Liquidation Strategies in a Long-Short Equity Portfolio

Patrik Nilsson

Department of Mathematics,
KTH, Stockholm,
October 2011

Abstract

Trading large volumes impact the price of the traded asset which implies a cost when the position is liquidated. Because of this large investors, such as hedge funds, need estimates of the expected market impact of their positions. I suggest a model for the market impact of trading and use this model to analyze and compare different liquidation strategies. I specifically consider liquidating large fractions of a long-short equity portfolio. I consider two common liquidation strategies and compare these to another strategy I introduce in this thesis; optimized liquidation which is the solution to an optimization problem. The results show that it is possible to reduce expected market impact costs from liquidation while keeping the remaining portfolio within pre-specified risk limits.

Acknowledgements

I want to thank my supervisor Filip Lindskog for guidance and support. I also want to thank Sam Nylander and Alexander Argiriou for their constant feedback and for making my time a Brummer & Partners one of the most interesting and fun parts of my studies at KTH.

Contents

1 Introduction

Financial assets are similar to any other traded good in that prices are determined by supply and demand. At any given price, there are only a limited number of potential buyers and sellers and thus a limited volume that can be traded at that price. If an investor wants to buy (or sell) a larger volume he will only be able to trade a fraction of that volume at the original market price and as he proceeds to buy (or sell) the rest of his position he will have to accept less favorable prices in order to attract counterparties.

The way trading affects the price of the traded asset is often referred to as the *market impact of trading*. Market impact is generally not relevant for small investors. However, for investors that trade in larger volumes market impact implies an additional cost of trading and needs to be considered.

One of the larger investors in the market is the hedge fund industry. A hedge fund is a collective investment scheme that gets its capital from other investors, both private investors and institutions. The aggregate capital allows the fund to invest aggressively in positions they consider profitable. Although buying and selling large positions always require consideration about market impact, it is especially relevant when the fund for some reason needs to liquidate a large fraction of its positions within a relatively short time frame. A typical example of this is when one of the fund's large investors withdraws his money. With a broad variety of investment strategies, deciding which positions to liquidate is not trivial. On one hand, the portfolio manager wants to sell the most liquid positions to avoid the cost of market impact. On the other hand, if only the most liquid positions are sold, then this will often change the properties of the portfolio in an unfavorable way for the remaining investors.

To balance these opposites, a portfolio manager needs to have a *liquidation strategy* ready for when this situation arises. This is not only important from a managing point of view. To keep and attract investors, a fund needs to be able to provide information about what costs investors will face when they withdraw their capital.

In this thesis I suggest a model for the market impact of trading and use this model to compare different strategies for liquidating a large fraction of a portfolio. Specifically, I consider long-short equity portfolios.

1.1 Outline

- In Section 2 I introduce definitions relevant for this thesis.
- In Section 3 I suggest a model for the market impact of trading by presenting results from a large number of empirical studies.
- In Section 4 I show what this model implies for the cost of trading.
- In Section 5 I formalize the liquidation problem that was briefly mentioned in the introduction.

- In Section 6 I present two case studies where I implement the models that have been introduced.
- In Section 7 I present results.
- In Section 8 I conclude.

2 Definitions

In this section I define concepts that will be important in this thesis.

2.1 General framework

The investment universe available to a hedge fund is vast. I limit this universe by assuming the fund can only invest in stocks and in a money asset which is perfectly liquid. The money asset is supposed to reflect any positions where the hedge fund stores capital rather than invest it. Typically, the risky positions in the portfolio are active investment decisions whereas some capital is kept in assets whose expected returns are secondary to them being as liquid and risk free as possible. Because of this, I assume the interest of the money asset is zero.

Consider a portfolio that have positions in n stocks. The number of shares held in each stock is given by the weights

$$\mathbf{h}(t) = (h_1(t), h_2(t), ..., h_n(t))$$

The weights can be negative in which case they are called *short* positions (positive weights are called *long* positions). The spot prices of the stocks are

$$\mathbf{S}(t) = (S_1(t), S_2(t), ..., S_n(t))$$

The portfolio further consists of a position $h_0(t)$ in a perfectly liquid money asset $S_0(t)$ with zero interest rate, i.e. $S_0(t) = 1$ for all t.

The mark-to-market value of this portfolio at time t is:

$$V(t) = h_0(t) + \sum_{i=1}^{n} h_i(t) S_i(t)$$

Furthermore, the net asset value (NAV) of the portfolio is equal to its mark-to-market value. The NAV is an important concept that will be used frequently in this thesis.

2.2 Risk measures

At Brummer $\&$ Partners, risk is calculated in the standard way in units of some currency. However, it is often presented as a fraction of the NAV of the portfolio considered. Below I will give formal definitions of the various risk measures to be used in this thesis but in later sections I will discuss risk in relative rather than absolute terms. To emphasize this difference I include two simple examples below.

2.3 Net Exposure

The net exposure of an equity portfolio measures any long or short bias of the portfolio. It only includes the exposure of the stock positions, any risk free position is excluded.

Definition 1. *Consider a portfolio $V_p(t)$ consisting of the weights **h** (t) = $\big(h_1(t), ..., h_n(t)\big)$ in stocks with spot prices **S** (t) = $\big(S_1(t), ..., S_n(t)\big)$ and the weight $h_0(t)$ in the risk free asset $S_0(t) = 1,\ \forall t$. The net exposure of the portfolio is defined as*

$$\textit{Net Exposure } [V_p(t)] = \sum_{i=1}^{n} h_i(t)S_i(t)$$

Example 1. *Consider two stocks, $S_1(t)$ and $S_2(t)$, both with spot prices 100 SEK. Furthermore, consider a portfolio $V_p(t)$ consisting of a **short** position of 100 shares in $S_1(t)$, a **long** position of 200 shares in $S_2(t)$ and 10000 SEK in the risk free asset $h_0(t)$. The absolute net exposure of this portfolio is*

$$\textit{Abs. Net Exposure } [V_p(t)] = \sum_{i=1}^{n} h_i(t)S_i(t) = -100 \cdot 100 + 200 \cdot 100 = 10000$$

The NAV of the portfolio is given by

$$V(t) = h_0(t) + \sum_{i=1}^{n} h_i(t)S_i(t) = 20000$$

The relative net exposure (relative the NAV) is given by

$$\textit{Rel. Net Exposure } [V_p(t)] = \frac{\sum_{i=1}^{n} h_i(t)S_i(t)}{h_0(t) + \sum_{i=1}^{n} h_i(t)S_i(t)} = 0.5$$

If the portfolio did not hold any capital in the risk free asset (i.e. $h_0(t) = 0$), then the NAV would be equal to 10000 and relative net exposure of the portfolio would be

$$\textit{Rel. Net Exposure } [V_p(t)] = \frac{\sum_{i=1}^{n} h_i(t)S_i(t)}{h_0(t) + \sum_{i=1}^{n} h_i(t)S_i(t)} = 1$$

2.4 Gross exposure

The gross exposure is a measure of the leverage of the portfolio. It is defined in a way similiar to the net exposure.

Definition 2. *Consider a portfolio $V_p(t)$ consisting of the weights **h** (t) = $\big(h_1(t), ..., h_n(t)\big)$ in stocks with spot prices **S** (t) = $\big(S_1(t), ..., S_n(t)\big)$ and the weight $h_0(t)$ in the risk free asset $S_0(t) = 1,\ \forall t$. The gross exposure of the portfolio is defined as*

$$Gross\ Exposure[V_p(t)] = \sum_{i=1}^{n} |h_i(t)|S_i(t)$$

Example 2. *Consider the same portfolio as in Example 1. The absolute gross exposure of this portfolio is*

$$\textit{Abs. Gross Exposure } [V_p(t)] = \sum_{i=1}^{n} |h_i(t)|S_i(t) = 100 \cdot 100 + 200 \cdot 100 = 30000$$

The NAV of the portfolio is the same as in Example 1

$$V(t) = h_0(t) + \sum_{i=1}^{n} h_i(t) S_i(t) = 20000$$

The relative gross exposure (relative the NAV) is given by

$$\textit{Rel. Gross Exposure } [V_p(t)] = \frac{\sum_{i=1}^{n} |h_i(t)| S_i(t)}{h_0(t) + \sum_{i=1}^{n} h_i(t) S_i(t)} = 1.5$$

If the portfolio did not hold any capital in the risk free asset (i.e. $h_0(t) = 0$), the NAV would be equal to 10000 and the relative gross exposure of the portfolio would be

$$\textit{Rel. Gross Exposure } [V_p(t)] = \frac{\sum_{i=1}^{n} h_i(t) S_i(t)}{h_0(t) + \sum_{i=1}^{n} h_i(t) S_i(t)} = 3$$

2.5 Value-at-Risk

Value-at-Risk is the value such that the probability of experiencing a loss larger than this value (over a given time horizon) is the given probability level. Many definitions include a loss variable L defined as

$$L^{\Delta} = -\big(V(t+\Delta) - V(t)\big)$$

where Δ is the time horizon of interest. Below is a (slightly modified) definition from [1].

Definition 3. *Given some confidence level $\alpha \in (0, 1)$ and time horizon Δ, the Value-at-Risk of a portfolio $V_p(t)$ at a confidence level α is given by the smallest number l such that the probability that the loss L exceeds l is not larger than (1 - α). Formally,*

$$\textit{VaR}_{\alpha,\Delta} = \inf\{l \in \mathbb{R} : P(L^{\Delta} > l) \leq 1 - \alpha\}$$

There are a vast number of ways to estimate the VaR of a portfolio. A common approach is to assume that the loss variable L^{Δ} is normally distributed. This assumption does not agree with empirical findings (the lack of heavy tails in the normal distribution underestimates large price fluctuations). Nonetheless, this assumption is frequently used in the industry. The reason for this is that it is a reasonable approximation and that it significantly simplifies computations. In practice, the Value-at-Risk of a position would be calculated in a number of different ways to address this problem. The normal assumption leads to a very simple expression for the Value-at-Risk which is presented in the following corollary.

Corollary 1. *If the loss variable L above is normally distributed, $L \in N(\mu_L, \sigma_L^2)$, then the Value-at-Risk is given by*

$$\textit{VaR}_{\alpha,\Delta} = \mu_L + \sigma_L \Phi(1 - \alpha)$$

where Φ is the standard normal cumulative distribution function.

Proof 1. *The proof is left for the appendix.*

2.6 Beta

Beta is a measure of the relation between the log returns of a given asset (or portfolio) and the overall stock market. It can be described as a volatility adjusted correlation. In practice, beta is typically measured against a relevant index rather than against all existing stocks.

Definition 4. *Consider log returns R_p of a portfolio $V_p(t)$ and log returns R_b of the overall stock market (or possibly an index). The beta of the portfolio V_p is defined as*

$$\beta_p = \frac{Cov(R_p, R_b)}{Var(R_b)}$$

2.7 Tracking Error

Tracking Error (TE) is a measure of how closely one portfolio "tracks" another portfolio. It is most often used as an ex post measure of the performance of some portfolio relative a benchmark portfolio, such as an index. In that case it compares a series of realized returns over some time interval. In this thesis I will instead consider the less used ex ante measure; the *expected* tracking error.

Definition 5. *Consider a portfolio $V_p(t)$ and a benchmark portfolio $V_b(t)$ with (random) log returns $R_p(\Delta)$ and $R_b(\Delta)$ respectively over the time interval $(t, t+\Delta)$. Furthermore, assume the marginal distribution functions of the stochastic variables $R_p(\Delta)$ and $R_b(\Delta)$ are $f_b(r)$ and $f_p(r)$ respectively. The expected tracking error of the portfolio $V_p(t)$ relative the portfolio $V_b(t)$ is given by*

$$TE = \sqrt{\int_{\mathbb{R}} \left(f_b(r) - f_p(r)\right)^2 \, dr}$$

3 The market impact function

Consider a trader liquidating a large position during a relatively short time interval (say hours or days). It can be argued that during this short time interval, the long term drift of the asset due to fundamentals only has a second order effect on price movements. Instead, the price movements are largely due to the change of supply and demand in the market resulting from the large order. For notational clarity I assume, without loss of generality, that the trader starts liquidating his positions at time 0 and I propose the following price process for the asset

$$S(t) = S(0)\Big(1 + \epsilon f(x(t); \Omega) + \xi\Big) \tag{1}$$

Here, S is the (mid) price of the asset a time t, ϵ denotes the sign of the trade (positive for buy trades and negative for sell trades[1]) and $f(x(t); \Omega)$ is a (not yet specified) function describing the market impact of trading. The function depends on the volume x that has been transacted up until time t and also a set of other variables, for now denoted by Ω. With this notation it is obvious that the market impact function is dimensionless. The last term, ξ is an error term with zero mean. It can be interpreted as reflecting the rest of the market activity during trading.

