Liquidity Risk in Limit Order Book Markets

Inauguraldissertation
zur Erlangung des Doktorgrades

der Wirtschafts– und Sozialwissenschaftlichen Fakultät
der Universität zu Köln

2006

vorgelegt von
Dipl.-Kfm. Daniel Mayston
aus Köln

Referent: Prof. Dr. Alexander Kempf

Korreferent: Prof. Dr. Heinrich Schradin

Abgabetermin der Arbeit:

02. Mai 2006

Contents

List of Figures

List of Tables

List of Abbreviations

BE	Book equity
CAPM	Capital Asset Pricing Model
CET	Central European time
DAX	Deutscher Aktienindex
ECN	Electronic communications network
EEX	European Energy Exchange
Euribor	European Interbank Offered Rate
Fibor	Frankfurt Interbank Offered Rate
FOK	Fill or kill
FSE	Frankfurt Stock Exchange
GARCH	Generalized Autoregressive Conditional Heteroscedasticity
GLS	Generalized Least Squares
IOC	Immediate or cancel
LR	Litzenberger Ramaswamy
LSE	London Stock Exchange
MDAX	MidCap DAX
ME	Market equity
NASDAQ	National Association of Securities Dealers Automated Quotations
NSC	National Security Council
NYSE	New York Stock Exchange
OLS	Ordinary Least Squares
PCA	Principal component analysis
PIN	Probability of Informed Trading
SETS	Stocks Exchange Trading System (in London)
SUR	Seemingly unrelated regressions
TecDAX	Technology DAX

List of Symbols

Symbols in Chapter 2

A, B	Ask and bid side subscripts
E	Subscript for buy or sell executions ($+E$, $-E$)
I, P	Subscripts for patient and impatient traders
Q_l	Quantity of limit orders in the book at price l (cumulated)
T	Time until order execution
U	Utility
X	True security value
Y	Future asset payoff
b	Marginal rate of substitution
c	Consumption
d	Innovation of new information
f	Index of the patience of traders
g	Number of spreads in equilibrium
h	Index of equilibrium spreads
i	Firm or stock index
j	Index of tick difference between market and limit orders
k	Number of prices on the ask and bid side ($+k$, $-k$)
l	Index of prices on the ask and bid side ($+l$, $-l$)
m	Market order volume
n	Equilibrium spread
p	Price
q	Quantity of limit orders with time and price priority
s	Bid-ask spread
t	Time index

α	Price impact of market orders
β	Time preference
δ	Waiting cost per unit of time
γ	Order processing cost
λ	Class parameter (for Poisson distribution)
μ	Drift component
π	Profits
Δ_h	Spread improvement from spread n_h to n_{h-1}
Φ	Conditioning variables
Θ_f	Proportion of traders f in the population

Symbols in Chapter 4

A, B	Ask and bid side subscripts
AR	Order arrival rate
BF	Beta Factor
DEP	Depth
END	End of the trading day
IP	Informed trader profits
L_i, L_M	Stock and market liquidity
MC	Market capitalization
MO	Market order volume
MQ	Midprice
nCA	Number of cancellations
nLO	Number of limit orders
nMO	Number of market orders
OI	Order imbalance
PAT	Proportion of Patient Traders
SPR	Bid-ask spread
TRV	Trading volume
$UNXV$	Unexpected volatility
VOL	Return volatility

h	Conditional volatility
i	Firm index
k	Index of ticks in the order book (starting at $k = 0$)
l	Index of prices in the limit order book (starting at $l = 0$)
n	Number of shares
P	Price
r	Return
t	Time index
$\alpha, \beta, \delta, \gamma$	Regression coefficients
ϵ	Error term
κ	Mean reversion parameter
μ	Conditional mean of the return
φ	Mean reversion coefficient
Δz	Stochastic increment
Θ	Long-run mean of liquidity

Symbols in Chapter 5

A, B	Ask and bid side subscripts
$DEP_{A,B}$	Depth at the best bid and ask price
L_i, L_M	Stock and market liquidity
M	Market subscript
MQ	Midprice
PC	Principal component
PI	Price impact (relative half-spreads)
RS	Relative spread
VOL	Conditional volatility (squared returns)
d	Index of the trading day
h	Index of the time of day
i	Firm index
l	Index of prices in the limit order book

n	Number of shares
p	Price
r	Return
t	Time index
w	Index of rolling ten-day windows
x	Volume in the limit order book
$\alpha, \beta, \delta, \eta, \Psi, \xi$	Regression coefficients
ε	Error term
μ_h	Time-specific mean
σ_h	Time-specific standard deviation
λ	Linear price impact coefficient
ρ	Quadratic price impact coefficient

Symbols in Chapter 6

BE	Book equity
BM	Book-to-market factor
C	Variance covariance matrix
$DEP_{A,B}$	Depth at the best bid and ask price
DEP_k	Depth at the k-th tick (cumulated)
HS	Half spread
L	Liquidity
$LLEV$	Level of liquidity
$LRES$	Resiliency factor
$LSYS$	Commonality factor
M	Subscript for the market
ME	Market equity
R	Excess return
R	Monthly return
$SIZE$	Size factor
SPR	Bid-ask spread
T	Index of test years

VOL	Conditional volatility (= squared return)
W	Weighting matrix
i	Firm index
j	Index of pricing coefficients
k	Tick index
l	Index of prices in the limit order book
p	Portfolio index
r_i	Stock return
r_f	Risk-free rate of return
r_M	Market return
s	Superscript of stacked vectors and matrices
t	Time index
α, β, δ, γ, ξ	Regression coefficients
$\hat{\beta}_i$	Pre-ranking stock beta
$\hat{\beta}_k$	Portfolio beta
λ	Linear price impact coefficient
φ	Mean reversion coefficient
ε	Error term

Symbols in the Appendix

C	Variance covariance matrix
X	Data matrix
Z	Standardized X data matrix
i	Index of columns in the data matrix
j	Index of eigenvectors and eigenvalues
x	Column in X
z	Column in Z

λ	Eigenvalue
Λ	Diagonal matrix of sorted eigenvalues
γ	Eigenvector
Γ	Matrix of sorted eigenvectors

Operators and Functions

E	Expectation
$h(\cdot)$	Price impact function
$I(\cdot)$	Indicator function
int+	Next higher integer value
max	Maximum
min	Minimum
Pr	Probability
Δ	First difference
σ	Standard deviation
Σ	Sum

Chapter 1

Introduction

The liquidity of a financial security characterizes the speed and ease with which any quantity can be purchased or sold. A liquid asset can be traded quickly and without large price effects. From this point of view, liquidity is desirable for any investor who wishes to buy or sell a security. While liquidity is widely accepted as a component of transaction costs, the risk that arises from illiquidity has received very little attention in the literature of finance. Most valuation models like the standard CAPM or the Black Scholes option pricing formula even assume frictionless markets with perfectly liquid assets. However, this assumption can be very dangerous. For example, in the 1987 stock market crash and the 1998 bond market failure (see Figure 1.1) liquidity drained from the market dramatically, which made it difficult to trade at all. The fall of the Long-Term Capital Management (LTCM) hedge fund, which was strongly exposed to liquidity risk in bond markets, highlights the impact that liquidity risk can have on portfolios.

Figure 1.1: Stock and Bond Market Crashes

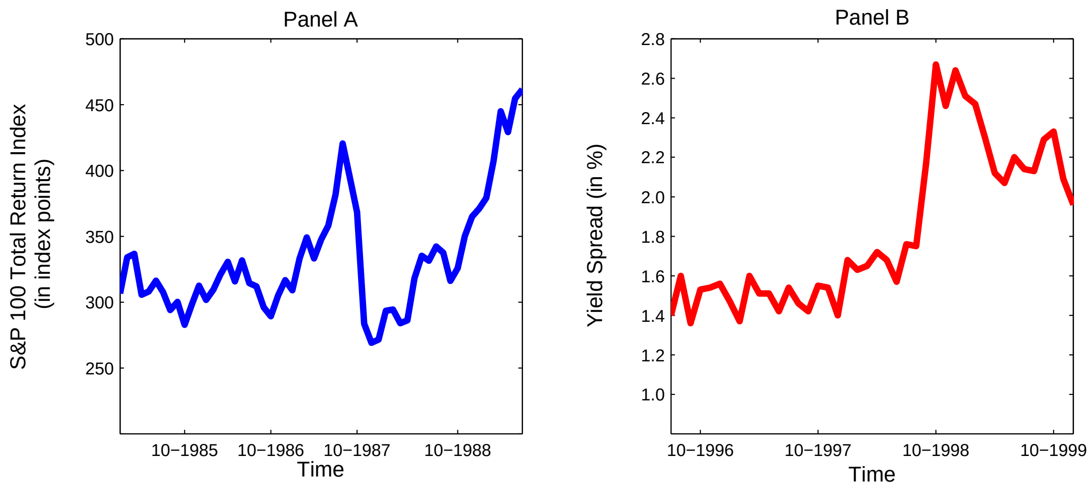

Figure 1.1 shows the stock and bond market crashes in 1987 and 1998. Panel A plots the value of the S&P100 Total Return Index (in index points) which fell by almost 50% in the autumn of 1987. Panel B plots the yield spread of Moody's BAA bond index over US government bonds with a constant 10 year maturity (in %). As a result of Russia's default, the spread increased dramatically.

Against this background I investigate the properties, magnitude and importance of liquidity risk in today's electronic limit order book markets.

