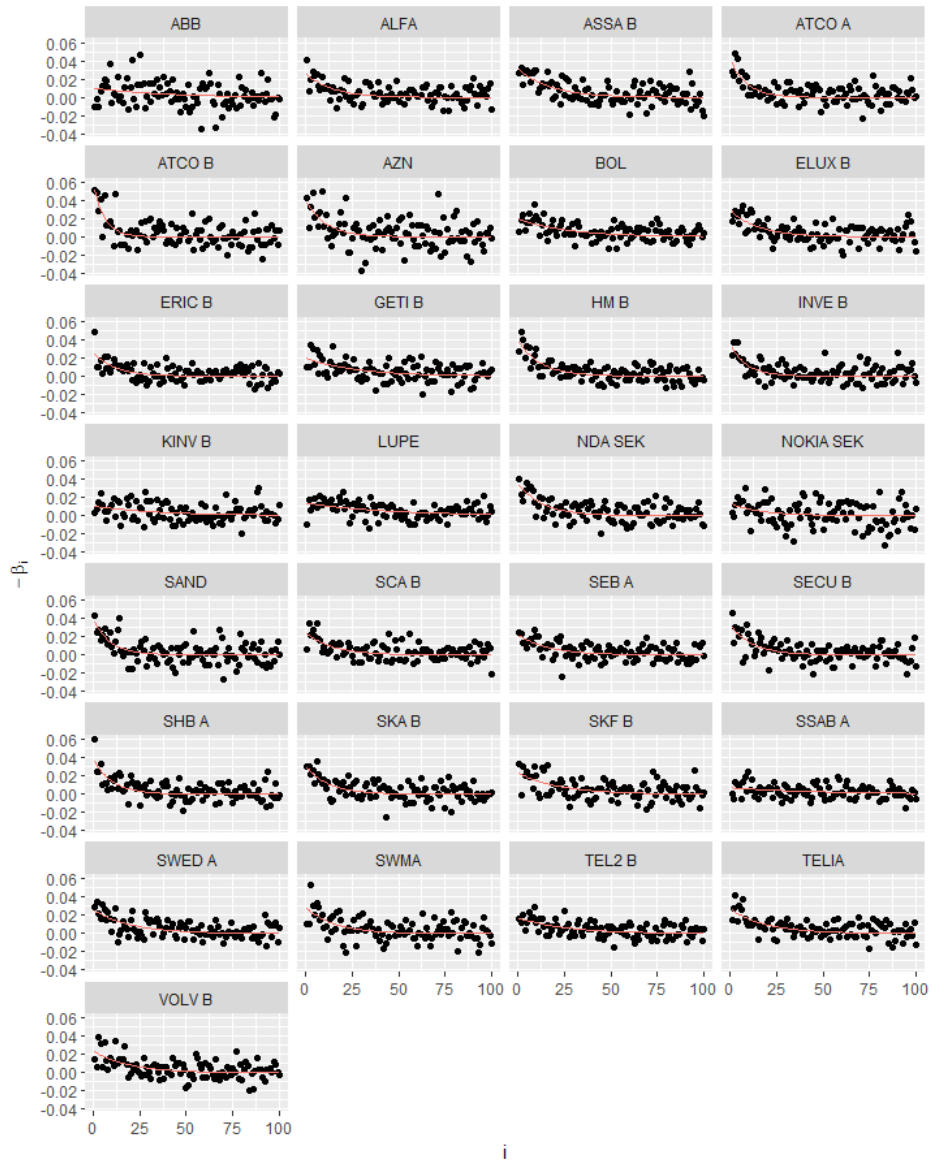


Figure C.6: Fitted $G(\cdot)$ over $\hat{\beta}$:s for all Stocks.