Trivially,

$$E\big[S(t)\big] = S(0)\Big(1 + \epsilon f(x(t); \Omega)\Big)$$

The intuitive way of deriving the market impact function is by performing statistical tests on empirical data. However, this approach is problematic. To perform such a study, detailed data about a significant number of large orders is needed. The funds at Brummer & Partners refrain from trading very large orders exactly because of the market impact that I intend to model. Fortunately, there exist a number of empirical studies where researchers have been given access to large data sets of large orders. In this thesis I will review a number of these studies and combine their results with intuition and theoretical arguments to derive a model for market impact.

3.1 Differences between assets

Speaking of *the* market impact function is misleading. Different assets have completely different properties and likely impact the market in different ways. Some assets are standardized and traded on centralized exchanges whereas other assets are traded off-exchange in private agreements (often referred to as over-the-counter assets). Finding market impact functions for all different assets would obviously be preferable but

[1]It is intuitive that a large sell (buy) order will increase the supply (demand) of a stock in the market and thus decrease (increase) prices. However, since every trade has a buy and a sell side it is not obvious what is meant by "sell orders" and "buy orders". It is common practice to make the distinction based on which side that initiates the trade. The side that initiates the trade is called liquidity taker whereas the passive side is called liquidity provider. Moro et al. [10] compares market impact between orders that are executed by liquidity takers and liquidity providers. They find that both strategies give a positive market impact but impact is larger for the liquidity taker. The major difference between these two strategies is that only liquidity taking guarantees that the order is actually executed. Thus, large orders are typically executed by taking liquidity.

the large number of assets makes this impossible.

This thesis will focus on the impact of trading stocks, which is the asset that has, by far, been studied the most in academic papers. However, even for stocks there is not necessarily a universal impact function. Stocks can be traded in different ways and this might affect the impact of the trade.

The advances of computer technology during the last decade have allowed a significant growth of the electronic financial market. Today, most stock exchanges keep an electronic trading platform where agents can trade in a limit order book (LOB). In a LOB any agent in the market is allowed to place limit orders that offer to buy (or sell) a certain number of shares of a stock at a maximum (or minimum) price. The highest buy offer is called the best bid price and the lowest sell offer is called the best ask price. Any agent in the market is also allowed to place market orders that execute immediately against the best price in the LOB. If the volume of the market order is larger than the volume offered at the best price, the remaining volume will execute against a less favorable price and this will lead to a price change of the stock. Agents placing limit orders are often referred to as liquidity providers whereas agents placing market orders are referred to as liquidity takers.

If an order is small enough it can be executed in one transaction at the best price. However, the available volume at the best prices is typically small. A quick look at the Nasdaq Nordic Stock Exchange shows that it is typically less than a percent of the daily volume for large cap stocks and a few percent of the daily volume for mid and small cap stocks. To reduce market impact it has become common practice to split orders (even of modest sizes) and execute them incrementally over time.

Splitting orders introduces a time dependence on the impact of trading. Intuitively, if an order is executed slowly it allows liquidity providers to enter the market during the execution. This increases the supply (or demand) of the asset being bought (or sold). Thus, a more patient trading strategy is expected to decrease the impact of the order. However, slow execution has a disadvantage. Holding the position over a longer time interval implies a larger market risk. Because of this, there is a tradeoff between the expected cost of liquidation and the market risk of trading slowly.

Besides trading in the LOB, agents can also execute their positions outside the electronic market in a privately negotiated transaction, often referred to as a *block trade*. This is typically done for very large orders but can also be done if the initiator of the trade wants to avoid the market risk associated with order splitting.

There are obviously other ways stocks can be traded but this generalization captures the great majority of trades. Since stocks can be traded in different ways, empirical studies also differ in what they measure, the variables they consider and the results they get. However, there is a consensus about the dependence on the traded volume. More or less all studies show that impact is a concave function of the traded volume. A concave function has a negative second derivative. The functions suggested in the studies also have a positive first derivative. These two properties imply that the impact is an increasing function of the traded volume but that it increases slower and slower. This is true for individual orders, split orders and block trades. An example of such a function is plotted in Figure 1. Some studies argue that market impact is a logarithmic

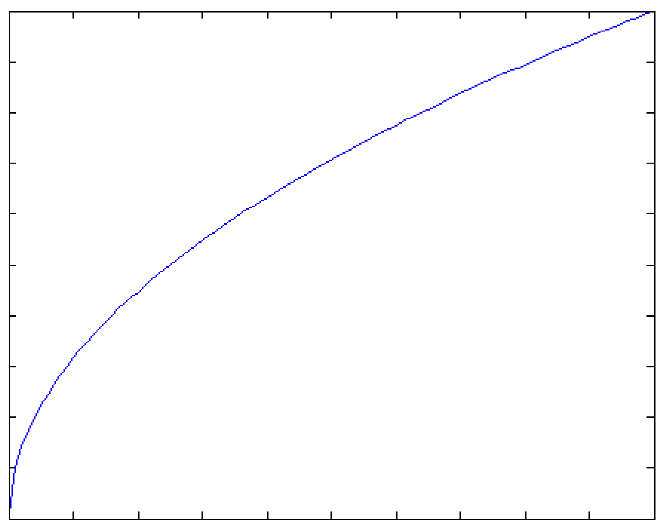

Figure 1: Graph of a concave function with positive first derivative

function but the majority propose a power law

$$f(x(t); \Omega) \sim x(t)^{\varphi} \tag{2}$$

with the notation introduced above. As mentioned $x(t)$ is the volume executed up until time t. The function is concave if the exponent φ is between zero and one.

3.2 Empirical studies

The majority of the studies about market impact considers individual trades, split orders or block trades. In this section I will summarize some of the studies and fit the results to equation (2). As mentioned, all studies have different approaches to the subject and fitting all results to one equation is naturally an approximation.

3.2.1 Individual trades

Hausman, Lo and MacKinlay [3] consider six stocks from the New York and American Stock Exchanges from January 4 to December 30 of 1988. The data set consists of time-stamped trades, trade size and bid/ask quotes. They find a concave impact function for all stocks but do not try to fit the results to a function.

Lillo, Farmer and Mantegna [4] used the Trade and Quote (TAQ) database to analyze roughly 113 million transactions and 173 million quotes from the 1000 largest companies on the New York Stock Exchange between 1995 and 1998. They find $\varphi \approx 0.5$ for smaller volumes and 0.2 for large volumes.

Farmer, Patelli and Zovko [5] use a data set from the London Stock Exchange between August 1998 and April 2000, which includes a total of 434 trading days and roughly six million events. They find $\varphi \approx 0.3$.

Hopman [6] considers a data set from the Paris Bourse between January 1995 to October 1999. The data set contains all the order submitted to the exchange and all the best quotes available at any time. He considers 30 minute intervals and finds $\varphi \approx 0.4$ depending on the urgency of the order.

Potters and Bouchaud [7] refer to the study by Lillo, Farmer and Mantegna [4] and argue that the lower value of φ for larger volumes indicate that a logarithmic function might fit the impact function better. They analyse three stocks from the Paris Bourse and fit the results to a logarithmic function.

3.2.2 Split orders

The BARRA Market Impact model is considered somewhat of an industry standard and was derived by Torre and Ferrari [8]. They consider TAQ data from the NYSE, aggregate trades on a half hour time interval and find $\varphi \approx 0.5$.

Gabaix et al. [9] have used TAQ data from the New York, London and Paris stock markets and found $\varphi \approx 0.5$, although for relatively short time intervals (around 15 minutes).

Moro et al. [10] considers transactions on the London Stock Exchange between January 2002 and December 2004 as well as transactions on the Spanish Stock Exchange from January 2001 to December 2004. They find that $\varphi \approx 0.5$ for the Spanish Stock Exchange and a slightly larger value for the London Stock Exchange.

Almgren et al. [11] use a data set consisting of almost 700,000 US stock trade orders executed by Citigroup Equity Trading desks. They consider almost 700,000 US stock trade orders executed by Citigroup Equity Trading desks and divide impact in a permanent part that changes as the order is executed and a temporary part that decays immediately after a transaction has been executed. The permanent part is to be interpreted as the market impact of trading and the temporary part as representing fixed costs and possibly any part of the impact that decays after the trade. For the temporary part they find $\varphi \approx 0.6$ and for the permanent part they find $\varphi \approx 0.9$.

Engle, Ferstenberg and Russell [2] consider a sample of more than 200,000 orders executed by Morgan Stanley and measure, not market impact, but transaction cost. They find transaction cost (relative the volume liquidated) increases approximately as a square root function of volume. Under the assumption that fixed costs (mainly the bid-ask spread) are linear in volume, this implies a square root market impact function.

3.2.3 Block trades

Block trades in the upstairs market have been studied by Keim and Madhavan [12]. They use trading history from an investment management firm which contains trade dates, trade prices, number of shares traded and commissions paid for all block trades in which the firm participated between 1985 and 1992. They find a concave dependence on volume but do not try to fit the results to a function.

3.3 Why is the market impact function concave?

There are a number of plausible explanations why market impact is concave. For individual trades, it has been suggested that the information of small trades hold approximately the same information as large trades. This implies that relative their volume, small trades impact the market more than large trades. Farmer et al. [14] suggests concavity is due to selective liquidity taking. This means that traders are more likely to place large market orders when the available liquidity is favorable, implying that large volumes do not have a smaller relative impact per se, but rather that large volumes are only traded when they do not have a large impact.

Another explanation can be derived from the empirical finding that order flow[2] is autocorrelated which was found independently by Lillo and Farmer [15] and Bouchaud [16]. The fact that order flow is autocorrelated implies that future order flow is predictable and this raises an obvious question; if buying pushes prices up, selling pushes prices down and the sign of future orders are predictable, how do returns remain unpredictable? The return on a tic-by-tic scale is equal to the expected impact of buys and sells weighted with the respective probability of their occurrence. Farmer and Lillo argue that impact is not fixed but depends on the markets predictions of the future order flow. More specifically, a buy order followed by a series of buy orders must have a smaller impact (on average) for returns to remain unpredictable. It is suggested that this is caused by liquidity providers changing their quotes but it is emphasized that this is only one possible explanation of many.

It should be pointed out that the market impact function is not concave for very small volumes. Any volume that can be executed against available liquidity at the best price will not have any impact on the spot price. Thus, the market impact function should be slightly S-shaped. The majority of studies considered in this section argue that impact is concave even for very small volumes. The fact that the available liquidity at the best prices is typically less than a percent of the daily volumes is an indication of

[2] The order flow refers to the series of signed trades in the market. The sign of a trade is positive if it is initiated by the buyer and negative if it is initiated by the seller.

this. In the other end of the spectrum, for volumes corresponding to large fractions of the market capitalization of a company, it might be assumed that a concave function does not model impact very well. However, such volumes imply responsibility for the strategic decisions of the company and are more relevant for venture capitalists than hedge funds.

3.4 Time dependence

It is intuitive that the time taken to liquidate a position (often referred to as the trading velocity) affects market impact. Still, few papers have attempted to empirically fit the time dependence to a function. As mentioned previously, trading over time introduces a dependence on the way trading is done. Is the trader liquidating her positions with a constant trading velocity? Is she trying to match the available liquidity with her orders? Is she using some other strategy known only to herself? There are essentially endless ways in which a large order can be executed over a given time interval.

The majority of studies assume traders use a constant trading velocity. For instance, this is assumed in the study by Almgren et al. [11] previously introduced. They introduce a time dependence in their market impact model by assuming that impact should not be measured as a function of the volume alone, but rather as a function of the fraction of the volume that is "normally" traded during the relevant time interval. Thus, they measure the dependence, not on volume, but on the variable

$$\frac{X}{VT}$$

where X is the traded volume, V is the daily volume and T is the time interval considered. By introducing this variable they implicitly assume that trading over an *n* times longer time interval is equivalent to trading in an *n* times deeper market. The reasoning behind this argument is fairly intuitive. However, the expression implies that trading one days volume over one day has the same impact as trading a very small fraction of the daily volume during an equally small time interval. It is apparent that this expression gives unreasonable results if used for very short time intervals.