1.1 Key Issues and Relevance

Nowadays all large stock exchanges like New York, London or Frankfurt and all large electronic communication networks (ECNs) such as Island or Instinet are organized as electronic trading facilities. While some trading venues also have market maker features, they all operate on the basis of an open limit order book. Therefore I focus on an investor's liquidity risk in the context of electronic limit order book markets.

To understand the importance of liquidity risk for investors, imagine that an investor who has a stock portfolio suddenly needs to liquidate some positions to meet unexpected cash requirements. If the securities are illiquid at the time, their liquidation will be very costly. Therefore the investor will have to sell more than intended originally to meet his liabilities. This scenario shows that liquidity risk typically constitutes the danger that prices deteriorate heavily in response to trades. In such situations, securities can only be traded at unfavorable prices if at all. Motivated by the demonstrated importance of liquidity risk, I address the following three aspects of liquidity risk in my thesis:

1. How fast does a limit order book refill after liquidity has been taken away?

2. How strongly does liquidity risk spill over across different stocks?

3. To what extent does liquidity risk enter stock prices as a priced factor?

The first question is usually referred to as resiliency. It describes the extent to which new liquidity flows back to the market after the order book has been cleared. The second question addresses the co-movement of liquidity over time, generally referred to as commonality in liquidity. The third question investigates the link between liquidity risk and asset pricing to establish whether liquid assets realize higher prices than their less liquid counterparts. Together, the resiliency, commonality and pricing dimension give a comprehensive picture of liquidity risk in limit order book markets.

To illustrate the economic relevance of liquidity risk, let us have a brief look at the LTCM case. After a strong decline in liquidity due to the Russian debt crisis, LTCM's portfolio value dropped dramatically (see Figure 1.2). In turn, this triggered off margin calls that had to be met. The fund was forced liquidate large parts of the portfolio in an environment in which it was very expensive to sell off assets. Firstly, the markets for the individual securities were not resilient, which meant that new liquidity did not flow back into the market. The unwinding of large positions was accompanied by strong adverse price movements. Secondly, LTCM found itself amidst a market-wide liquidity crisis. The strong commonality of liquidity prohibited any protection through diversification effects of liquidity risk across instruments. The combination of low resiliency and market-wide illiquidity forced LTCM to its knees so strongly that a group of financials institutions led by the Federal Reserve Bank of New York bailed the hedge fund out to avoid complete bankruptcy. The downfall of the LTCM is an acute illustration that investors are well advised to integrate liquidity risk into their risk management. Before the downfall investors earned well from investing in LTCM which reflects that the market compensates investors for taking on liquidity risk. However, high liquidity risk also implies a higher probability of losses – something that occurred very dramatically in the case of LTCM.

The relevance of liquidity risk at the market-wide level comes from the fact that liquidity serves as the lifeblood of financial markets. It enables the translation of information into order flow and prices and thereby promotes the stability of the trading environment. A sudden drop of liquidity in a certain segment or

Figure 1.2: Rise and Fall of $1 Invested with LTCM

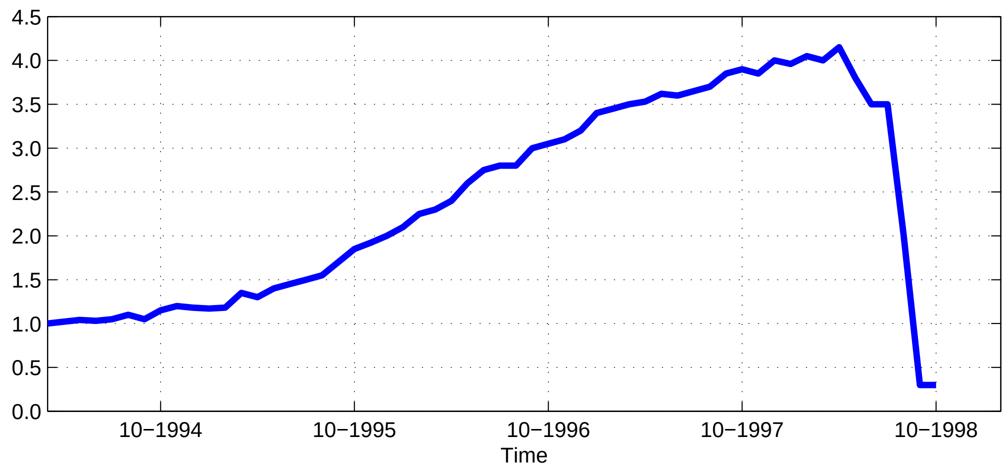

Figure 1.2 shows the gross value of one dollar invested with LTCM from March 1994 to October 1998. LTCM's funds had a value of $ 4.8 billion in April 1998 of which more than $ 4 billion got wiped out in just a few months because liquidity dried out. The figure follows Lowenstein (2000, p. XV) who gives a detailed account of the rise and fall of the infamous hedge fund.

geographic region can potentially spill over to further segments and countries to lead market instability on a larger scale. Past crises like in Russia or Indonesia show that the liquidity of whole regions and markets can dry out and lead to very destabilizing effects. Therefore, the less resilient and the more systematic liquidity risk is across stocks, the larger the potential for market-wide disruptions. A deeper understanding of liquidity risk will be very desirable for legislators, regulators, exchanges and financial institutions to enhance the stability and smoothness of trading in financial markets.

1.2 Contribution to the Literature

The initial interest in the microstructure of security markets and the liquidity of financial assets is often traced back to Demsetz (1968) and Garman (1976). From then on, the literature has produced an abundance of models of liquidity in security markets. While too numerous to list exhaustively, the most prominent approaches include the inventory models of Stoll (1978) and Ho and Stoll (1981), the asymmetric information models of Glosten and Milgrom (1985) and Easely and O'Hara (1987) and the models of strategic trading as in Kyle (1985). Their theoretical extensions and empirical tests also make up a large body of the liquidity literature. While these approaches address the emergence of liquidity, their extension to liquidity risk – models and tests in which liquidity is stochastic – has only just begun.

The first pillar of liquidity risk that I consider is the resiliency of liquidity. According to Garbade (1982) a market is resilient if price changes that result from high order volumes quickly attract new limit orders which, in turn, pull the price back again. In empirical applications, Holthausen, Leftwich and Mayers (1987) and Chordia, Roll and Subrahmanyam (2005) study the resiliency of prices, yet they do not consider liquidity or use liquidity measures, either. Coppejans, Domowitz, Madhavan (2003) and Gomber, Schweickert and Theissen (2004) analyze the dynamic properties of the limit order book, yet they only focus on the level of liquidity or half-life measures. Degryse, Jong, Ravenswaaij and Wuyts

(2005) capture resiliency indirectly by the aggressiveness of orders, yet they do not model the refreshment process explicitly. Large (2005) sets up a model of the probability that prices return to pre-trade levels, yet does not distinguish the refreshment process that takes place in the limit order book. In contrast, I extend the literature by directly implementing the Garbade (1982) definition of resiliency. I set up a mean reversion model of liquidity that captures the change in current liquidity in response to past liquidity. Furthermore, I interact the liquidity changes with microstructural determinants to examine their impact on the resiliency mechanism of the limit order book.

The second pillar of liquidity risk in my thesis is the stochastic covariation of liquidity across assets. Chordia, Roll and Subrahmanyam (2000) introduced the idea of market-wide liquidity in an empirical study of US quote data. They provide evidence that market liquidity has a significant impact on individual stock liquidity. Brockman and Chung (2002) apply this approach to intraday data from Hong Kong's order-driven stock market with very similar results. Halka and Huberman (2001) document correlation in the liquidity of different stock portfolios. Hasbrouck and Seppi (2001) on the other hand find little evidence of commonality once deterministic time-of-day effects have been removed. The mixed results and low levels of commonality do not make the current literature very persuasive.[1] An obvious shortcoming is the confinement to very narrowly

[1]The theoretical literature has developed several mechanisms through which the liquidity supply of different stocks is linked. In these models, the correlation of liquidity preferences, intermediary behavior or informational shocks create contagion effects in the liquidity of different stocks (see Allen and Gale (2000), Kyle and Xiong (2001), Gromb and Vayanos (2002), Fer-

defined measures of liquidity that do not do justice to liquidity risk in limit order markets. Therefore I advance the study of commonality by modeling the liquidity of a limit order book. The section is most closely related to the work of Bauer (2004) and Domowitz, Hansch and Wang (2005). Bauer (2004) performs a principal component analysis of the liquidity in the limit order book across stocks. Domowitz, Hansch and Wang (2005) investigate the influence of order flow and order type correlations on liquidity commonality. The most important difference is that I focus on how commonality in liquidity depends on how deep I look into the limit order book.

The third pillar of liquidity risk that I examine is its impact on expected returns. The first empirical studies examined the relation between the level of liquidity and expected stock returns (see Amihud and Mendelson (1986), Amihud and Mendelson (1989), Eleswarapu (1997), Brennan and Subrahmanyam (1996), Brennan, Chordia and Subrahmanyam (1998) or Amihud (2002)). The next generation of studies focused on the pricing of liquidity risk as opposed to the level alone. They include Pástor and Stambaugh (2003), Gibson and Mougeot (2004) and Acharya and Pedersen (2005). Together these studies provide positive evidence that the level and risk of liquidity get priced. They all consider liquidity movements over long horizons (mostly monthly frequencies). However, liquidity adjustments in limit order book markets are phenomenona that take place and can

nando (2003), Watanabe (2003) and Brunnermeier and Pedersen (2005)). In empirical studies, Coughenour and Saad (2004) relate commonality in liquidity to common market maker behavior. Domowitz, Hansch and Wang (2005) argue that the correlation of order type choice is the reason for the correlation of liquidity.

only be observed acurately within hours or even within minutes. Furthermore, most studies do not differentiate between level and risk effects in liquidity. I therefore extend the literature to the pricing effects of liquidity measures that are based on limit order book data. In addition, I separate level and risk effects of liquidity and estimate their impact on returns simultaneously.