C.6 Plots of Price Impact for Chi-X, Nasdaq and Turquoise

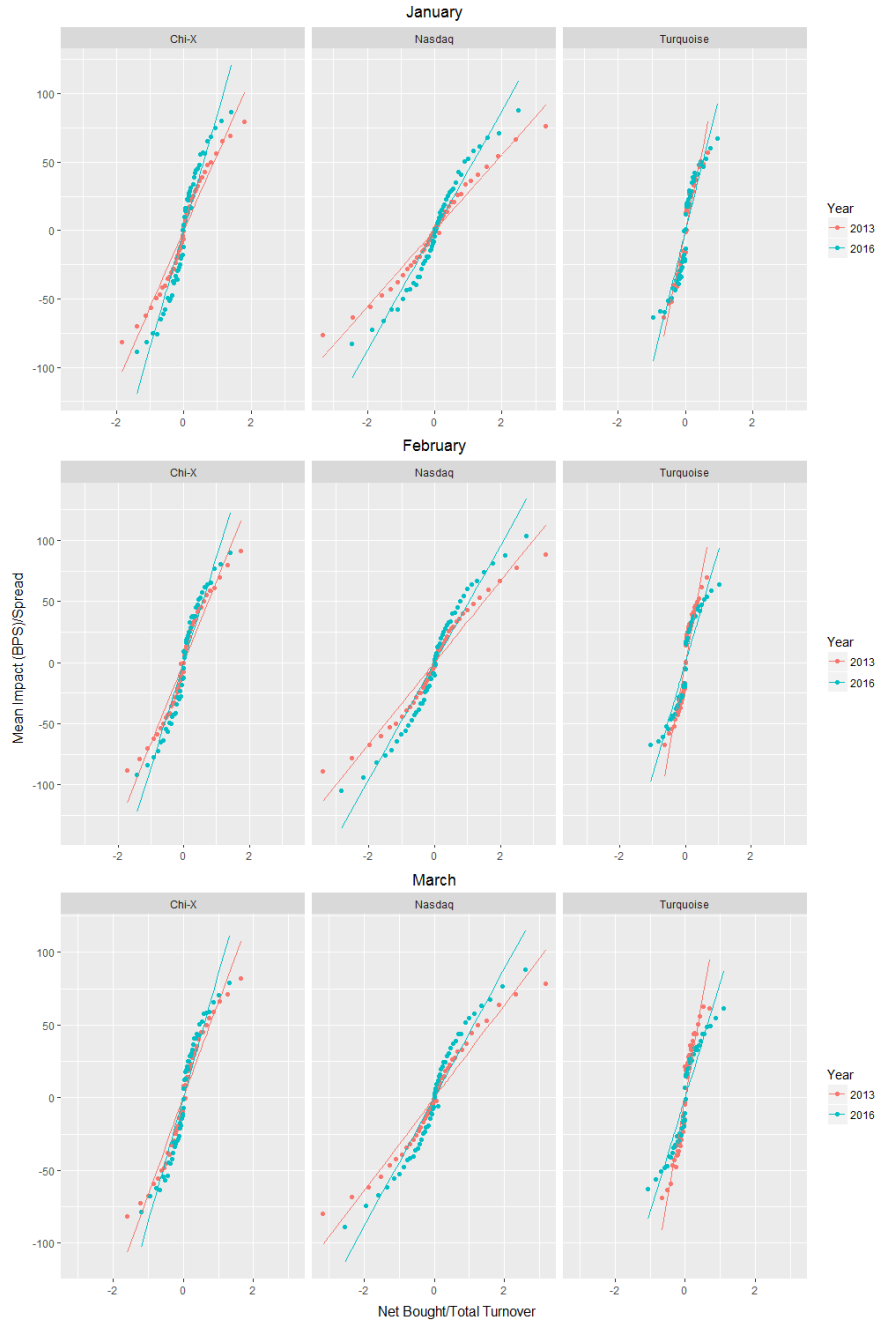


Figure C.7: Monthly plots of price impact for Chi-X, Nasdaq and Turquoise for the first quarter. The x-axis is divided by the total turnover over all venues in order to simplify comparisons, and in order to neutralize the effect of spread the y-axis is divided by spread.