In contrast, Engle, Ferstenberg and Russell [2] find only a weak dependence on trading velocity when measuring transaction costs in their extensive study.

Studies about price manipulation investigate specific trading strategies in more detail. The purpose is to investigate if models for market impact allow for price manipulation strategies with positive expected profits. A model introduced by Bouchaud et al. [16] and generalized by Gatheral [18] hypothizes that the impact of trading continuously reverts back to the pre-execution price level. In one way this is reasonable since it assumes trading only temporarily impacts the price from some fundamental value. However, it is difficult for the market to tell whether an order is executed because the trader holds superior information about the company or simply because he wants to liquidate his position to raise cash. Furthermore, it is not unlikely that price manipulation strategies with positive expected payoffs exist. The fact that there exist regulations in this area is indicative that the strategies might be potentially profitable.

Seemingly, there is no consensus about how market impact depends on the trading velocity other than that a more rapid execution increases costs.

3.5 Other variables explaining market impact

Other than the volume and the trading velocity, the most used variable is the volatility. In the BARRA market impact model [8] Torre argues that higher volatility implies larger market impact but he does not fit any functional form. Grinold and Kahn [19] introduce a model for transaction costs, seperate market impact from other costs and scale impact between stocks with the volatility. They state that their model fits nicely with the trader rule of thumb that "it costs one day's volatility to trade one day's volume". Almgren et al. [11] argue that by trading a stock over time a trader participates in the "normal" motion of the stock and that the volatility is a good proxy for this. They do not present statistical results but argue that volatility is the most important scaling variable between stocks.

The bid-ask spread has an obvious effect on the transaction cost but not necessarily on market impact. It is true that the spread tends to widen when markets become illiquid in crises but this effect is also captured in the volatility which also generally increases in stressed markets. Introducing a dependence on both the bid-ask spread and the volatility might be superfluous.

There are also arguments for conditioning impact on the market captialization. This dependence is to some extent captured in the dependence on the daily volume but the market capitalization also shows the hidden liquidity. Almgren [11] finds a weak dependence on the market capitalization, but with large error bars. Bikker et al. [20] finds that the dependence on market capitalization pales in comparison to volatility.

3.6 LOB modeling

I will briefly mention another approach from recent studies that are able to measure impact more accurately than the studies considered so far. Instead of conditioning impact on the volume traded, impact is conditioned on a microscopic variable; the *order flow imbalance*. This variable is a measure of the imbalance between supply and demand in the order book. Cont et al. [13] considers a data set consisting of one calendar month (April, 2010) of trades and quotes (TAQ) data for 50 stocks and claim that the order flow imbalance is able to explain impact to a far higher degree than conditioning it on macroscopic variables. This is intuitive since microscopic variables describe the actual structure of the limit order book at any time. However, from a risk management perspective there is a major disadvantage with using microscopic variables. A condition for the model in this thesis is that it should be based on variables that are easily available and relatively easy to measure - this is generally not the case with the order flow imbalance.

3.7 Modeling market impact

The previous section shows that there is overwhelming evidence that market impact is a concave function of the traded volume. The exact functional is unknown but empirical studies indicate that it is a power law with an exponent close to 0.5. This is also in line with various transaction cost models that propose that market impact contributes to transaction costs as a square root function of the traded volume.

There is less consensus about the dependence on the trading velocity. Some models have been proposed but they either give unreasonable results (Almgren), consider microscopic variables (LOB modeling) or are to a large extent theoretical (price manipulation). The lack of robust empirical results concerning the time dependence means that including it in the model will significantly decrease the validity of the results.

I avoid this by considering a fixed time interval for the liquidation, thereby eliminating the time dependence from the model. This is not necessarily a significant limitation of the model. A reasonable assumption is that the volume determines the absolute impact of a trade and that the trading velocity only scales the impact. If the scaling is equal for all stocks this implies that deriving the optimal weights to liquidate for *any* time interval will give the optimal weights to liquidate for *all* time intervals. The time interval chosen will only scale the total cost of the liquidation.

In the beginning of this section I introduced the following model for the stock price as a function of the market impact of our trading.

$$S(t) = S(0)\big(1 + \epsilon\, f\big(x(t); \Omega\big) + \xi(t)\big)$$

With the arguments and empirical results presented in the section above I propose the following model for market impact.

$$f(x(t); \Omega) = \sigma_d \sqrt{\frac{|x(t)|}{V}}$$

Here σ_d is the daily volatility, V is the daily volume and $x(t)$ is the volume executed up until time t. Thus, market impact is a square-root function of the traded volume and it is scaled with the volatility and daily volume of the stock. A more volatile stock is expected to have a larger market impact and a stock that is more frequently traded is expected to have a smaller impact.

Introducing the market impact function in the model for the stock price gives

$$S(t) = S(0)\big(1 + \epsilon\, \sigma_d \sqrt{\frac{|x(t)|}{V}} + \xi(t)\big)$$

and trivially

$$E\big[S(t)\big] = S(0)\big(1 + \epsilon\, \sigma_d \sqrt{\frac{|x(t)|}{V}})$$

Figure 2 shows the expected future spot price as a function of the number of daily volumes liquidated during the time interval $(0, T)$. The stock has a spot price of 100 and a daily volatility of 1%.

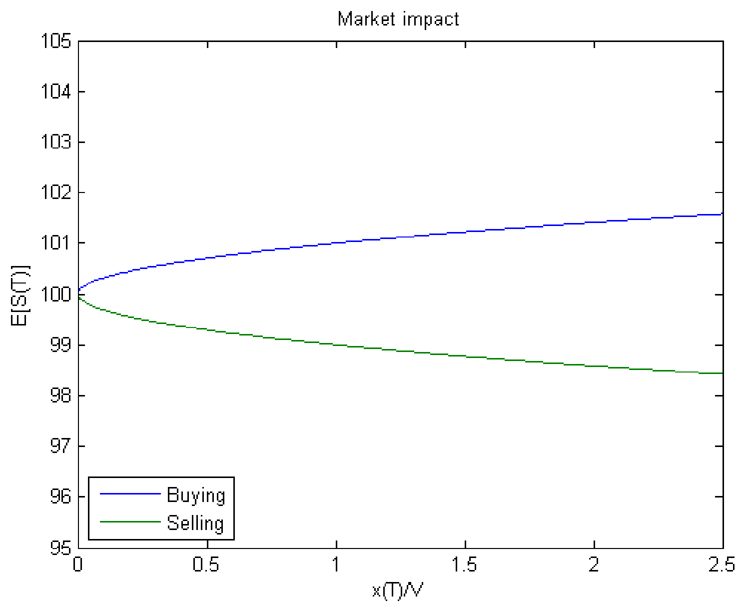

Figure 2: The expected spot price of a stock at time T as a function of the number of daily volumes bought and sold during the time interval $(0, T)$

4 Transaction costs

Most transaction cost models split costs up in different parts. Generally,

Transaction cost = Commission + Bid-ask spread + Market impact costs

The commission term reflects the fixed costs to be paid to a broker if trading is outsourced. This term can be excluded if the portfolio manager places orders himself. The second term addresses the fact that a stock is never traded at the mid price, but rather at the best bid or best ask price. This thesis will disregard the first two terms and only focus on transaction costs coming from the market impact of our trading. The argument for this is that commission and spread costs are linear in trade size and easily added to the cost resulting from market impact. Also, the market impact cost will tend to dominate for large orders.

4.1 Transaction costs due to market impact

Consider liquidating a volume X (which can be negative) by splitting it up and executing it over the time interval $t \in (0, T)$. Let $x(t)$ denote the volume that has been transacted up until time t (where $x(0) = 0$ and $x(T) = X$). At each infinitesimal time interval dt, the volume $dx(t)$ will be executed at the spot price $S(t)$. Thus, the cash raised from using a given liquidation strategy $x(t)$ is

$$\text{Cash raised} = \int_0^T S(t)\ dx(t)$$

To be able to model this expression as an integral I assume that continuous trading is possible. This is obviously practically impossible since it requires splitting orders in infinite small parts. However, continuous trading is a good approximation if orders are split up in many pieces. The obvious benchmark to compare this value with is the value of the position before the liquidation, i.e. $X \cdot S(0)$. I define the *liquidation cost* as the difference between the mark-to-market value of the position before trading and the cash that is actually raised from the liquidation.

$$\text{Liq. cost} = X \cdot S(0) - \int_0^T S(t)\ dx(t) = \int_0^T \Big(S(0) - S(t)\Big)\ dx(t)$$

As mentioned, I only consider the market impact part of transaction costs. The liquidation cost, as defined here, is a part of this cost. However, as will be shown later in this section, there is also an indirect cost related to market impact.

4.2 Square root market impact

The expression above is an ex post measure of the liquidation cost. The *expected* liquidiation cost can be modeled by introducing a model for the stock price. I introduce the market impact model derived in Section 3.

$$\begin{aligned}
\text{E[Liq. Cost]} &= \int_0^T \Big(S(0) - E[S(t)]\Big)\, dx(t) \\
&= \int_0^T \Big(S(0) - S(0)(1 + \epsilon f(x(t)))\Big)\, dx(t) \\
&= -\int_0^T S(0)\; \epsilon\; f(x(t))\, dx(t) \\
&= -\;\epsilon \int_0^T S(0)\sigma_d \sqrt{\frac{|x(t)|}{V}}\, dx(t) \\
&= \frac{\sigma_d S(0)}{\sqrt{V}} \frac{2}{3} |X|^{3/2}
\end{aligned}$$

which is given in units of some currency. Notably, according to this model, the liquidation cost is independent of the liquidation strategy $x(t)$. There are an infinite number of ways to liquidate a given volume during a given time interval and different liquidation strategies likely imply different impact on the market. I assume, like many other studies, that the trader uses a constant trading velocity. This is in line with how trading is actually done in practice[3].

The derivation above shows that the absolute cost increases as a power law with exponent $3/2$. However, the cost relative the value of the position liquidated (more specifically its mark-to-market value when trading begins) is given by

$$\frac{1}{|X| \cdot S(0)} \frac{\sigma_d S(0)}{\sqrt{V}} \frac{2}{3} |X|^{3/2} \sim |X|^{1/2}$$

This is in line with the “square root” cost function suggested in literature [19], [21].

4.3 Indirect cost

The situations considered so far have only focused on positions that are liquidated completely. However, in many situations a portfolio manager will liquidate only a fraction of a position. If trading impacts spot prices, this is going to change the mark-to-market value of the shares that remain in the portfolio after the liquidation. More specifically, consider holding a volume Y in an asset and then liquidating a smaller volume X during the time interval $t \in (0, T)$. I define the *indirect cost* as the decreased mark-to-market value of the remaining shares.

$$\text{Indirect cost} = \big(Y - X\big)\big(S(0) - S(T)\big)$$

[3]Two common ex post benchmarks for the cash raised from a liquidation is the volume weighted average price (VWAP) and the time weighted average price (TWAP). They are calculated at the end of the day when the exchange closes. A broker that executes an order on behalf of a customer can often guarantee the customer the VWAP or TWAP of his position. TWAP corresponds to a constant trading velocity (in time) and WVAP corresponds to matching the intradaily volume.

Introducing the market impact function derived in the previous section gives the expected indirect cost as

$$\begin{aligned}
\text{E[Indirect cost]} =& \big(Y - X\big)\big(S(0) - E[S(T)]\big)\big) \\
=& \big(Y - X\big)\big(S(0) - S(0)\big(1 + \epsilon\, \sigma_d \sqrt{\frac{|x(t)|}{V}}\big)\big)\big) \\
=& \big(Y - X\big)S(0)\big(-\epsilon\, \sigma_d \sqrt{\frac{|x(t)|}{V}}\big)
\end{aligned}$$

The indirect cost can also be negative, which is equivalent to a profit. For example, a manager who holds a large volume in a stock can purchase further shares to increase the mark-to-market value of his position. This is not surprising and there are examples of this happening in practice. As a matter of fact, this strategy is called price manipulation and is to some extent regulated. Secondly, this does not imply a profit in a strict sense. To actually realize the profit the manager will have to liquidate the positions, again impacting the market. Still, this cost will be relevant for this thesis since hedge funds use spot prices to estimate the net asset value of their portfolios.