In all I bring together the diverse aspects of stochastic liquidity in an attempt to give a comprehensive view of liquidity risk. All of the above issues are applied to limit order book data from a purely quote-driven limit order book market. I use three months of Frankfurt Stock Exchange's (FSE) electronic protocol which keeps record of all events that took place in their Xetra trading system. With the help of some substantial computer programming that implemented the trading rules of the Xetra system I was able to reconstruct the limit order book from the raw data for any point in time.[2] As the blue-chip segment of FSE has no additional liquidity supply by specialists and does not face any notable competition from regional exchanges, the order book data enables a clinic view of the liquidity risk of financial securities in electronic limit order book markets.

[2] I thank Deutsche Boerse for the electronic trading protocols and initial order book. The data was supplied and initially prepared in SAS. The subsequent reconstruction programming took place Gauss. I gratefully acknowledge the support of Helena Beltran-Lopez (Université Catholique de Louvain), Joachim Grammig and Stefan Frey (Universitaet Tuebingen) who shared large parts of the programming sequences. The later data construction and econometric programs were written in Matlab and Eviews. They are disclosed in part in Appendix C and available upon request.

1.3 Main Results and Procedure

In all, the study of the resiliency of the liquidity supply, the commonality in liquidity and the pricing effects of liquidity risk yield the following results:

- **The refill mechanism of the limit order books is strong**. The findings suggest that in general resiliency is high. Empty order books are refilled promptly, which restores the normal level of liquidity reasonably fast. Resiliency is stronger if trading is high, yet if volatility is high, liquidity only increases for investors who wish to buy whereas sales become more difficult. Informed trading has a weak impact: evidently, liquidity suppliers cannot anticipate the information content of trades in anonymous limit order book markets as they cannot identify the traders behind individual trades.

- **The liquidity co-movement in the order book is substantial**. Comprehensive measures of limit order book liquidity as opposed to measures at the best bid and ask price exhibit substantial co-movement: systematic movements across stocks make up about 20% of their overall movement for measures beyond best prices, while the systematic component for measures at the best price is only 2%. These figures underline that market-wide liquidity movements are too large to be neglected. Furthermore, commonality is strongly time-varying: while commonality is lower in rising markets, it increases considerably in falling markets.

- **Liquidity levels and liquidity risk are priced factors**. Both liquidity levels and liquidity risk get incorporated into stock returns. The higher the level of stock liquidity is, the lower the expected return that an investor will receive from buying that stock. Likewise, the higher the liquidity risk of the stock is, the higher the return paid on the investment. The main implication of these results is that, evidently, investors pay attention to the tradability of stocks.

I proceed as follows: in Chapter 2, I present the basic microstructure theory of securities markets and model the liquidity supply in a limit order book market. In Chapter 3, I give some details on the market structure, the data that I use in the empirical sections and some descriptive statistics. In Chapter 4, I examine the resiliency of the limit order book. Chapter 5 addresses systematic liquidity risk by focusing on the common movement of liquidity over time. Chapter 6 deals with the pricing dimension of liquidity levels and liquidity risk. It incorporates liquidity factors into a Fama-MacBeth framework to study their impact on stock returns. Chapter 7 concludes.

Chapter 2

Liquidity in Limit Order Book Markets

The first section of this chapter gives an overview over the microstructure of financial markets. In the second section I present a static of model of the limit order book, while the third section develops a dynamic model of the limit order book. In the fourth section I conclude how my further empirical research builds on the theoretical literature.

2.1 Market Microstructure

Short-run price dynamics are an important force that drives the liquidity of assets and markets. To understand this process fully, the following section explains price formation on security markets. It shows how short-term price dynamics are embedded in a more long-run valuation process and how short-lived deviations from information efficiency are related to liquidity. Finally I apply these notions to the limit order book.

2.1.1 Price Formation with Market Frictions

One of the most prominent fields of modern finance is investment theory. Investment models deal with the equilibrium value of financial assets. They are mostly set in a world in which markets are frictionless and efficient. In such markets, asset prices reflect all available information and thus correspond to their expected true values at any point in time. However, a second implication of frictionless markets is that assets are perfectly liquid. Unfortunately, that is not the case in the real world. This discrepancy is what the literature of market microstructure addresses.

Market microstructure owes its name to Garman (1976) who defines a market's microstructure as the interaction of individual exchange actions along time that, in aggregation, make up the market. Loosely speaking, we can think of market microstructure theory as the field that deals with the actual mechanics of financial markets. In particular, it explicitly introduces market frictions. A central feature is that most microstructure models are characterized by a multitude of prices: agents who offer to trade propose bid prices for sales and ask prices for buys. Market participants who trade against these offers realize so-called transaction prices. Price data in the media and press usually reports midquotes, the midpoint between the best bid and ask price.

A useful way of distinguishing the investment and microstructure view is by their time horizons. The investment view considers long-run price dynamics. It focuses on the asset's value in the long run which it derives from fundamental factors about the company. In contrast, the microstructural view considers short-run price dynamics which involve elements of the trading process itself such as the order size, trading activity or market mechanism. In a general sense, microstructural price dynamics are short-run price disturbances around a long-term value process – Hasbrouck (2004) uses the term overlay component.[3]

To illustrate the link between market efficiency and microstructure prices, let us use a slightly more formal representation. Asset pricing theory states that, for a market to be efficient, prices have to be martingales.[4] The martingale property implies that returns are serially uncorrelated and that prices reflect expectations at all points in time. In a consumption-based asset pricing model, Cochrane (2001) shows how under what conditions prices are martingales. Let c_t denote an agent's consumption in t and let u denote utility from period consumption and let U denote overall utility. An agent's overall utility then depends on current and future consumption,

$$U(c_t, c_{t+1}) = u(c_t) + \beta \cdot u(c_{t+1}), \tag{2.1}$$

where β is a time preference parameter. If p_t denotes an asset's price, Y_{t+1} is its future payoff value. The future payoff corresponds to the sum of the price in $t+1$ and dividends d, $Y_{t+1} = p_{t+1} + d_{t+1}$. If agents maximize utility, the price of the

[3] The view of microstructure as a temporary noise component over a more fundamental value process might explain why it is often considered as a second order effect and why microstructure risk was neglected for a long time.

[4] A martingale is a stochastic process for which, conditional on today's information, today's price is the best prediction for tomorrow's price. The random walk model is a popular version of such a martingale process.

asset is

$$p_t = E_t[\beta \, \frac{u'(c_{t+1})}{u'(c_t)} \, Y_{t+1}]. \tag{2.2}$$

An important assumption is that for short time horizons time preference is negligible, $\beta = 1$, and there are no dividends, $d_{t+1} = 0$. Denoting the marginal rate of substitution as $b_t := \frac{u'(c_{t+1})}{u'(c_t)}$, we can under these assumptions set $Y_{t+1} = p_{t+1}$ and write the expected future price as

$$p_t = E_t[b_t p_{t+1}]. \tag{2.3}$$

Under the risk-neutral probability measure $b_t = 1$ and $E_t[p_{t+1}] = p_t$. Thus, under a fairly simple set of assumptions, prices in frictionless markets are martingales.[5] However, when we have several prices (for example for buy and sell orders), the price process is not a martingale any longer. Hasbrouck (2004) shows that for bid and ask prices, Equation 2.3 becomes

$$p_t^B \leq E_t[b_t Y_{t+1}] \leq p_t^A \tag{2.4}$$

and that the bid-ask spread prevents returns from being serially uncorrelated. In other words, prices are not equal to expectations at all points in time and, subsequently, the market is not informationally efficient. Roll (1984), Stoll (1989), George, Kaul, Nimalendran (1991), Huang and Stoll (1997) and Madhavan and Sofianos (1997) present empirical evidence for short-run return predictability that

[5] A large part of the central predictions of asset pricing theory about risk-neutral valuation goes back to the work of Ross (1976) and Ross (1978). Harrison and Kreps (1979) extended this to a more formalized theory embedded in a martingale grounding.

arises from microstructure phenomena. This result seems to imply that differing bid and ask prices lead to return predictability and therefore are not consistent with the notion of market efficiency which requires that prices are not predictable. However, most models in the microstructure literature reconcile market efficiency and microstructural frictions by including an unobservable true value of an asset. Market participants do not observe the true value, but they receive signals Φ_t which they use to form expectations of the true value, $E_t[Y_{t+1}|\Phi_t]$. These expectations are martingales again. Short-run prices can deviate from expectations, but in the long run expectations ensure that prices are efficient.

The decomposition of microstructure prices into information efficient and non-efficient components leads to another important concept of the microstructure studies: liquidity. The fact that expected values and actual transaction prices need not be the same means that some traders pay more or receive less than the asset is really worth. These costs can vary from stock to stock. Liquidity summarizes differences in prices of the same asset that arise from short-run divergences of transaction prices from the expected true value of the asset. In very general terms, assets are considered liquid if reasonable quantities can be traded at prices p_t^A and p_t^B that are not too far from the true value, Y_t. However, as O'Hara (1995) and many others point out, this is a very blurry concept. What is a reasonable quantity? When is a price too far from the true value?