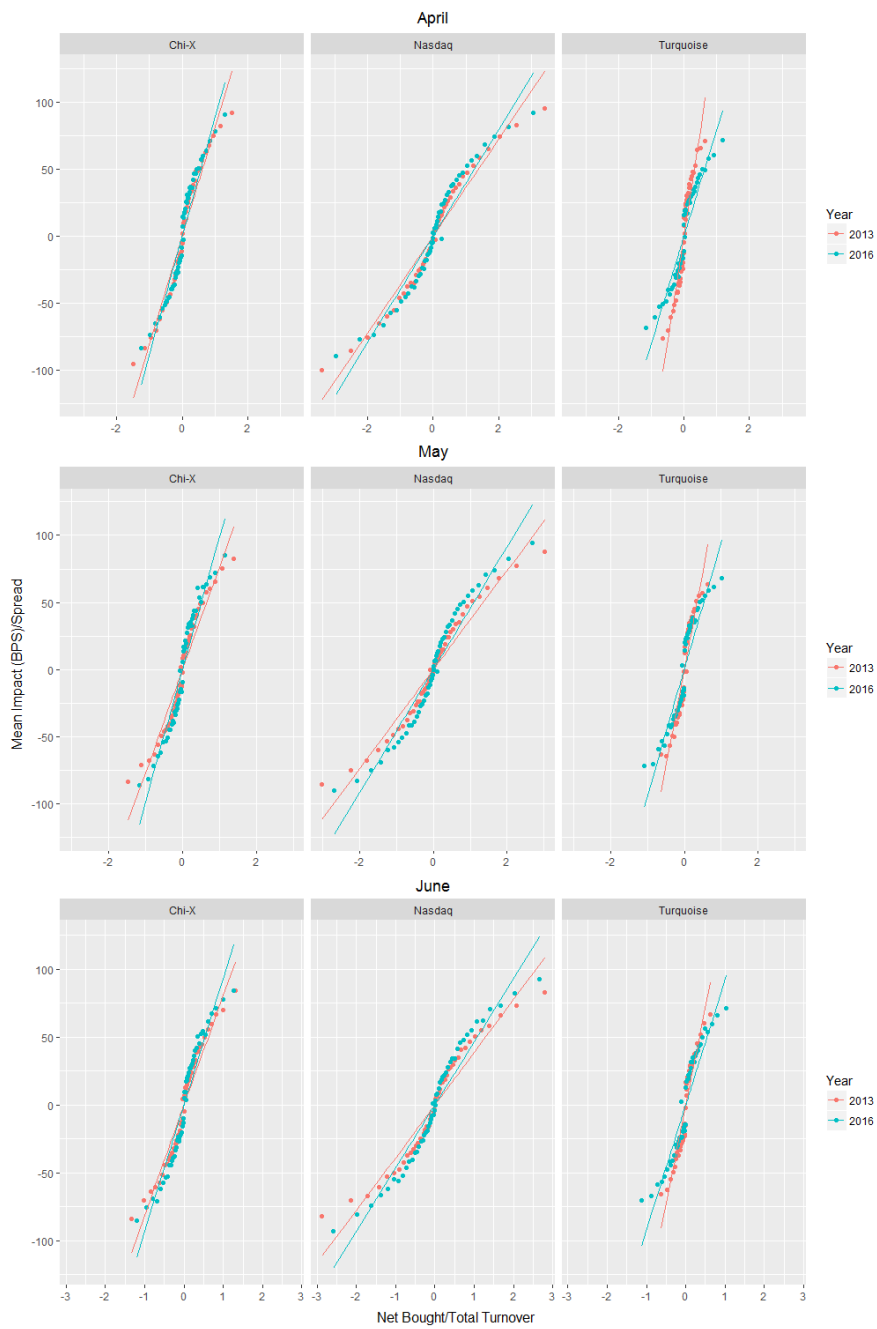


Figure C.8: Monthly plots of price impact for Chi-X, Nasdaq and Turquoise for the second quarter. The x-axis is divided by the total turnover over all venues in order to simplify comparisons, and in order to neutralize the effect of spread the y-axis is divided by spread.