4.4 Total liquidation cost

The total liquidation cost in this setting will be sum of the liquidation cost and the indirect cost.

$$\text{E[Total liq. cost]} = \frac{\sigma_d S(0)}{\sqrt{V}} \frac{2}{3} |X|^{3/2} + \big(Y - X\big)S(0)\big(-\epsilon\, \sigma_d \sqrt{\frac{|x(t)|}{V}}\big)$$

This expression explains the decrease in the NAV of the portfolio from the liquidation. It is different from the market impact part of the general transaction cost model introduced in the beginning of this section. No previous model takes the indirect cost into account when measuring transaction costs. The indirect cost is not really a cost, but it does decrease the NAV of the portfolio and I include it in the model.

5 The liquidation problem

The market impact framework can be extended to a portfolio setting. I will consider the situation where a hedge fund considers liquidating a relatively large fraction of its positions due to investor withdrawals. Since trading involves transaction costs the manager would prefer to liquidate assets with little market impact or, even better, avoid trading altogether. However, liquidating only some of the positions will change the relative weighting and overall properties of the remaining portfolio. Also, paying investors without liquidating any risky positions will generally increase the risk of the remaining portfolio. Thus, which liquidation strategy to use is not trivial.

5.1 Formalizing notation

Consider a portfolio at time 0 (without loss of generalization) with the weights

$$\mathbf{h}(0) = (h_1(0), h_2(0), ..., h_n(0))$$

in stocks with spot prices

$$\mathbf{S}(0) = (S_1(0), S_2(0), ..., S_n(0))$$

The portfolio also consists of a position $h_0(0)$ in a perfectly liquid money asset $S_0(0)$ with zero interest rate. I.e. $S_0(t) = 1$ for all t.

Next, consider liquidating some fractions of the risky positions during the time interval $t \in (0, T)$.

$$\Delta\mathbf{h}(0) = (\Delta_1 h_1(0), \Delta_2 h_2(0), ..., \Delta_n h_n(0))$$

The domain of Δ_i is

$$\begin{aligned}
&\Delta_i \in [0, h_i(0)]\ if\ \ h_i(0) > 0 \\
&\Delta_i \in [-h_i(0), 0]\ if\ \ h_i(0) < 0
\end{aligned}$$

I will motivate this below. At time T, the new weights in the risky positions are given by

$$\begin{aligned}
\mathbf{h}(T) = &\big((h_1(0) - \Delta_1 h_1(0)), (h_2(0) - \Delta_2 h_2(0)), ..., (h_n(0) - \Delta_n h_n(0))\big) \\
&\big((1 - \Delta_1)h_1(0), (1 - \Delta_2)h_2(0), ..., (1 - \Delta_n)h_n(0)\big)
\end{aligned}$$

In the following sections I summarize how withdrawals and a possible subsequent liquidation affect various properties of the portfolio.

5.2 The new spot price

The expected spot price of stock i, at time T is given by

$$E\big[S_i(T)\big] = S_i(0)(1 + \epsilon\, \sigma_i \sqrt{\frac{|\Delta_i h_i(0)|}{V_i}}), i \in 1, ..., n$$

5.3 Cost of trading

The cost of trading was derived in the section above. With the notation introduced in this section, the expected total liquidation cost is

$$\begin{aligned}\text{E[Total liq. cost]} = &\sum_{i=1}^{n} \frac{\sigma_i S_i(0)}{\sqrt{V_i}} \frac{2}{3} |\Delta_i h_i(0)|^{3/2} \\ &+ \big(h_i(0) - \Delta_i h_i(0)\big) S(0) \big(- \epsilon\, \sigma_i \sqrt{\frac{|\Delta_i h_i(0)|}{V_i}} \big)\end{aligned}$$

5.4 The NAV of the portfolio

The expected NAV of the portfolio immidiately after the liquidation is equal to the NAV before the liquidation less the total liquidation cost. I denote this $V_-(T)$. At this point, a fraction k of the NAV will be paid out to investors. I denote the NAV of the remaining portfolio by $V_+(T)$.

$$E[V_-(T)] = V(0) - E[\text{Total liq. cost}]$$

and

$$E[V_+(T)] = (1-k) \cdot E[V_-(T)]$$

From the expression for $E[V_-(T)]$ it is apparent that the total liquidation cost is equvalent to the decrease of the NAV of the portfolio due to the liquidation.

5.5 Portfolio risk measures

The risk measures I will consider are net exposure, gross exposure and Value-at-Risk.

5.5.1 Net Exposure

The expected net exposure of the remaining portfolio is given by

$$\begin{aligned}\text{Net Exposure} &= \frac{\sum_{i=1}^{n} h_i(T) E[S_i(T)]}{E[V_+(T)]} \\ &= \frac{\sum_{i=1}^{n} (h_1(0) - \Delta_1 h_1(0)) S_i(0) (1 + \epsilon\, \sigma_i \sqrt{\frac{|\Delta_i h_i(0)|}{V_i}})}{\big(1-k\big)\big(V(0) - E[\text{Total liq. cost}]\big)}\end{aligned}$$

5.5.2 Gross Exposure

The expected gross exposure of the remaining portfolio is given by

$$\begin{aligned}\text{Net Exposure} &= \frac{\sum_{i=1}^{n} |h_i(T)| E[S_i(T)]}{E[V_+(T)]} \\ &= \frac{\sum_{i=1}^{n} |h_1(0) - \Delta_1 h_1(0)| S_i(0) (1 + \epsilon\, \sigma_i \sqrt{\frac{|\Delta_i h_i(0)|}{V_i}})}{\big(1-k\big)\big(V(0) - E[\text{Total liq. cost}]\big)}\end{aligned}$$

5.5.3 Value-at-Risk

I use the normal assumption to compute the (expected) Value-at-Risk of the remaining portfolio. The VaR will depend on the new spot prices and the new weights in the portfolio. I assume that the covariance matrix of the stocks at time T is best estimated by the covariance matrix at time 0. Arguments can be made that the liquidation should affect the volatilities and correlations of the stocks. However, I assume that the covariance matrix captures the behavior of the stocks under normal market conditions and that the liquidation only temporarily disrupts these conditions. With these assumptions, the Value-at-Risk of the remaining portfolio, at a confidence level α and over the time interval $(T, T+\Delta)$, is given by

$$VaR_{\alpha,\Delta} = \frac{\mu_L + \Phi(1-\alpha)\sigma_L}{E[V_+(T)]} = \frac{\mu_L + \Phi(1-\alpha)\cdot\sigma_L}{\big(1-k\big)\big(V(0) - E[\text{Total liq. cost}]\big)}$$

Here, μ_L and σ_L are the expected value and standard deviation of the loss variable L introduced in Section 2. To have use of this expression, the values of these two parameters need to be computed. At this point it should be mentioned that it is common practice to disregard the expected value in the expression above. For short time intervals the expected value is typically small compared to the standard deviation. Another reason is that the expected value exhibits more seasonality than the volatility, it is less robust over time and more difficult to estimate.

It still remains to calculate the volatility of the loss variable L. Stock prices are typically modeled as geometric stochastic processes. This implies the following expression for L^{Δ}

$$\begin{aligned} L^{\Delta} &= -\big(V(T+\Delta) - V(T)\big) \\ &= -\big(\sum_{i=1}^{n} h_i(T)S_i(T)e^{R_i(\Delta)} - \sum_{i=1}^{n} h_i(T)S_i(T)\big) \\ &= -\sum_{i=1}^{n} h_i(T)S_i(T)\big(e^{R_i(\Delta)} - 1\big) \end{aligned}$$

Here, $R_i(\Delta) = R_i(T, T+\Delta)$ is the log return of stock i over the interval $(T, T+\Delta)$. For short time intervals, this expression can be approximated with a *linearization*. Noting that $e^x \approx 1 + x$ for small x the following approximation can be made

$$\begin{aligned} L^{\Delta} &= -\sum_{i=1}^{n} h_i(T)S_i(T)\big(e^{R_i(\Delta)} - 1\big) \\ &\approx -\sum_{i=1}^{n} h_i(T)S_i(T)\big(1 + R_i(\Delta) - 1\big) \\ &= -\sum_{i=1}^{n} h_i(T)S_i(T)R_i(\Delta) \end{aligned}$$

By assumtion, the returns R_i are normally distributed; $R_i(\Delta) \in N(\mu_i, \sigma_i^2)$. The variance of the loss variable is given by

$$
\begin{aligned}
\sigma_L^2 =& \text{Var}[L^{\Delta}] \\
=& \text{ Var}[-\sum_{i=1}^{n} h_i(T)S_i(T)R_i(\Delta)] \\
=& \sum_{i=1}^{n}\sum_{j=1}^{n} Cov\big(h_i(T)S_i(T)R_i(\Delta), h_j(t)S_j(t)R_j(\Delta)\big) \\
=& \sum_{i=1}^{n}\sum_{j=1}^{n} h_i(T)h_j(T)S_i(T)S_j(T)\sigma_i\sigma_j\rho_{i,j}
\end{aligned}
$$

Here $\rho_{i,j}$ is the correlation coefficient between return $R_i(\Delta)$ and $R_j(\Delta)$. This is the variance of the portfolio at time T. When we start the liquidation at time 0, the stock prices $S_i(T)$ are stochastic and the expected variance is given by

$$
Var[L^{\Delta}] = \sum_{i=1}^{n}\sum_{j=1}^{n} h_i(T)h_j(T)E[S_i(T)]E[S_j(T)]\sigma_i\sigma_j\rho_{i,j}
$$

The expected VaR of the portfolio at T is given by introducing the variance in the expression for the VaR above (and excluding the expected value).

5.6 Constraints on the weights to liquidate

I also impose restrictions on the weights that can be liquidated. More specifically, I do not allow increasing the absolute volume in any stock, nor liquidating *more* than the volume we hold. The reason for this is, as mentioned, that this is equivalent to price manipulation. Price manipulation is a vast area of its own and would make the investment universe significantly larger. It will also likely give extreme solutions since it would allow negative costs. Formally,

$$
\begin{aligned}
&\Delta_i \in [0, h_i(0)] \; if \;\; h_i(0) > 0 \\
&\Delta_i \in [-h_i(0), 0] \; if \;\; h_i(0) < 0
\end{aligned}
$$

5.7 Different liquidation strategies

Choosing which liquidation strategy to use is a balance between minimizing the cost of trading and keeping the properties of the portfolio intact for the remaining investors. I will consider three approaches; the *naive approach*, *proportional liquidation* and *optimized liquidation*.

The naive approach is to simply pay investors without liquidating any risky positions. The obvious advantage with this approach is that it does not involve any trading costs.

However, paying out money without liquidating any risky positions decreases the NAV of the portfolio without decreasing the market exposure. This increases the risk of the portfolio. If a withdrawal increases the risk of the portfolio for the remaining investors, this encourages investors to withdraw their money before anyone else. This is obviously not a preferable consequence of a liquidation strategy.

Proportional liquidation, on the other hand, means the manager scales the portfolio down by liquidating equal fractions of each position. This strategy has the advantage that the properties of the remaining portfolio remain more or less unchanged[4]. The fact that these properties have been deliberately chosen by the (supposedly skilled) manager is further argument for a proportional liquidation strategy. However, this is still not necessarily the optimal liquidation approach. Firstly, proportional liquidation is not always even possible; assets cannot be divided into arbitrarily small parts. Secondly, I will find that proportional liquidation is relatively expensive compared to other liquidation strategies. Thus, it can be argued that it is motivated to deviate from proportional liquidation if this implies a lower expected cost.

How much to deviate from proportional liquidation depends on the preferences of the portfolio manager or the risk manager. By formulating an optimization problem with relevant constraints it is possible to find solutions that imply smaller transaction costs without changing the properties of the portfolio beyond some pre-specified limit. I define *optimized liquidation* to be the result of this optimization problem. It is not obvious which constraints to use. From a risk manager point of view it is important that the remaining portfolio satisfies all specified risk limits. For a portfolio manager, however, it is probably more important to keep a certain "profile" of the portfolio. This could for example mean that the portfolio should have a certain exposure to a specific asset, exposure to companies in a certain geographical area or that the portfolio manager is unwilling to liquidate certain positions that he considers profitable. However, it would be difficult to generalize this in mathematical terminology. I will instead formulate the constraints from a risk manager point of view. To address the portfolio manager's point of view, I will compare the tracking error of the remaining portfolio relative the pre-liquidation portfolio. This is a (somewhat crude) measure of how similar the two portfolios are.