While the above models of price dynamics and liquidity serve as a good framework, the actual price process, asset liquidity and subsequent trading costs depend on further issues such as the institutional environment of trading. The most common trading mechanisms nowadays are auction markets. Since the introduction of electronic trading platforms, traders are able to trade continuously when the exchange is open. This is a so-called continuous double auction system. It can be either quote-driven – a market maker supplies liquidity – or order-driven – limit order traders are the only source of liquidity. In practice, markets will often have hybrid structures which combine different elements and systems.

2.1.2 Elements of a Limit Order Book Market

In practice, nearly all markets operate on the basis of an electronic limit order book. Most markets have some distinct features with regard to priority of order execution, competition or transparency. For example, NYSE operates a hybrid system of a limit order book combined with specialists. Paris Bourse and LSE only disclose a fixed number of orders in the book. FSE on the other hand displays each single price in the order book and has no market makers. However, all limit order markets share a set of basic construction elements and rules. Firstly, they all build on the archetypical order forms of limit and markets orders:

- Market orders carry no price limit and are executed immediately against orders in the order book. In the rare case that the book is not liquid enough, they are executed partially and the remaining order lot is executed at the next possible point in time.

- Limit orders carry a pre-specified limit for execution. They are added to the limit order book and are executed against new incoming market orders as soon as possible. If they are larger than the market order, they are executed partially and the non-executed lot remains in the book.

The most important difference is that market orders have execution certainty, while the execution of limit orders is uncertain. Handa and Schwartz (1996) point out that submitting a limit order has similarities with an option. For example, an investor who places a buy order writes a free put option to the market. Let us assume that the investor places a buy limit order at 100 Euros. If the share price falls below 100 Euros, any trader in the market can hit the limit order at 100 Euros and the option will be executed. This implies that the writer of the option buys the stock above the market price. The lower the market price is, the higher are the losses that result from the limit order. On the other hand, if the market price rises, no trader will execute against the limit order because the stock can be sold at a higher price. Therefore, the profits from the limit order are zero if the stock price rises above 100 Euros. In all, the payoff is negative if the market price falls below the limit price and zero if the market price rises above the limit price. This payoff is identical to a put option. As market participants decide whether to execute against the option, limit order placement is equivalent to a short position in the put. By the same analogy, an investor who places a sell limit order writes a free call option to the market.

The option characteristic of limit orders shows that limit orders supply liquidity by enabling market participants to buy or sell at the prices in the order book. Limit order traders take over the role that the market maker has in a dealer market. However, some important differences remain between the monopolistic market maker model and a limit order market: firstly, limit order traders compete amongst each other for the supply of profitable liquidity. Secondly, they know neither the size of incoming market orders nor the identity of the trader. Therefore, limit order traders have less power to infer the information content of trades than a market maker has in dealer market where he alone sets the quotes.

The differences between market maker quotes and limit orders have implications for the objectives with which limit orders are used. Traditionally, liquidity supply only takes place through dealer quotes who make profits by quoting higher sell prices than buy prices. Traders buy or sell securities because they have either superior information or exogenous cash requirements. In limit order markets, the use of limit orders is not confined to liquidity suppliers who submit trades on both sides of the order book to capitalize the spread. For example, a trader who has no need for immediacy might submit a limit order instead of a market order to ensure a good price. On the other hand, not every market order requires superior information or cash demands. A liquidity supplier might spot a stale limit order and pick it off quickly by means of an immediate market order. In other words, the choice of a liquidity-supplying limit order depends more strongly on whether the trader can afford uncertainty of execution – in which case a limit order is more likely – or whether the trader's strategy requires immediacy – in which case a limit order is less likely. Consequently, the notion of limit order traders as liquidity-supplying dealers is not sufficient. In the following section I develop a formal framework of profit-maximizing limit order placement.

2.2 Static Models of the Limit Order Book

An early model of order placement was developed by Ho and Stoll (1983) where dealers compete for the next incoming order to hit their quotes. Glosten (1994) is the seminal paper on the theory of the limit order book market. Rock (1996) and Seppi (1997) extend the model to include discretized prices and a time priority rule for orders with the same limit price. Sandås (2001) and Frey and Grammig (2005) relax some assumptions of the theoretical model and provide empirical evidence. In the following sections I summarize the Sandås (2001) model. I do not extend the theoretical framework. Instead, the aim is to use this model to show why it is profitable to submit limit orders, why limit orders are submitted at different prices and how this leads to an equilibrium limit order book.

2.2.1 Model Assumptions

The market consists of two types of agents: liquidity suppliers and traders. Liquidity suppliers submit limit orders to the order book. They are risk-neutral and profit-maximizing. Traders submit market orders that consume liquidity. They trade either because they have private information or because they wish to satisfy liquidity requirements.

The agents trade a risky asset whose true value at time t is denoted by X_t.

This fundamental value of the asset is conditional on all publicly available information. The law of motion is given by

$$X_t = X_{t-1} + d_t, \tag{2.5}$$

where d_t is a random innovation. The increment d_t contains both new information of trades and new non-trade information in t.

Trading occurs over periods indexed by t. In each period, liquidity suppliers can submit new limit orders until no liquidity supplier wishes to supply a new order anymore. Then a trader arrives and submits a market order that consumes some of the liquidity in the order book. Finally, the new true value of the asset is announced and the procedure starts again.

Let the price vector $\{p_{+1}, p_{+2}, ..., p_{+k}\}'$ denote the ask prices in the order book where a positive index indicates the ask side. It is ordered from the best price, p_{+1}, to the $k-th$ best price p_{+k}. Let $\{Q_{+1}, Q_{+2}, ..., Q_{+k}\}'$ denote the volumes that correspond to the prices of the same index. Likewise, $\{p_{-1}, p_{-2}, ..., p_{-k}\}'$ and $\{Q_{-1}, Q_{-2}, ..., Q_{-k}\}'$ represent bid side prices and quantities. The incoming market order volume is denoted by m_t where positive quantities, $m > 0$, are buy orders and negative volumes correspond to sell orders, $m < 0$. Market orders arrive at the ask side and the bid side with equal probability. Their volume is exponentially distributed with parameter λ. The distribution of m is:

$$f(m) = \begin{cases} \frac{1}{2\lambda} e^{-\frac{m}{\lambda}} & \text{if } m > 0 \text{ (buy order)} \\ \frac{1}{2\lambda} e^{+\frac{m}{\lambda}} & \text{if } m <= 0 \text{ (sell order)}. \end{cases} \tag{2.6}$$

Figure 2.1: Density and Probability Function of Market Orders

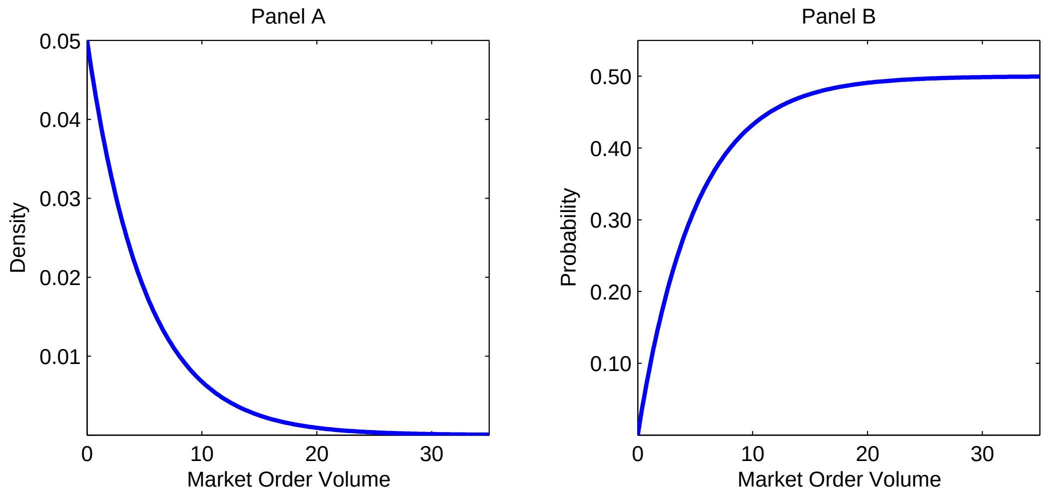

Figure 2.1 shows the density function (Panel A) and probability function (Panel B) of the market order volume specified in Equation 2.6. The figure assumes $\lambda = 5$.

Fig. 2.1 illustrates the density and probability function of the market order volume. Furthermore, each order incurs an order processing cost γ that is quantity invariant and equal for buy and sell orders.

2.2.2 Equilibrium Outcome

Limit order traders have no knowledge about the value of the random innovation d_{t+1}, yet they know that market order traders might be informed. Market order volume is informative about the future value of the asset. The relation between market order quantity and the change in the fundamental value of X is defined by a non-decreasing price impact function, $h(m)$. Limit order suppliers update their beliefs about the future value of the asset subsequent to the volume of the market order:

$$E[X_{t+1}|X_t, m] = X_t + h(m) \tag{2.7}$$

The specification of $h(\cdot)$ is assumed linear in the size of the incoming market order m. This leads to the following revision of beliefs:

$$h(m) = \alpha m \tag{2.8}$$

A market buy order leads to an increase in revised beliefs and a sell order leads to lower expectations of the asset value. All other things being equal, a larger value of α corresponds to stronger price impacts and a higher revision in beliefs. On the other hand, if $\alpha = 0$ the price impact function is horizontal. Then beliefs are not revised at all and the expectation of the future value is not influenced by the size of the incoming market order.