[4]Actually the cost of trading will change the properties of the portfolio. However, only slightly.

6 Case studies

To illustrate the liquidation problem I consider two case studies. First, I form a toy portfolio consisting of two stocks that I use to show graphically how trading different combinations of the positions affect the properties of the remaining portfolio.

I also consider a portfolio consisting of four stocks from the Nasdaq Omx Nordic. They have been chosen to represent a portfolio that consists of some liquid and some more illiquid stock positions. I use that portfolio to make a more thorough analysis of the problem.

6.1 A two stock portfolio

Consider a portfolio consisting of two stocks and a cash postition. The stocks are weighted so that the portfolio is perfectly net neutral. The properties of the portfolio is given in Table 1. The risk measures are calculated as specified in Section 5. Furthermore, the covariance matrix Σ of the stocks daily log returns is given by

	Stocks A	Stock B	Cash
Spot price	100	125	1
Daily volume	1000	1250	–
Weights	-2500	2000	25000
% of Gross	50	50	-
% of NAV	-100	100	100
NAV	25000		
Net Exp.	0 %		
Gross Exp.	200 %		
Value-at-Risk	3.8 %		

Table 1: Properties of two stock portfolio

$$\Sigma = 10^{-4} \cdot \begin{pmatrix} 1.44 & 1.20 \\ 1.20 & 6.250 \end{pmatrix}$$

This implies that the daily volatility of the stocks (the square root of the variances in the covariance matrix) is 0.12 and 0.25 respectively.

6.1.1 Total liquidation cost

The cost of trading depends on the market impact function which scales with the volatility and the daily volume. Since the stocks have otherwise similar properties Stock B is expected to be more expensive to trade since it is more volatile. This is seen in Figure 3 which shows the cost of liquidation as a function of the number of shares of Stock A and Stock B liquidated.

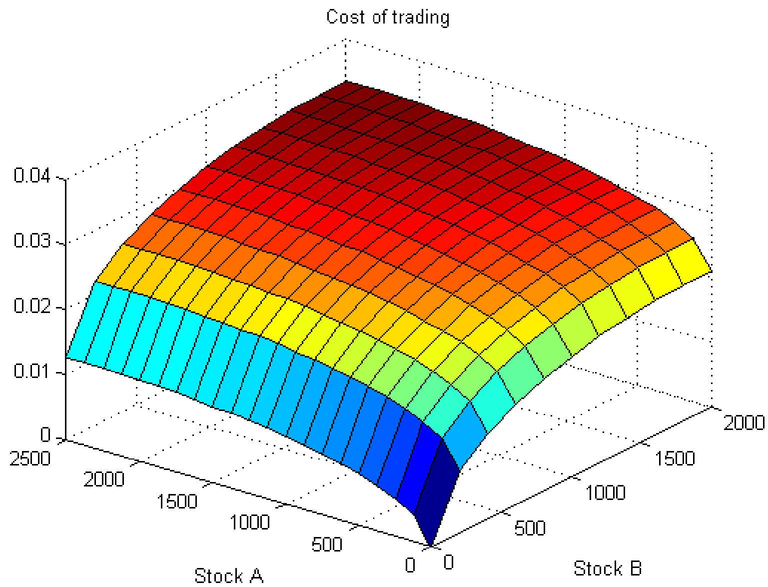

Figure 3: Cost of trading: The figure shows the total liquidation cost (relative the NAV of the pre-liquidation portfolio) as a function of the number of shares liquidated of Stock A and Stock B. Stock B has a higher volatility and is therefore more expensive to liquidate.

One thing that is apparent from the figure is that it generally is cheaper to liquidate only one of the positions rather than a combination of the two. This is because the concavity of the cost function implies a decreasing marginal cost. This can be illustrated graphically. Assume a portfolio manager wants to decrease the gross exposure of his portfolio by 50%. Figure 2 shows the cost of doing this for all linear combinations of Stock A and Stock B. I.e. the figure shows the cost of liquidating

$$\alpha S_A(0) + (1-\alpha)S_B(0) = 25000, \alpha \in [0,1]$$

The figure clearly shows that it is generally less expensive to only liquidate one position than to liquidate a combination of two positions. Notably, proportional liquidation (corresponding to the value 0.5 on the x-axis) is relatively expensive. The fact that it generally is less expensive to liquidate only one stock will become apparent in the next case study, where I solve an optimization problem.

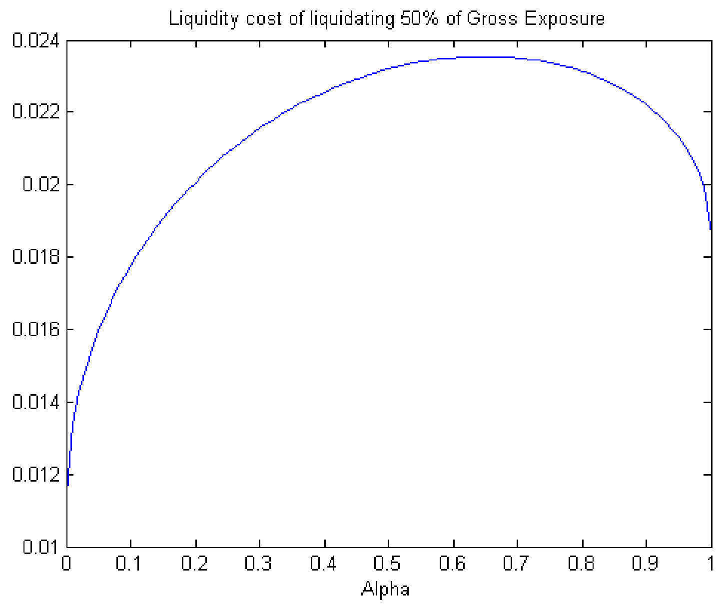

Figure 4: Comparing liquidity cost: The figure shows the total liquidation cost (relative the NAV of the pre-liquidation portfolio) of liquidating 50 percent of the mark-to-market value of the stock positions for all linear combinations of Stock A and Stock B. Notably, proportional liquidation (corresponding to 0.5 at horisontal axis) is relatively expensive.

6.1.2 Risk measures

I consider three risk measures, the Value-at-Risk, the net exposure and the gross exposure. Figure 5 shows the Value-At-Risk of the portfolio after we liquidate a given fraction of our stock positions. The VaR also depends on the fraction of the NAV that is paid out to investors after the liquidation. In this figure, I consider the situation where this amount is zero. If some fraction is paid out, the VaR-surface would have the same shape, only scaled.

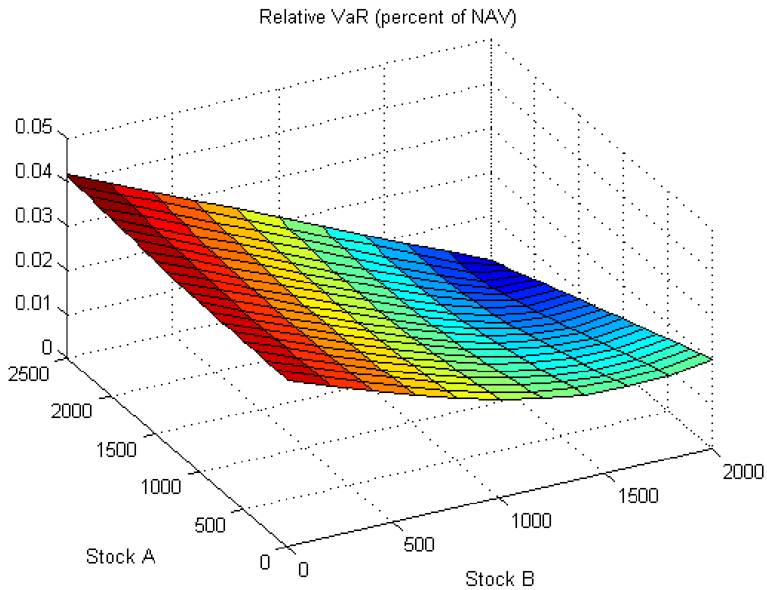

Figure 5: Value-at-Risk: The figure shows the VaR of the remaining portfolio when a given fraction of Stock A and Stock B is liquidated. Liquidating only Stock B will decrease the VaR since the remaining portfolio is then left with a higher relative weight in Stock A, which is less volatile.

The figure shows that if all shares are sold, then the VaR becomes zero since we do not have any risky positions left. Figure 6 shows the same graph from another angle.

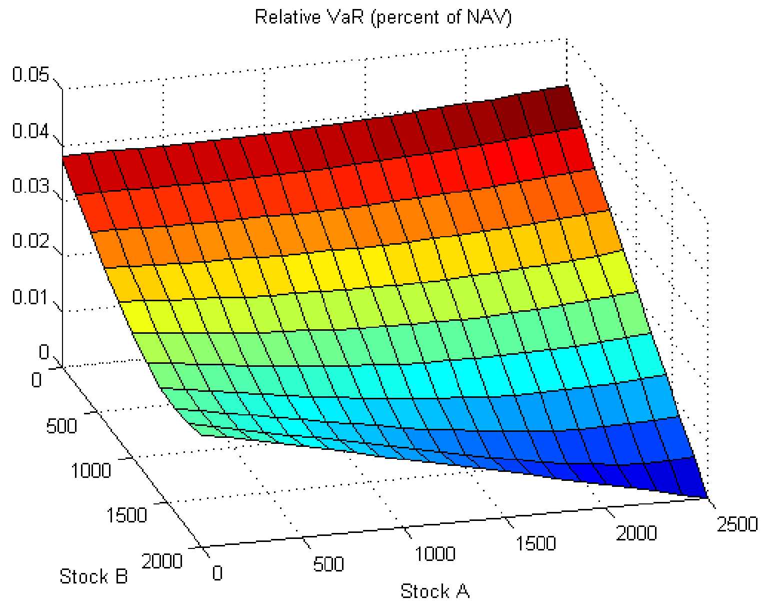

Figure 6: Value-at-Risk: This is the same figure as Figure 4, only from a different angle.

Figure 7 and Figure 8 show the net and gross exposure of the remaining portfolio. Again, I consider the situation where no fraction of the NAV is paid out to investors after the liquidation.

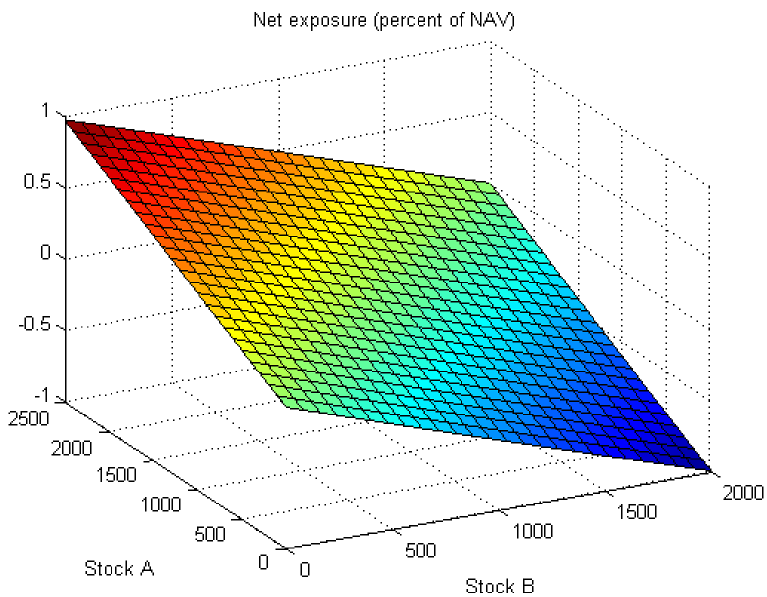

Figure 7: Net Exposure: The figure shows the net exposure of the remaining portfolio when a given fraction of Stock A and Stock B is liquidated. Liquidating Stock A will leave the remaining portfolio net positive since the portfolio holds a short position in Stock A

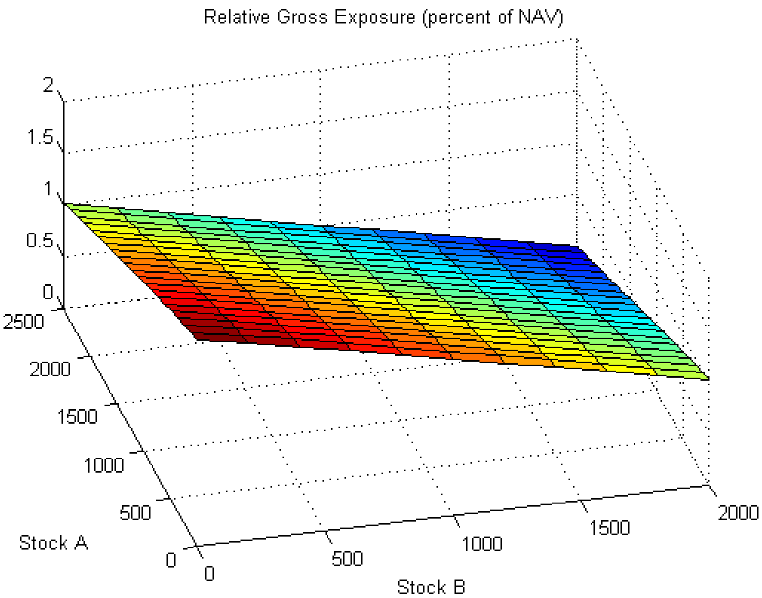

Figure 8: Gross Exposure: The figure shows the gross exposure of the remaining portfolio when a given fraction of Stock A and Stock B is liquidated. Liquidating shares in any stock will decrease the gross exposure.