Let us turn to the liquidity supplier's decision problem whether it is profitable to submit a limit order to the order book. If a limit order that has a price of p_{+1} is executed, it generates an expected profit that depends on the size of the subsequent market order:

$$p_{+1} - E[X_{t+1}|X_t, m] - \gamma = p_{+1} - X_t - \alpha m - \gamma \tag{2.9}$$

The above equation shows that the expected profit depends on the size of the market order: large market orders are a signal of private information on the part of the trader and lower the expected profits of limit order suppliers.

Most importantly however, the execution of a newly submitted limit order is uncertain. Let q denote the cumulative quantity of all limit orders that have price and time priority. Any infinitesimally small new limit order only gets executed if the incoming market order volume is at least as high as q, $m \geq q$. Let $I(m \geq q)$ denote an indicator function that yields $I = 1$ for order execution and zero otherwise. At the best price level p_1, the expected profit of a liquidity supplier conditional on the execution of the limit order is:

$$
\begin{aligned}
E[p_{+1} - X - \alpha m - \gamma \mid I = 1] &= \int_q^{\infty} (p_{+1} - X_t - \alpha m - \gamma) f(m) dm \\
&= \int_q^{\infty} (p_{+1} - X_t - \alpha m - \gamma) \frac{1}{2\lambda} e^{-\frac{m}{\lambda}} dm \\
&= -e^{-\frac{q}{\lambda}} (X_t + \gamma + \alpha(q + \lambda) - p_{+1}) \qquad (2.10)
\end{aligned}
$$

A limit order trader is indifferent to adding another order to the book at p_{+1} if the expected profit is zero. Thus, equating the above profit to zero yields the equilibrium quantity Q_{+1}:

$$
Q_{+1} = \max \left\{ \frac{p_{+1} - \gamma - X_t}{\alpha} - \lambda; 0 \right\} \qquad (2.11)
$$

The zero profit condition can be extended to the quantity that will be offered at the next best price, p_{+2}, in the same way. Order execution is now dependent on $m \geq q + Q_1$ which implies

$$
E[p_{+2} - X_t - \alpha m - \gamma \mid I = 1] = \int_{Q_{+1}+q}^{\infty} (p_{+1} - X_t - \alpha m - \gamma) f(m) dm \qquad (2.12)
$$

and yields the following equilibrium quantity Q_{+2}:

$$Q_{+2} = \max\left\{\frac{p_{+2} - \gamma - X_t}{\alpha} - Q_{+1} - \lambda; 0\right\} \tag{2.13}$$

Let l be an index of the prices in the order book (with $l = 1, 2, ..., k$) and let π denote marginal profits for the case that the limit order gets executed ($I = 1$). In the general case, we obtain the following marginal profits for the order book:

$$\begin{aligned} E[\pi_{+l}] &= e^{-\frac{q+Q_{+(l-1)}}{\lambda}}\left[(p_{+l} - X_t) - \alpha(q + Q_{+(l-1)} + \lambda) - \gamma\right] \quad \text{(ask)} \\ E[\pi_{-l}] &= e^{+\frac{q+Q_{+(l-1)}}{\lambda}}\left[(p_{+l} - X_t) - \alpha(q + Q_{+(l-1)} - \lambda) + \gamma\right] \quad \text{(bid)} \end{aligned} \tag{2.14}$$

The corresponding equilibrium quantities are as follows:

$$\begin{aligned} Q_{+l} &= \max\left\{\frac{p_{+l} - X_t - \gamma}{\alpha} - Q_{+(l-1)} - \lambda; 0\right\} \quad \text{(ask)} \\ Q_{-l} &= \max\left\{\frac{X_t - p_{-l} - \gamma}{\alpha} - Q_{-(l-1)} - \lambda; 0\right\} \quad \text{(bid)} \end{aligned} \tag{2.15}$$

The above equations summarize the status of the limit order book when it is in equilibrium and limit order traders have exploited all profit opportunities that yield a positive expected payoff.

2.2.3 Implications

In this section I highlight the intuition of Equation 2.14 and Equation 2.15 by providing a numerical example. Then I discuss the implications that the model has for the later sections of my thesis.

Figure 2.2 compares the price schedule and the expected payoffs to limit orders

Figure 2.2: Order Book Schedule and Profit Opportunities

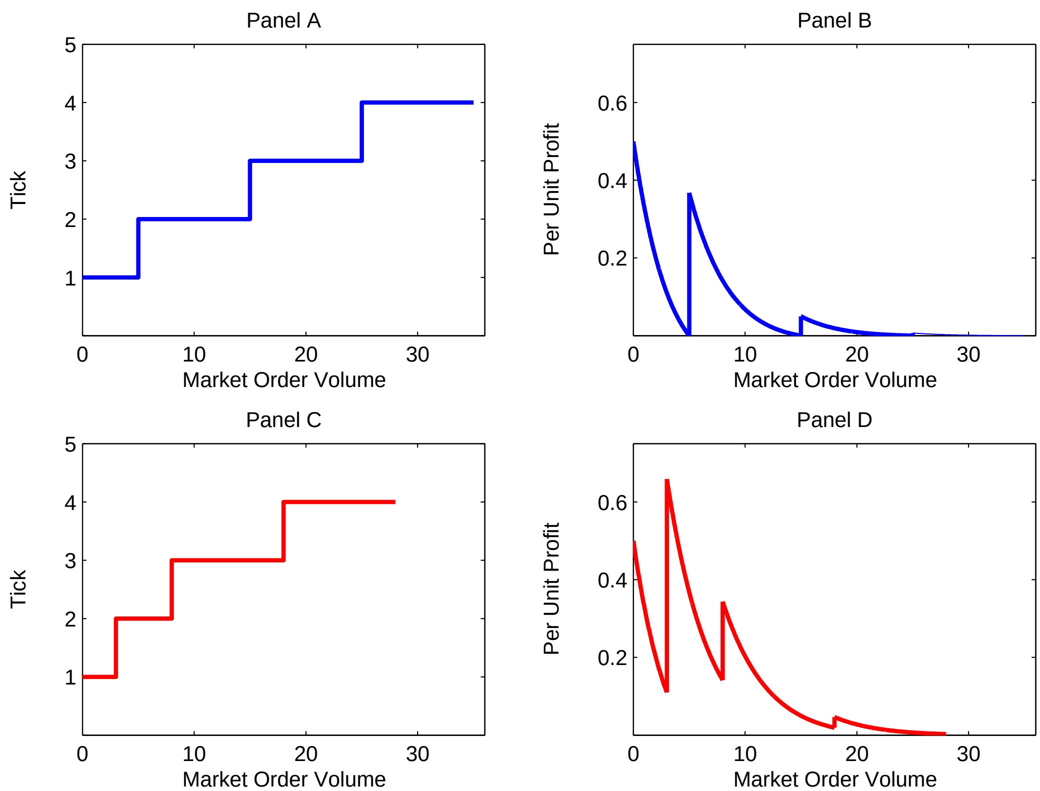

Figure 2.2 shows the price schedule and profit opportunities of two limit order books. Panel A and B (blue graphs) correspond to an order book that is in equilibrium, while Panel C and D (red graphs) belong to a book with unexploited profit opportunities. The figures are numerical examples of Equation 2.14 and Equation 2.15 computed for $X_t = 100$, $\alpha = 0.1$, $\gamma = 0$ and $\lambda = 5$. The y-axis is in ticks where the first tick is the first possible ask price $p_{+1} > X_t$.

traders for two different order books. Panel A and B show a limit order book and payoff for the parameter constellation $X_t = 100$, $\alpha = 0.1$, $\gamma = 0$ and $\lambda = 5$. The book is in the equilibrium state implied by the Equations 2.14 and 2.15. In contrast, Panel C and D show the snapshot of a hypothetical limit order book which is not in equilibrium.

The most striking difference between the two order books is that the submitted limit order quantities of the non-equilibrium order book are smaller. Panels A and C show the cumulative quantity in the book against their respective prices. In equilibrium, the marginal break-even quantity at the first tick is 5 and at the second tick it is 10. However, the bottom order book offers only 3 shares at the first tick and only 5 shares at the second tick. Subsequently, it is still profitable to offer 2 more shares at the first tick and 5 more at the second tick. The unexploited profit opportunities are highlighted in Panels B and D. If we compare the two panels we see that the graph in Panel B always falls to zero before it jumps back up again, while the jumps of the graph in Panel D occur for non-zero values. This underlines the fact that in the non-equilibrium book, marginal profits are still positive.

The equilibrium properties of the Sandås (2001) model help us understand why limit order traders submit liquidity to the limit order book. In particular, we can derive implications that we use in the later sections for the empirical study of the limit order book. The model formalizes the following behavior and incentives of limit and market order traders:

1. **Profitability of a Limit Order**: Liquidity suppliers submit buy limit orders at prices which are below their expectation of the future price (and vice versa for sell limit orders). These limit orders get executed against market orders and generate a profit for limit order traders. This implies that, ceteris paribus, a higher amount of market order trading offers higher profit opportunities for liquidity suppliers.

2. **Position in the Limit Order Queue**: There are profit opportunities along the whole price-quantity schedule of the limit order book. Limit orders near the best tick have higher execution probability and lower conditional profits, while limit orders further away have lower execution probability but higher conditional profits. This implies that limit order flows can take place deeper in the book even if there is no change at the best price.

3. **Transaction Costs of Market Orders**: The equilibrium limit order book shows that large market orders get executed against several different limit order prices. The larger the market order, the larger the marginal costs that it incurs. Large market orders take away depth from the limit order book and shift the price-quantity schedule. This implies that liquidity beyond best prices becomes relevant for large orders.

4. **Information Signals**: Limit order traders take the size of trades as an informative signal of the true value of the asset which they take into account in their future liquidity supply. They interpret large orders as highly informative and revise their beliefs particularly strongly. This implies that for limit order flows in any period we have to distinguish whether they took place in an information-intensive or low-information environment.

The model implications might only highlight simple mechanics of limit order book markets, but they are a good starting point to compare the empirical features of the limit order book against. Yet before we proceed any further, let us cast a brief look at the weaknesses of the model. Firstly, when an order book is in equilibrium, the model does not allow the submission of new liquidity at the best prices. New limit orders only replace liquidity at the furthest tick.[6] A second weakness is that market orders are assumed to be exogenous. In practice, however, the dynamics of the limit order book are far more complex.

2.3 Dynamic Models of the Limit Order Book

In the following section I present a dynamic model of a limit order book market. The theoretical literature includes models by Parlour (1998), Foucault (1999), Parlour and Seppi (2003)), Goettler, Parlour and Rajan (2004) and Foucault, Kadan and Kandel (2005). I present a simplified version of Foucault, Kadan and Kandel (2005) who endogenize the decision between limit and market order placement.[7] While the previous model provided implications why limit order traders trade and where they submit their limit orders, I present the following model to obtain hypotheses for the dynamic properties of limit order flow.

[6] However, Hasbrouck (2004) points out that this problem is overcome when uncertainty is introduced, albeit at the cost of simple analytic solutions.

[7] In contrast, Parlour (1998) models how the order placement decision depends on the depth of the limit order book at the best quotes. Foucault (1999) addresses the risk that limit order strategies lose against agents who have better information. Parlour and Seppi (2003) set up a dynamic model of different exchanges that compete against each other for order flow. Goettler, Parlour and Rajan (2004) model limit order trading as a stochastic game that takes place sequentially. I present the Foucault, Kadan and Kandel (2005) framework because it is based on the immediacy of market orders versus the delayed execution of limit orders. This framework is well-suited to derive hypotheses for the time series behavior of liquidity flows as I will be doing in the following sections.

2.3.1 Model Assumptions

The market is organized as a limit order book in which market participants trade one single security. The highest sell price of the security is p_A^{max} and the lowest buy price is p_B^{min} (with $p_A^{max} > p_B^{min} > 0$). The best ask price is denoted by p_A and the best bid price by p_B.[8] The model investigates the dynamics of the bid-ask spread, $s := p_A - p_B$, within the maximum price range, $p_A^{max} - p_B^{min}$. At the endpoints of the price interval traders offer to sell and buy an unlimited amount of shares.

Market participants arrive at the market following a Poisson process with parameter $\lambda > 0$. The model has an infinite horizon and the times between trade arrivals are exponentially distributed with an expectation of $\frac{1}{\lambda}$. Market participants are risk-neutral. Upon their arrival they submit either a market order, which is executed immediately, or a limit order, which is executed later, but at a better price. Furthermore, each agent bears waiting costs δ (per unit of time) for the time until order execution. Traders either belong to the class of impatient traders or patient traders: impatient traders value fast trade execution and have high waiting costs of δ_I, while patient traders have low waiting costs of δ_P (with $\delta_I > \delta_P$). The proportion of patient traders to impatient traders in the population is Θ_P; the proportion of impatient traders is $\Theta_I = 1 - \Theta_P$.

The trading mechanism of the market is a centralized limit order book in

[8] For simplicity I assume a unit tick size. The results are qualitatively identical if a tick size variable is included, yet they make the model unnecessarily complicated as it is not my aim to study tick size effects.

which all limit orders are executed according to price priority. To facilitate the analysis of the equilibrium outcome, Foucault, Kadan and Kandel (2005) make the following assumptions:

- Traders arrive only once, submit a market or limit order of size 1 and exit the market. Orders that have been submitted cannot be cancelled or revised.
- Limit orders always improve the current best price and reduce the spread by at least one tick. They cannot queue at the same price.
- A buy order is always followed by a sell order and sell orders are always followed by buy orders. The probability that the first order is a buy is 0.5.

I denote the execution price of buy and sell orders by p_{+E} and p_{-E}. A buyer either submits a market order for which he pays the lowest ask price, p_A, or submits a limit order on the bid side. The limit order enters the order book at a new best bid price, p_B. Likewise, a seller either trades at the best bid price or enters a new best ask price. Hence, execution prices can be written as

$$p_{+E} = p_A - j \tag{2.16}$$

$$p_{-E} = p_B + j \tag{2.17}$$

where j denotes the number of ticks from the best price in the limit order book against which the order could be executed (with $j \in \{0, ..., s-1\}$). For a market order $j = 0$, the time to execution is zero and therefore the waiting costs are zero, too. For a limit order it must hold that $j \in \{1, ..., s-1\}$ depending on the spread that it creates. Its time to execution is $T(j)$ and its expected waiting costs are $\delta_f T(j)$ with $f \in \{I, P\}$. The payoff from the choice of order type – market order versus j-limit order – can be written as the difference between the bid-ask spread and waiting costs, $\pi_f(j) := j - \delta_f T(j)$. For a market order it is zero, $\pi_f(0) = 0$, while for a limit order it is either positive or negative. A traders order placement is optimal if it solves

$$\max_{j \in \{0,...,s-1\}} \quad \{\pi_f(j) := j - \delta_f T(j)\} \tag{2.18}$$

for buyers and sellers alike.[9] In equilibrium, a trader's strategy solves Equation 2.18 for waiting costs calculated under the same strategy.

2.3.2 Equilibrium Outcome

As the payoff of a market order is zero, a trader will only submit a limit order if his price improvement j offsets his waiting costs $\delta_f T(j)$. A limit order trader has to wait at least one period for the execution of the limit order. Since the average time between orders is $\frac{1}{\lambda}$ the smallest expected waiting costs of a trader of type f are $\frac{\delta_f}{\lambda}$. Hence, the smallest spread j_f^* that a trader will create – which we call the reservation spread – must be the next highest integer above $\frac{\delta_f}{\lambda}$. I assume that the reservation spread of patient traders is smaller than the reservation spread of

[9] Foucault, Kadan and Kandel (2005) assume that a trader who is indifferent between two orders submits the order that creates the higher spread.

impatient traders, $j_P^* < j_I^*$, and that both are smaller than $p_A^{max} - p_B^{min}$.[10] The following section presents the resulting equilibrium and the dynamics of the limit order book in equilibrium.[11]

In equilibrium there exists a cut-off spread s_c which, together with patient traders' reservation spread, defines three different spread regions of the limit order book. In the region $< 1, j_P^* >$ patient and impatient traders submit market orders. If the spread is in the region $< j_P^* + 1, s_c >$ only patient traders will submit a price-improving limit order, while impatient traders will submit a market order. In the region $< s_c + 1, p_A^{max} - p_B^{min} >$ both patient and impatient traders submit limit orders which narrow the spread. An important point to note is that impatient traders sometimes submit market orders even if the spread is higher than their reservation spread (when $s_c > s > j_I^*$). The reason is that the expected waiting costs of an impatient trader will, in general, exceed his improvement in execution price. For ease of presentation I consider the case when $s_c = p_A^{max} - p_B^{min}$. This assumption has no impact on the results, yet it shortens the presentation.

The optimal order of a trader depends on the current bid-ask spread in the limit order book. In equilibrium, there exist exactly g such spreads which I order from lowest to highest, $n_1 < n_2 < ... < n_g$. The lowest spread is equal to the

[10] The equilibrium state can be derived more easily is we assume that the reservation spread of a patient and impatient trader are equal. However, it is more realistic that the reservation spread of the impatient patient trader is higher, because the impatient trader attaches more importance to fast execution.

[11] I only present the equilibrium results, but do not prove them explicitly. My main intention is to convey their economic intuition. The technical details are in Foucault, Kadan and Kandel (2005), 1178-1184 and 1209-1215.

patient traders' reservation spread ($n_1 = j_P^*$) and the highest spread n_g equals $p_A^{max} - p_B^{min}$. Consequently $n_g = s_c$ and impatient trader always submit market orders irrespective of the current spread. In contrast, a patient trader submits either a market or a limit order depending on the current spread: if the spread equals the patient trader's reservation spread ($s = j_P^*$), the patient trader submits a market order. If the spread is higher than the patient trader's reservation spread ($s > j_P^*$) he submits a price-improving limit order.