The figures in this section illustrate the properties of the problem that is to be solved. Figure 3 and Figure 4 show that the cost of liquidating only one stock is generally less expensive than to liquidate a combination of two stocks with the same mark-to-market value. However, Figures 6-8 show that liquidating shares in only one of the stocks increases the risk more than liquidating a combination of two stocks. Thus, the figures illustrate the balance that has to be made between reducing the market impact cost of trading and managing the risk of the remaining portfolio. This problem will be discussed in more detail in the following section, where I consider a portfolio of four stocks.

6.2 A four stock portfolio

I consider four stocks from Nasdaq Omx Nordic; H&M, Danske Bank, Eniro and DSV. I use data from 2011-02-01 to 2011-08-20 to estimate the daily volume, volatility and covariance matrix of the stocks. I estimate the daily volume with a 30 day average. The properties of the stocks and the portfolio is given in Table 2.

The portfolio is approximately net neutral with a leverage slightly above two. The stocks are supposed to represent an equity portfolio consisting of a mix of stocks that are traded more and less frequently and have different volatilities. The covariance ma-

	H&M	Danske Bank	Eniro	DSV	Cash
Spot price	188	77.0	16.7	100	1
Daily Volatility	1.54%	2.33%	3.33%	1.70%	0
Daily Volume	5.3 M	2.3 M	0.9 M	1.2 M	-
Portfolio properties					
Weights (in million shares)	3	-5	4.25	-2.5	$6 \cdot 10^8$
Weights (in daily vol.)	0.566	-2.17	4.72	-2.08	-
Exposure (percent of gross)	44.4	30.3	5.60	19.70	-
Exposure (percent of NAV)	94.0	-64.2	11.8	-41.7	100
NAV (of portfolio)	$6.0 \cdot 10^8$				
Risk measures (percent of NAV)					
Net Exposure	0 %				
Gross Exposure	217 %				
Value-at-Risk	3.29 %				

Table 2: Properties of the four stock portfolio

trix of the one day log returns of the stocks is given by

$$\Sigma = 10^{-4} \cdot \begin{pmatrix} 2.39 & 0.85 & 1.30 & 0.89 \\ 0.85 & 5.43 & 2.84 & 1.99 \\ 1.30 & 2.84 & 11.1 & 2.10 \\ 0.89 & 1.99 & 2.10 & 2.87 \end{pmatrix}$$

6.2.1 The liquidation problem

I summarize the liquidation problem that has been discussed throughout the thesis. There are a number of properties that are preferable for a liquidation strategy.

1. The total liquidation cost should be as small as possible. This cost is equal to the decrease in the NAV of the portfolio.

2. From a risk manager point of view, the risk of the remaining portfolio must not exceed any pre-specified risk limits.

3. From a portfolio manager point of view, the remaining portfolio should have a profile similar to the original portfolio.

The properties are to some extent incompatible and the problem can be formalized as an optimization problem. I will consider the total liquidation cost as the objective function to minimize and take the risk manager point of view and introduce various risk measures as constraints. To address the portfolio manager, I will analyze the result by comparing the tracking error of the portfolio before and after the liquidation. The resulting optimization problem can be expressed, heuristically

$$\begin{array}{ll} \text{minimize} & \text{expected total liquidation cost} \\ \text{s.t.} & \text{Value-at-Risk} \leq \text{pre-specified limit} \\ & \text{Gross exposure} \leq \text{pre-specified limit} \\ & \text{Net Exposure} \leq \text{pre-specified limit} \\ & \text{Weights liquidated} \in \text{Allowed domain} \end{array}$$

The objective function to minimize is intuitive since the total liquidation cost (as defined) is equal to the expected decrease in the NAV of the portfolio from the liquidation. Since investors are paid after the NAV of the portfolio has been re-estimated, this cost affects both investors withdrawing their capital and the investors that remain.

The objective function and risk measures have been explicitly expressed in Section 5. For this case study I will use the following risk limits (all relative the NAV of the portfolio).

- Value-at-Risk: 4 %
- Net exposure: +/- 50 %
- Gross exposure: 250 %

Furthermore, I will consider the liquidation over a time interval of one day, i.e. $T = 1$. As mentioned in Section 3.7 it is necessary to consider a fixed time interval since there are no robust results regarding the time dependence of market impact. The reason the time interval chosen is one day is because that, at Brummer & Partners, that time interval is representative of the time intervals over which large orders are typically liquidated.

6.3 Method

The optimization problem is highly non-linear; both the objective function and the constraints. There are a number of different algorithms that can be used to solve optimization problems of this type. I have used *sequential quadratic programming* (SQP).

SQP is one of the most effective methods for solving non-linear constrained optimization problems [22]. It is an iterative algorithm which solves, at each iteration, a quadratic programming problem. I have used the SQP-algorithm that is implemented in the software MATLAB. In the results section I plot the cost function and various constraints as a function of the fraction of the NAV that is to be paid out to investors. Because of this, I have solved the optimization problem for a large number of such values.

The optimization problem has several local minima which makes the outcome from the algorithm highly dependent on the first guess. Because of this, a lot of time has been spent trying different start guesses and comparing results. Fortunately, the local minima are not found arbitrarily in the feasible region. The minima are typically found when one stock is liquidated to the largest possible extent and one (or more) other stock is liquidated further just so that the result satisfies the risk constraints. To find the global minima I used start guesses for different combinations of stocks that have this property, and compared the results. I also used random start guesses to find local minima that do not have this property in the unlikely event that any of these local minima are actually the global minimum (which they were not).

The difficulty in finding the global minimum will lead to problems when generalizing this framework to a real portfolio that typically consists of a significantly larger number of positions. However, it is not necessarily of outmost importance to find the global minimum for problems of this type. A portfolio manager will likely be satisfied if he can find any reasonable liquidation strategy that has a significantly lower expected cost than some other strategy he considers.

7 Results

The different liquidation strategies considered in this thesis (the *naive approach*, *proportional liquidation* and *optimized liquidation*) give significantly different results. In this section I analyse the total liquidation cost, different risk measures and tracking error of the remaining portfolio when the different strategies are used.

7.1 Total liquidation cost

The total liquidation cost of the naive approach is identically zero for all fractions of the NAV paid out to investors. It should be expected that proportional liquidation implies significant liquidation costs, whereas the costs for optimized liquidation is smaller.

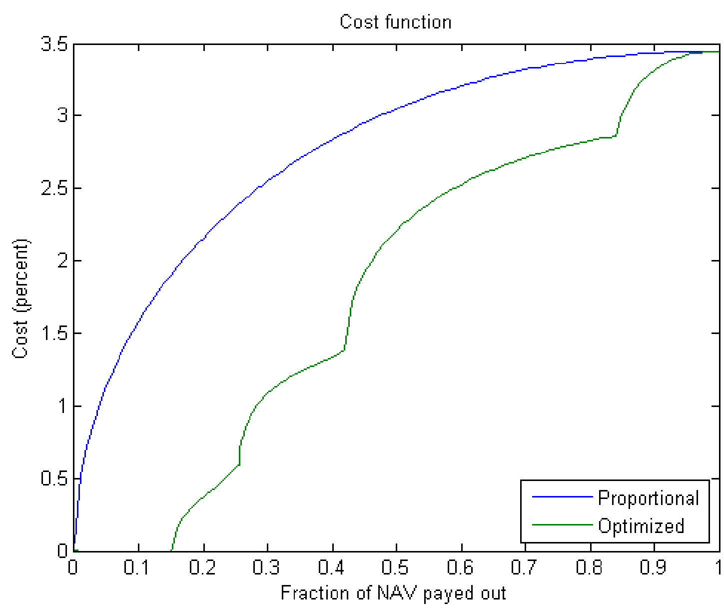

Figure 9: Total liquidation cost: The figure shows the expected total liquidation cost for different liquidation strategies as a function of the fraction of the NAV that is to be paid out to investors. Optimal liquidation is always less expensive than proportional liquidation. The cost of the naive approach is identically zero and omitted from the graph.

7.2 Risk measures

It should be expected that the risk from proportional liquidation changes only slightly due to the transaction costs that decreases the NAV. The optimized liquidation on the other hand will avoid trading until the first constraint becomes active. Finally, for the naive approach, risk will increase out of control as the NAV is decreased while the market exposure remains unchanged. The only risk measure that is left unchanged is the net exposure. This is because the *absolute* net exposure is always zero. However, only a slight deviation from zero of the absolute net exposure will lead to a significant *relative* net exposure. The plots of the risk measures can help explain the shape of the optimized cost in Figure 9. The general tendency of the optimization algorithm is to liquidate shares in only one stock to the furthest possible extent. At some point, a risk constraint becomes active and the algorithm will find that it needs to liquidate shares in another stock to satisfy the constraints. After presenting plots of the risk measures I will summarize which contraints that are active for each fraction of the NAV and what this implies for the optimized liquidation strategy.

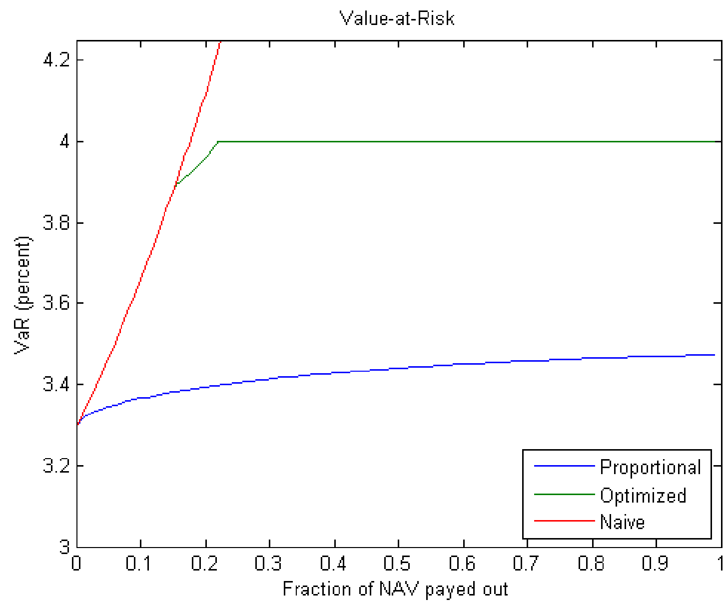

Figure 10: Value-at-Risk: The figure shows the expected VaR of the remaining portfolio as a function of the fraction of the NAV that is to be paid out to investors. Proportional liquidation implies only a slight increase in the VaR, coming from the fact that transaction costs will decrease the NAV. Optimal liquidation will never increase the VaR above the prespecified limit of 4 per cent. The naive approach of not trading at all will imply increasing risk as money is being paid out from the NAV. The risk from this approach is asymptotically infinite.

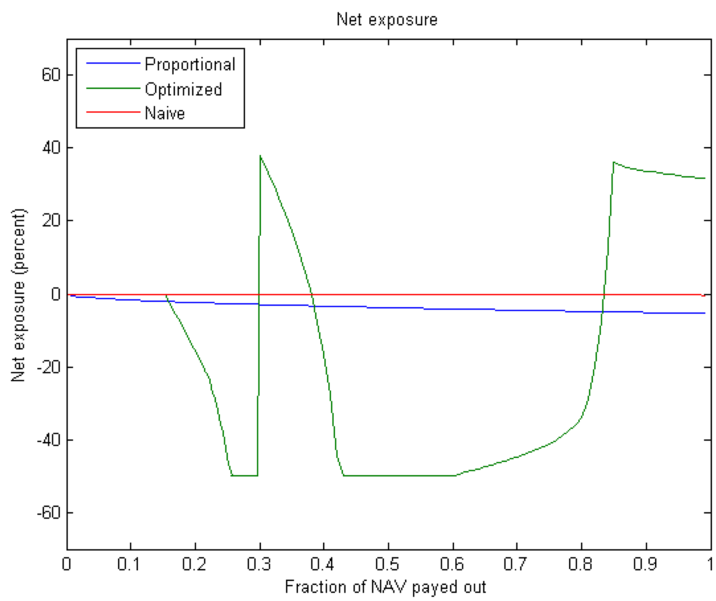

Figure 11: Net Exposure: The figure shows the expected Net Exposure of the remaining portfolio. The expected net exposure of the naive approach is zero since the expected *absolute* risk exposure is zero. However, the portfolio is highly leveraged and a small deviation in any of the stock prices will lead to a significant net exposure of the portfolio. The expected net exposure of the proportional liquidation is slightly non-zero due to the difference in impact of the different stocks. The behavior of the optimized liquidation will be discussed in the next section.

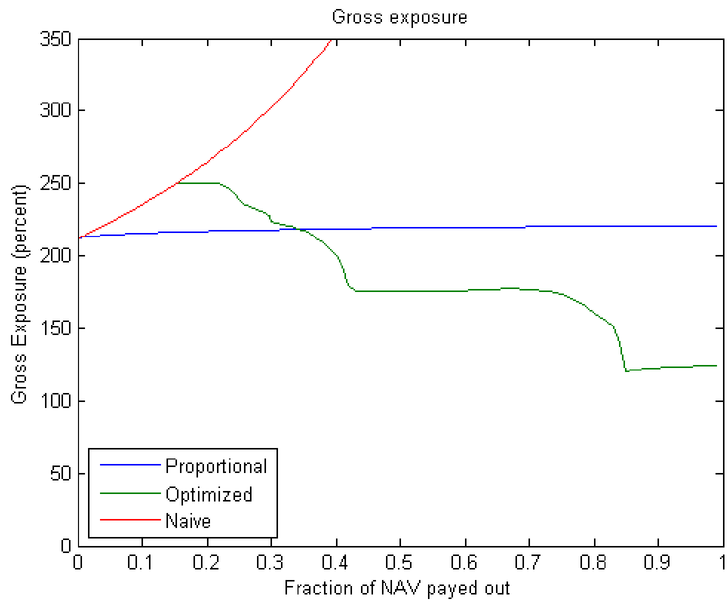

Figure 12: Gross Exposure: The figure shows the expected Gross Exposure of the remaining portfolio. The figure will be discussed in more detail in the following section

7.3 Order of liquidation

Figure 13 shows the fractions of each stock position that is liquidated, as a function of the NAV to be paid out to investors. This figure can be used, together with the figures showing the risk measures, to analyze what the optimization algorithm suggests for each fraction of the NAV.

For small volumes, no risk constraint is active and the algorithm suggests simply paying investors cash without liquidating any of the stock positions. For a fraction of the NAV of about 15 %, the gross exposure constraint becomes active. At this point the algorithm suggests liquidating H&M which is the stock with the smallest expected market impact.

The next constraint that becomes active is the Value-at-Risk at about 22 %. At this point there are two ways to go; either liquidating a small volume in another stock or keep liquidating H&M. However, in order to liquidate more shares in H&M it is necessary to liquidate disproportionally larger volumes in order to reduce the market exposure relative the NAV. Simply put, in order to be able to pay out an additional 100 SEK, it might be necessary to liquidate stock positions worth 150 SEK to keep the risk constraints satisfied. If another stock is liquidated, the latter sum might only be 125

SEK but that might still imply a larger cost due to market impact. Indeed, the algorithm suggests liquidating H&M further and the increasing derivative of H&M in the graph at this point indicates that it really is liquidated disproportionally (as suggested).

At about 25 % the net exposure constraint becomes active. At this point, it is necessary to liquidate a short position to keep this constraint satisfied. The algorithm suggests liquidating DSV which has the lowest marginal expected total liquidation cost of the short positions.

At about 30 % the algorithm suggests changing strategy drastically. It finds that it is less expensive to liquidate all of DSV than to keep liquidating in a similar manner. This result is to some extent expected since this strategy allows liquidating a smaller absolute volume and still satisfy the Value-at-Risk constraint. After this, the suggested strategies are rather predictable. The concavity of market impact implies a decreasing marginal (expected) total liquidation cost. Because of this, the algorithm suggests liquidating shares in as few stocks as possible. Thus, the stocks with larger expected market impact (specifically Eniro) will not be considered for anything but for very large fractions of the NAV.

The analysis shows that the concavity of market impact (combined with the constraints) can lead to some extreme results. Notably, for about 25 % of the NAV, the algorithm suggests not liquidating any fraction of DSV. However, for about 30 % it suggests liquidating the DSV position completely. Considering the relatively non-robust model for market impact this is naturally an undesirable property. Nonetheless, the difference between the value of the objective function at different local minima is typically rather small. Both values are also significantly better than that of the benchmark strategy (proportional liquidation). Thus, the optimization shows that it is possible to significantly reduce the expected total liquidation cost by considering liquidation strategies different from proportional liquidation.

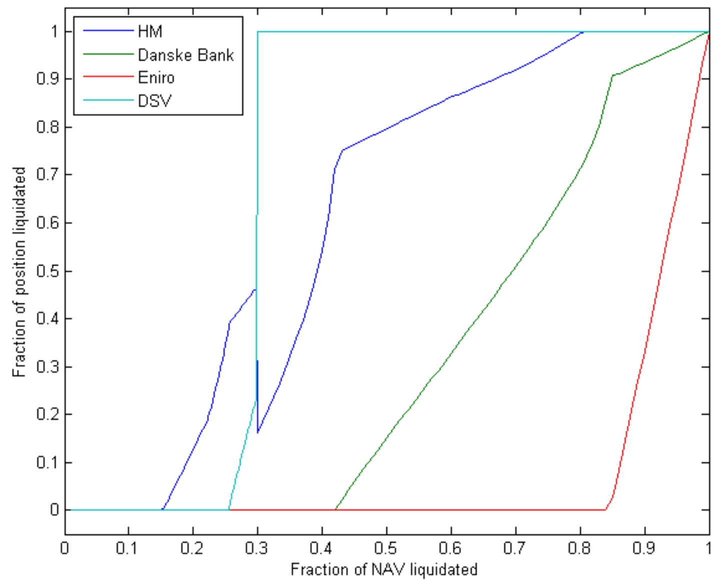

Figure 13: Order of liquidation: The figure shows the fraction of each stock position that the optimization algorithm suggests should be liquidated. It is plotted as a function of the fraction of the NAV that is to be paid out to investors.

7.4 Tracking error

In the case study above, a portfolio $V(0)$ at time 0 is, to some extent, liquidated during the time interval $(0, T)$ so that the new portfolio $V(T)$ has potentially very different properties. More specifically, any change in prices of the stocks and weighting in the portfolio will change the distribution of the (log) return of the portfolio.

At time T, the future return of the portfolio can be modeled as

$$V(T + \Delta) = V(T)e^{R^V(\Delta)}$$

where $R^V(\Delta)$ is the log return of the portfolio during the time interval $(T, T + \Delta)$. In the case study above I have considered performing the liquidation during one day so that $T = 1$. Furthermore, at time T I consider the distribution of the portfolio over the next day so that $\Delta = 1$. However, in the derivation below I will use T and Δ for notational clarity.

An explicit expression for this stochastic variable is given by

$$R^V(\Delta) = ln\Big(\frac{V(T+\Delta)}{V(T)}\Big) = ln\Big(\frac{h_0(T) + \sum_{i=1}^n h_i(T)S_i(T)e^{R_i(\Delta)}}{h_0(T) + \sum_{i=1}^n h_i(T)S_i(T)}\Big)$$

where $R_i(\Delta)$ is the log return of stock i. The expected stock prices $S_i(T)$ and portfolio weights $h_i(T)$ naturally depend on the specific liquidation strategy used during the time interval $(0, T)$.

Since the relative weighting (and thus distribution) of the pre-liquidation portfolio has been deliberately chosen by the portfolio manager, it is interesting to compare that distribution to that of the post-liquidation portfolio. More specifically, I consider two scenarios.

Scenario 1 (benchmark scenario): I consider the situation where no investors withdraw any money, no money is paid out and the portfolio weights remain unchanged. In this scenario no liquidation is performed and the changes in stock prices are only due to "normal" market movements. I make the assumption that, for short time intervals, the drift of any stock is negligable, i.e. $E[S_i(T)] \approx S_i(0)$. Thus, the log return $R^b(\Delta)$ of the *benchmark* portfolio $V_b(T)$ during the time interval $(T, T+\Delta)$ is given by

$$\begin{aligned} R^b(\Delta)|_{S_i(T)=S_i(0)} =& ln\Big(\frac{V_b(T+\Delta)}{V_b(T)}\Big) \\ =& ln\Big(\frac{h_0(0) + \sum_{i=1}^n h_i(0)S_i(0)e^{R_i(\Delta)}}{h_0(0) + \sum_{i=1}^n h_i(0)S_i(0)}\Big) \end{aligned}$$

Scenario 2 (liquidation scenario): In scenario 2, investors withdraw money, a fraction of the portfolio is liquidated (using some liquidation strategy) and stock prices and portfolio weights change from the liquidation. Denote by $S_i^p(T) = E[S_i(T)]$ the expected stock prices and by $h_i^p(T)$ the new portfolio weights with a given liquidation strategy. Then, the log return return $R^p(\Delta)$ of the portfolio $V_p(T)$ during the time interval $(T, T+\Delta)$ is given by

$$\begin{aligned} R^p(\Delta)|_{S_i^p(T), h_i^p(T)} =& ln\Big(\frac{V_p(T+\Delta)}{V_p(T)}\Big) \\ =& ln\Big(\frac{h_0^p(T) + \sum_{i=1}^n h_i^p(T)S_i^p(T)e^{R_i^p(\Delta)}}{h_0^p(T) + \sum_{i=1}^n h_i^p(T)S_i^p(T)}\Big) \end{aligned}$$

I compute the tracking error of the portfolio in scenario 2 relative the benchmark portfolio by using Definition 5. To do this, the distribution of the log returns $R^p(\Delta)$ and $R^b(\Delta)$ are needed. Furthermore, to estimate these distributions, the distributions of the stocks are needed. I make the assumption that historical log returns are representative of future log returns. More specifically, I consider the stochastic vector of log returns for the four stocks

$$\mathbf{R}(\Delta) = \big(R_1(\Delta), R_2(\Delta), R_3(\Delta), R_4(\Delta)\big)$$

and estimate the distribution of this vector with the historical outcomes

$$\mathbf{r}_i = \left(r_{1,i}, r_{2,i}, r_{3,i}, r_{4,i}\right)\ i = 1, ..., 145$$

I have 146 daily closing prices for the stocks and thus 145 daily log returns.

This is called an empirical distribution. The advantage with this approach is that it is easy to implement and that there is no need to statistically infer the distribution and dependence of the vector **R**. However, the success of the approach relies heavily on the quality and quantity of data. Introducing more data will increase the validity of any calculation made if the additional data is representative. However, increasing the data set typically means including outcomes from the more distant past, outcomes that might not have this property.