It is important to note that limit orders always reduce the current equilibrium spread to the next highest equilibrium spread. For example, if the current spread is n_g a patient trader will create a new spread of n_{g-1}. The spread improvement by this new limit order is therefore $n_g - n_{g-1}$. How large it is in terms of ticks depends on the values of the bid-ask spreads. Let Δ_h denote the spread improvement of a limit order that narrows the spread from n_h to n_{h-1}, $\Delta_h := n_h - n_{h-1}$. It can be shown that, in equilibrium, the spread improvement is determined endogenously by the ratio of patient to impatient traders, $\frac{\Theta_P}{\Theta_I}$, the current position of the spread, h, the waiting costs of patient traders, δ_P, and the expected time between order arrivals, $\frac{1}{\lambda}$,

$$\Delta_h = \text{int+}\left(2\left(\frac{\Theta_P}{\Theta_I}\right)^{h-1}\frac{\delta_P}{\lambda}\right) \qquad (2.19)$$

where $h \in \{2, ..., g\}$ and int+ indicates the next highest integer value. The spread improvement enables us to determine the set of equilibrium spreads as the sum of patient traders' reservation spread and subsequent spread improvements:

$$\begin{aligned}
n_1 &= j_P^* && (2.20) \\
n_h &= n_1 + \sum_{k=2}^{h} \Delta_k && (2.21) \\
n_g &= p_A^{max} - p_B^{min} && (2.22)
\end{aligned}$$

where $h \in \{2, ..., g-1\}$ and $g-1$ is the last integer for which $n_h < p_A^{max} - p_B^{min}$. Equations 2.20 to 2.22 define the set of all equilibrium spreads. The trading process is characterized by constant changes in the spread: at any current spread h a limit order trader reduces the spread by Δ_h to n_{h-1}. As long as only patient traders submit orders this process continues until their reservation spread n_1 is reached. This is the lowest possible equilibrium spread which we call the competitive spread. At the competitive spread, even patient traders will submit a market order that widens the spread. When impatient traders arrive they always submit market orders which widen the spread for all spreads from n_1 to n_{g-1}. At n_g it stays unchanged until a patient trader arrives who narrows the spread again.[12]

2.3.3 Implications

In equilibrium the spread follows a stochastic process whose values all belong to the set of equilibrium spreads. The movement along the equilibrium path remains stochastic because the patience of a trader is a random variable: an incoming

[12] The spread stays unchanged at n_g because of the assumption that a pool of traders stands ready to supply an unlimited amount of buy and sell orders at p_A^{max} and p_B^{min}.

order from a patient trader arrives with probability Θ_P and an order from an impatient trader arrives with probability Θ_I. A simple way of characterizing the stochastic process of the limit order book is to measure the probability with which the spread reverts to its competitive level before the next market order arrives. If the current bid-ask spread is n_h a complete reversal to the competitive spread requires $h-1$ consecutive limit orders. If we denote time with an index t the conditional probability can be written as follows:

$$Pr\ (s_{t+1} = j_P^* \mid s_t = n_h) \quad = \quad \Theta_P^{h-1} \tag{2.23}$$

for $h \in \{2, ..., q\}$.[13] It is important to note that h is a subset of the equilibrium number of spreads which is determined endogenously in the model. Therefore, the equilibrium reversal of spreads to their competitive level – the resiliency of the market – depends on all exogenous parameters.[14]

More specifically, we can use the equilibrium properties of the model to derive implications for empirical studies of the limit order book. They involve the relationship of the exogenous variables (the proportion of patient traders, the waiting costs, the order arrival rate) with the resiliency of the market. In particular, the model establishes the following hypotheses:

[13] The probability of spread reversal at the competitive level, $Pr\ (s_{t+1} = j_P^* \mid s_t = j_P^*)$, is zero by construction.

[14] An important point to note is that in Foucault, Kadan and Kandel (2005) definition of resiliency, the reversal of liquidity to its pre-shock mean only refers to the bid-ask spread. As all orders have a unit size in their model, they cannot study the reversal of depth.

1. **Order arrival rate**: All other things held equal, a higher order arrival rate implies a less resilient limit order book.

2. **Time between trades**: There is positive association between the average time between trades, conditional on the size of the spread, and market resiliency.

3. **Proportion of patient traders**: The higher the proportion of patient traders in the population, the higher the resiliency of the limit order book.

Foucault, Kadan and Kandel (2005) go on to argue that traders' impatience is likely to increase towards the end of the trading day. For this reason they conjecture that the proportion of patient traders falls at the end of the trading day. This yields a fourth hypothesis:

4. **Time of Day**: A limit order book is less resilient at the end of the trading day than in earlier trading periods.

Although the market structure of a limit order book market is modeled in a very stylized fashion, the model still yields rich implications for the dynamics of the limit order book. Future research will no doubt relax some of the assumptions to widen the scope of the theoretical results. In the next section I conclude how I build on the theoretical literature in my empirical research.

2.4 Conclusion

Static models of the limit order book address the equilibrium state of the limit order book under the assumption that liquidity is not risky. For example, Sandås (2001) shows that it is rational for limit order traders to submit new liquidity as long as their marginal profit is positive. All things being equal, the updating of beliefs determines the equilibrium quantities and expected profits for all prices in the book. The model pinpoints why limit orders are submitted at different ticks and how the typical shape of the limit order book arises. It has the implications the market order flow offers profit opportunities for limit order traders to exploit and that limit order flows can take place anywhere in the limit order book. Furthermore, it implies that we should measure liquidity beyond best prices and that limit order supply varies with the information content of trades.

Dynamic models of the limit order book address the evolution and equilibrium state of the limit order book over time. For example, Foucault, Kadan and Kandel (2005) develop a dynamic framework in which traders can choose to submit limit orders which improve the spread or market orders which clear the book and widen the spread. The order choice determines the evolution of the spread along time and its reversal to a competitive level. The model derives an equilibrium for the endogenous order choice and thereby also endogenizes the resiliency of the limit order book. In terms of liquidity dimensions, Foucault, Kadan and Kandel (2005) extend the theoretical literature to the resiliency of liquidity. However, they only model the resiliency of the bid-ask spread. They do not model the equilibrium levels of depth nor do they address the resiliency of depth.

The static and dynamic models in the previous sections show that the theoretical literature has already addressed many aspects of liquidity, yet that the current generation of models provides only partial views of a limit order book market. So far, the theoretical literature has not yet brought together the spread dimension, the limit order book's depth at different ticks and its resiliency mechanism. Therefore there is no consistent foundation for the study of more specific issues – for example the study of liquidity risk in limit order book markets – to build on. However, as these are very relevant issues and as, in practice, the missing elements in the theoretical models are important features of the market structure, it is not surprising that the empirical research has overtaken the theoretical literature. Likewise, I depart from the predictions in Sandås (2001) and Foucault, Kadan and Kandel (2005) to empirically extend their scope to a more rigorous treatment of liquidity risk in limit order book markets.

In particular, I use the concept of resiliency in dynamic equilibrium models to set up an econometric model that can be tested empirically. I extend the recovery of the spread to the recovery of depth and examine whether the implications in Foucault, Kadan and Kandel (2005) hold strong for the reversal of the spread and of depth. Finally, I explore the relationship between the resiliency of the spread and the resiliency of depth. The study of systematic liquidity risk across stocks requires a multi-asset setting. In the absence of multi-asset models for the limit order book[15] I build on Sandås (2001) to measure the liquidity supply of the price-quantity schedule. I then proceed empirically to estimate the extent of commonality in liquidity and the time variation of commonality. While the following chapters contribute to a comprehensive view of liquidity risk in limit order book markets, they also give insights for future modeling purposes.

[15] Domowitz, Hansch and Wang (2005) present empirical evidence that the correlation of order type seems to be a determinant behind the commonality in liquidity. Coughenour and Saad (2004) relate commonality to the correlated behavior of market makers and limit order placement. An incorporation of depth into the Foucault, Kadan and Kandel (2005) model would endogenize both of these suggested determinants of commonality and make it a good candidate to extend to a multi-asset context.

Chapter 3

Market Structure and Data

In the following sections I present the data that I use in the empirical studies in Chapters 4 to 6. First I give some details of the market model that is used at the Frankfurt Stock Exchange. Then I present the reconstruction of the limit order books from the raw data and finally give some descriptive statistics.

3.1 Market Structure

3.1.1 Xetra Market Model

In my thesis I use data from the electronic limit order book market at the Frankfurt Stock Exchange (FSE). The electronic system which is used in Frankfurt is called Xetra. Anyone with a computer that is connected to Xetra can trade stocks directly without going through further intermediaries. The same trading platform is also used at the stock exchanges in Vienna and Dublin as well as at the European Energy Exchange (EEX) in Leipzig. Xetra is the dominating trading venue for stocks that are listed in Frankfurt. There are some regional exchanges as well as a floor trading facility, however the Xetra system attracts more than 98% of all trading activity.

The exchange market model for stock trading defines the way in which agents can submit orders and how these orders are matched. It includes price determination, priority of orders as well as the scope of information that is disclosed and observable to traders. The market model implements legal and regulatory requirements of the stock exchange as well as the terms and conditions of trading in Frankfurt. The FSE is an order-driven exchange that allows market orders, limit orders, market-to-limit orders, stop orders and iceberg orders.[16] Let us start with a brief overview of some fundamental principles of the market model:[17]

1. A security can be traded continuously or only in auctions.

2. Continuous trading starts with an opening auction, can be interrupted by intraday auctions and ends with a closing auction at the end of the day.

3. During the auction phase, the order book is partially closed and during continuous trading, the order book is completely open to all agents.

4. Orders are executed according to price priority and then time priority.

5. Trading is anonymous: before a trade takes place agents do not know whom the orders in the book belong to and their own identity is not revealed either.

[16] Market-to-limit orders, stop orders and iceberg orders will be explained in more detail in the later parts of this section.