With this approach, the explicit expression for the tracking error is

$$\text{TE} = \sum_{i=1}^{145} \left(p_b(r) - p_p(r)\right)^2$$

where $p_b(r) = P(R_b(\Delta) = r)$.

Figure 14 shows the tracking error of the remaining portfolio relative the benchmark portfolio for the different liquidation strategies considered. I plot the tracking error as a function of the fraction of the NAV that is paid out to investors. One expects the proportional liquidation strategy to track the original portfolio perfectly. However, transaction costs will decrease the NAV and make the tracking error non-zero. The tracking error of the optimized liquidation strategy will be larger. Finally, the naive approach will have a very volatile distribution since log returns are calculated relative the NAV of the portfolio.

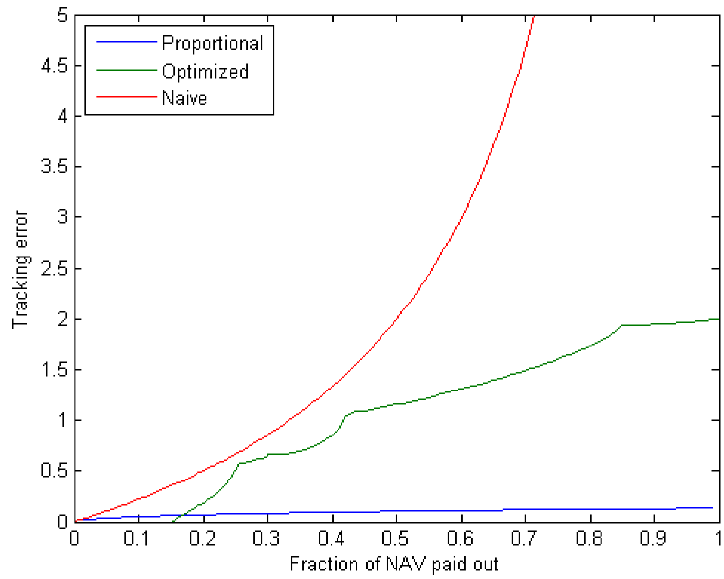

Figure 14: Tracking Error: The figure shows the tracking error of the remaining portfolio relative the original portfolio. The tracking error for the proportional approach becomes slightly positive due to transaction costs. The tracking error from optimized liquidation is significantly larger. The tracking error for the naive approach is asymptotically infinite.

7.5 Beta

I will also include an analysis of the beta of the remaining portfolio. Beta can be described as a volatility adjusted correlation between an asset (or a portfolio) and the overall stock market. The beta of the new portfolios are calculated relative the Nasdaq Omx Nordic index, using Definition 4. The returns of the remaining portfolios are calculated as in the section above and I use historical log returns for Nasdaq Omx Nordic index (of the corresponding day the corresponding day).

Figure 15 shows the beta of the remaining portfolio for the different liquidation strategies as a function of the NAV that is to be paid out to investors.

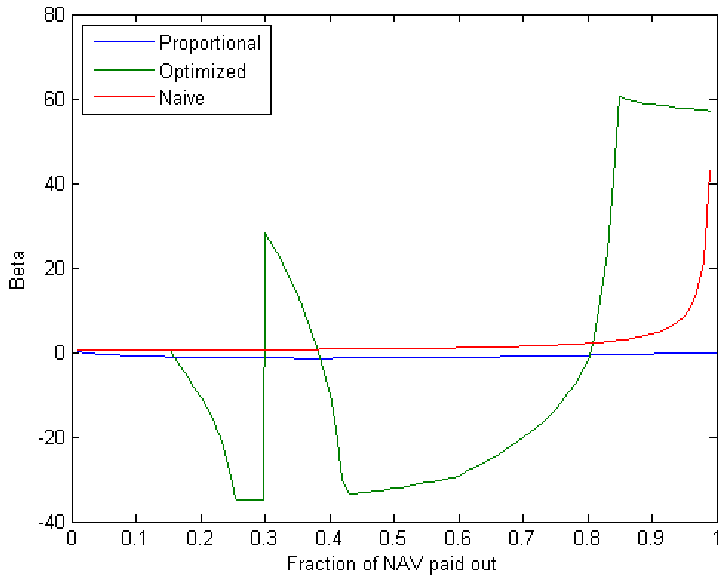

Figure 15: Beta: The figure shows the beta of the remaining portfolio (as a percentage of the NAV). The beta of the naive approach is asymptotically infinite. The beta of the optimized liquidation resembles that of the net exposure.

8 Conclusion

In this thesis I suggest a market impact function and use this function to compare different ways of liquidating a large fraction of a long-short equity portfolio. Specifically, I consider the situation where the portfolio belongs to a hedge fund and the liquidation is due to large withdrawals by investors. I have introduced three different liquidation strategies that I call the *naive approach*, *proportional liquidation* and *optimized liquidation*. The naive approach means paying investors without liquidating any stock positions, proportional liquidation means liquidating equal fractions of each stock position and optimized liquidation is the liquidation strategy found by the solving an optimization problem.

I find that the naive approach implies no market impact but significantly increases the risk of the remaining portfolio. Proportional liquidation keeps the distribution of the remaining portfolio intact (relative the pre-liquidation portfolio) but implies a relatively high cost due to market impact. Optimized liquidation has a smaller expected cost than proportional liquidation and keeps the risk of the remaining portfolio within pre-specified limits. However, optimized liquidation tends to change the distribution of the (log) return of the portfolio. The results show that the naive approach is undesirable unless very small volumes are considered. As for the other two approaches, the results in no way indicate that any one approach is superior to the other. Rather, which strategy to choose depends on the preferences of the manager in charge of the liquidation. The decision boils down to a choice between decreasing expected costs and keeping the relative weighting and overall properties of the portfolio intact.

The results depend heavily on the validity of the market impact function suggested in the thesis. In particular, I suggest that market impact is a concave function of the traded volume which implies a decreasing marginal cost of trading. This property further implies that proportional liquidation is relatively expensive and that the optimization algorithm tends to suggest trading in as few positions as possible. This naturally implies large changes in the relative weighting of the portfolio. There is ample evidence of the concavity of market impact and this part of the model (and the results derived) is robust.

Apart from the dependence on the traded volume there is little consensus about how to model market impact, in particular with which parameters to scale the function and how it depends on the trading velocity. Because of this the actual values suggested as expected costs, although reasonable, might not be perfectly accurate.

8.1 Future work

The lack of robust empirical results call for the need of more empirical studies. The possibilities to obtain relevant data will likely increase with the increased use of electronic trading. The dependence on the trading velocity is especially interesting. This will allow modeling the expected cost of trading as a function of time and then compare this with the increased market risk of holding the position over a longer time interval. Finally, investigating other ways of formulating the optimization problem would address portfolio managers with other preferences. In particular, using the tracking error as the objective function and the expected cost as a constraint would appeal to anyone who prioritizes keeping the relative weighting of the portfolio intact.

References

[1] Embrechts, P., R. Frey, and A. J. McNeil, 2005, Quantitative Risk Management, Princeton University Press, Princeton

[2] Engle, Robert F., Ferstenberg, Robert and Russell, Jeffrey, Measuring and Modeling Execution Cost and Risk (July 1, 2008). Chicago GSB Research Paper No. 08-09. Available at SSRN: http://ssrn.com/abstract=1211162

[3] Hausman, Jerry A., Lo, Andrew W. and Mackinlay, A. Craig Craig, An Ordered Probit Analysis of Transaction Stock Prices (October 1991). NBER Working Paper Series, Vol. w3888, 1991. Available at SSRN: http://ssrn.com/abstract=467658

[4] F. Lillo, R. Mantegna, J. D. Farmer, Single Curve Collapse of the Price Impact Function for the New York Stock Exchange, e-print cond-mat/0207428; Master curve for price-impact function, Nature, 421, 129 (2003)

[5] Farmer, J. Doyne, Patelli, Paolo and Zovko, Ilija I., The Predictive Power of Zero Intelligence in Financial Markets (February 9, 2004). AFA 2004 San Diego Meetings. Available at SSRN: http://ssrn.com/abstract=483603

[6] C. Hopman, Are supply and demand driving stock prices?, MIT working paper, Dec. 2002

[7] Potters, Marc and Bouchaud , Jean-Philippe , More Statistical Properties of Order Books and Price Impact (October 31, 2002). Available at SSRN: http://ssrn.com/abstract=507343

[8] N. Torre. BARRA Market Impact Model Handbook. BARRA Inc., Berkeley, 1997.

[9] Gabaix, Gopikrishnan, Plerou Stanley - A theory of power-law distributions in financial market fluctuations, Nature Vol. 423, 2003

[10] Moro, Esteban, Vicente, Javier, Moyano, Luis G., Gerig, Austin, Farmer, J. Doyne, Vaglica, Gabriella , Lillo, Fabrizio and Mantegna, Rosario N., Market Impact and Trading Prole of Large Trading Orders in Stock Markets (August 3, 2009). Available at SSRN: http://ssrn.com/abstract=1443122

[11] R. Almgren, C. Thum, E. Hauptmann, and H. Li, Direct estimation of equity market impact, Journal of Risk, 18 (2005), p. 57

[12] Keim, Donald B. and Madhavan, Ananth, The Upstairs Market For Large-Block Transactions: Analysis And Measurement of Price Effects (May 1994). Available at SSRN: http://ssrn.com/abstract=5321

[13] Cont, Rama, Kukanov, Arseniy and Stoikov, Sasha, The Price Impact of Order Book Events (March 1, 2011). Available at SSRN: http://ssrn.com/abstract=1712822

[14] J. D. Farmer, L. Gillemot, F. Lillo, S. Mike, and A. Sen. What really causes large price changes? Quantitative Finance, 4(4):383-397, 2004

[15] F. Lillo and J. D. Farmer. The long memory of the efficient market. Studies in Nonlinear Dynamics and Econometrics, 8(3), 2004

[16] Bouchaud , Jean-Philippe , Gefen, Yuval, Potters, Marc and Wyart, Matthieu, Fluctuations and Response in Financial Markets: The Subtle Nature of 'Random' Price Changes (July 14, 2003). Available at SSRN: http://ssrn.com/abstract=507322

[17] Almgren, R., and N. Chriss. 2000. Optimal Execution of Portfolio Transactions. Journal of Risk 3:539.

[18] Gatheral, Jim, No-Dynamic-Arbitrage and Market Impact (September 28, 2009). Quantitative Finance, Vol. 10, No. 7, pp. 749-759, 2010. Available at SSRN: http://ssrn.com/abstract=1292353

[19] Grinold, R.C. and R.N. Kahn (1999). Active Portfolio Management: A Quantitative Approach for Providing Superior Returns and Controlling Risk. McGraw-Hill Professional Publishin

[20] Bikker, J.A., L. Spierdijk, and P.J. van der Sluis (2007). Market Impact Costs of Institutional Equity Trades. Forthcoming Journal of International Money and Finance

[21] Scherer, B. (2004) Portfolio Construction and Risk Budgeting (2nd ed.). London: Risk Books

[22] J. Nocedal and S. J. Wright. Numerical Optimization. Springer, 1999

Appendix

Proof of Corollary 1

Under the assumption that the *loss variable L* is normally distributed, the expression for the Value-at-Risk can be simplified. The proof is based on introducing the standard normal variable Z and noting the following equality (in distribution)

$$L^{\Delta} = \mu_L + \sigma_L Z$$

Proof 1.

$$\begin{aligned}
\mathit{VaR}_{\alpha,\Delta} =& \inf\{l \in \mathbb{R} : P(L^{\Delta} > l) \leq 1 - \alpha\} \\
=& \inf(l \in \mathbb{R} : P(\mu_L + \sigma_L Z > l) \leq 1 - \alpha) \\
=& \inf(l \in \mathbb{R} : P(Z > \frac{l - \mu_L}{\sigma_L}) \leq 1 - \alpha) \\
=& \{l' = \frac{l - \mu_L}{\sigma_L}\} \\
=& \inf(\mu_L + \sigma_L l' \in \mathbb{R} : P(Z > l') \leq 1 - \alpha) \\
=& \mu_L + \sigma_L \inf(l' \in \mathbb{R} : P(Z > l') \leq 1 - \alpha) \\
=& \mu_L + \sigma_L \Phi(1 - \alpha)
\end{aligned}$$