[17] For more details on the legal framework and the specific features of the Xetra market model of stock trading see Deutsche Boerse Group (2004).

6. All order sizes can be traded, both round and odd lots.

7. There is no tick size requirement (the minimum tick size is 0.01 Euros).

8. Trading is interrupted if the potential price jumps to a price that lies outside a pre-defined range and an auction is initiated ("volatility interruptions").

9. A trade confirmation is disseminated automatically after a trade has taken place with price, quantity and counterparty information.

The stocks listed at the Frankfurt Stock Exchange are segmented into various different groups. The main criteria for segmentation are market capitalization, liquidity or industry affiliation. The best-known segment is no doubt the DAX 30 which comprises Germany's thirty largest blue-chip stocks. Further segments are the MDAX, TecDAX, Liquid Foreign Equities, Illiquid Small Caps and Illiquid Foreign Equities. All trading segments have in common that the trading of stocks in the same segment follows the same rules.

An important feature of a market segment is the organization of its liquidity supply. In general, equities require at least one Designated Sponsor to be accepted for trading in the market model of continuous trading. A Designated Sponsor is an investment bank or securities firm that increases a share's liquidity by offering to buy and to sell simultaneously. However, the blue-chip segment is considered liquid enough without any market-making. Consequently, the liquidity of German blue-chip stocks in Xetra depends on the competition of limit order traders alone without any further institutional market makers. In the following sections I will be using DAX 30 data, which enables me to study the properties of a purely order-driven limit order book market.

3.1.2 Order Types and Matching Rules

The core feature of any trading platform is the mechanism how orders get matched and transactions take place. In Frankfurt, trading is based on a continuous double auction mechanism where orders are matched automatically by the trading system. The opening hours are from 9.00 to 17.30 CET. During those hours, the limit order book openly displays all orders that have been submitted already, but not executed yet. Market participants who wish to trade will then choose from several different order types.

The basic order types allowed during continuous trading are conventional market orders, limit orders as well as market-to-limit orders. A market-to-limit order is treated as a market order and executed against the best price in the order book. However, the remaining part that cannot be executed at the best price is converted to a limit order at the transaction price of the market order part. All order types can further be restricted to immediate-or cancel (IOC) or fill-or-kill (FOK). IOC orders are executed immediately and fully or as fully as possible; non-executed parts are deleted directly. FOK orders are only executed immediately and fully or not at all. Finally, the validity of orders can be specified further by means of the restrictions good-for-day, good-till-date and good-till-cancelled. A good-for-day order is only valid for the current exchange trading day and is then automatically cancelled by the exchange. A good-till-date order is valid until a specified date which can be up to 90 calendar days from the current trading day. It is not automatically deleted at the end of the trading day, but can be taken out by the originator at any time if he wishes to do so. A good-till-cancelled order has no specified date and is valid until it is either executed, deleted by the originator of the trade or it reaches its maximum validity of 90 days and is then cancelled automatically.

Beside the basic order types, there are some more sophisticated order types that are allowed in the Xetra market model, in particular stop orders and iceberg orders. A stop order is allowed to support trading strategies that depend on the occurrence of certain price events. In particular, a stop market order is a market order that is placed in the order book as a market order as soon as the stop price is reached. A stop limit order is a normal limit order that is placed in the order book if the pre-specified stop price is reached. In contrast, iceberg orders do not depend on any price being reached. They enable the entrance of large orders without the submitting party having to disclose the total order volume all at once. An iceberg order is a special form of a limit order that has a limit price, and overall order volume and a peak. The peak enters the order book as a normal limit order and, once it has been hit and fully executed, is replaced by a new limit order of the same peak size with a new time stamp. Iceberg orders are only ever valid for one trading day and get cancelled automatically at the end of the day; if they are supposed to run for a second day, the trader has to re-enter the order the next morning.

Each new incoming order is immediately checked for execution against orders on the other side of the order book. Matching and price determination take place on the basis of price and time priority. The time criterion applies if orders share the same price limit – earlier orders take priority. Incoming market orders are executed at the highest bid limit or lowest ask limit and, in the case of large volumes, walk up the book. Therefore, they can be executed against one limit price or also against several limit prices. It also possible that market orders only get executed partially if the limit order book does not have enough liquidity, however this is a very rare event for blue-chip stocks. In such a situation, the market order then has to be executed immediately against the next incoming order. If it cannot be executed at the last transaction price, the price is determined by the next incoming limit order.

The transparency and the sole reliance on anonymous limit order submissions make Frankfurt Stock Exchange a well-suited trading platform for the study of the microstructure of markets and market participant behavior. In an international context, the Xetra trading platform – at least its blue-chip segment – is one of very few exchanges that operate as a pure limit order book market without any dealers. For example, London Stock Exchange (LSE) uses the electronic platform SETS which supports anonymous limit order trading, but it also offers a dealer market off the book, which, effectively, competes with limit order traders. The New York Stock Exchange (NYSE) as well as the technology platform NASDAQ also use hybrid systems of automated order-driven trading and off-the-book matching which can improve prices. The Paris Stock Exchange, operated by Euronext, uses an electronic system called NSC which, like Xetra, is a purely order-driven market. However, traders can only see a limited amount of orders in the book. The great advantage of using Xetra data for my analysis is that it offers a comprehensive picture of a pure limit order book market in which the limit order book is displayed openly and fully.

3.2 Data Set

3.2.1 Order Book Reconstruction

The data set which I use ranges from 1 January 2004 and to 31 March 2004. It comprises the thirty largest stocks listed in Frankfurt which make up the German blue-chip index DAX 30. Deutsche Boerse provided us with the entire trading protocol which is recorded automatically for each action that takes place in the Xetra system. The trading protocol keeps record of all order entries to the system, order revisions, order cancellations, order expirations and executions. The data I was provided with is the original output of the Xetra system which is produced when any order is processed. As such, all events which took place in Xetra have to be contained in the protocol by construction. The only modification of the data is that the stock exchange deleted all originator information to keep customer information confidential.

Table 3.1 shows an excerpt of the raw data. Each order is recorded with a unique order number, date and time stamp. The date column is coded in integer values which count the number of days after the 01 January 1960. The time also uses integer values to give the number of seconds after midnight. Further information that is recorded includes the direction of the trade, the order type and the event that it triggered (column 6 in Table 3.1). For example, the value 1 corresponds to order entries, 5 corresponds to partial fills and 4 corresponds to final fills. Furthermore, all relevant price and quantity specifications and processing information is recorded as well. If an order is revised or if it is partially executed, the order number stays the same so that all actions that refer to one and the same order can be traced. However, the matching of orders is not identified in the system: when a market order is entered and executed, the system only generates a final fill record. Likewise, a final fill is recorded for the limit order, however there is no identification that links these two orders. The matching of limit and market orders was only possible by implementing the complete set of trading rules of the exchange.

The first step of the data analysis is the reconstruction of each stock's order book. At each point in time, the volume of all limit orders which belong to the same order book side and which carry the same limit order price is aggregated. A new limit order event initiates an update of the order book: if an order is entered at a price that is already in the book, the volume at that price is increased by the volume of the new limit order. If there is no limit order in the book at that price,

Table 3.1: Example of Allianz's Trading Protocol

Order Stamp			Information			Price		Size		
ID	Date	Time	Dir.	Ty.	Ev.	Limit	Match	All	Proc.	Rem.
27746	16073	33485	B	L	1	100.25	0	892	892	892
22817	16073	33496	S	L	5	100.26	100.26	1700	100	1600
22817	16073	33500	S	L	5	100.26	100.26	1700	500	1100
28473	16073	33500	B	M	1	0	0	500	500	500
28473	16073	33500	B	M	4	0	100.26	500	500	0
22817	16073	33505	S	L	4	100.26	100.26	1700	1100	0
28650	16073	33505	B	M	1	0	0	1450	1450	1450
28650	16073	33505	B	M	5	0	100.26	1450	1100	350
28650	16073	33505	B	M	4	0	100.29	1450	350	0
22665	16073	33506	S	L	3	100.35	0	1300	1300	0
28708	16073	33506	S	L	1	100.42	0	1300	1300	1300
22731	16073	33506	S	L	3	100.31	0	376	376	0
28712	16073	33506	S	L	1	100.36	0	376	376	376
28712	16073	33507	S	L	3	100.36	0	376	376	0
28762	16073	33507	S	L	1	100.37	0	376	376	376

Table 3.1 shows an excerpt of the raw data set. Each row corresponds to one event that was recorded in the Xetra trading system. The columns contain all necessary information to identify the event. Each event gets a unique order number (ID), date and time stamp (columns 1-3). Some general information is recorded with regard to the direction (Dir) and type of the trade (Ty) as well as the processing events (Ev) in columns 4-6. Columns 7 and 8 contain the limit and matching prices (Limit and Match). Columns 9-11 represent the original order size (All), the processed size (Proc) and the remaining volume (Rem). This example is only a small excerpt for a specific stock (Allianz). The complete data set has one matrix for each stock. The number of rows varies between 867,369 and 4,728,368.

a new price limit with corresponding volume is added to the book. Likewise, limit order cancellations reduce the book's volume, while revisions reduce the volume at the old price and increase the volume at the new price. The execution of market orders also reduces the volume of the limit order book. As Deutsche Boerse Group recorded the initial order book for each stock, I use a program that, starting from the initial state, updates the limit order book continuously for each event in the original data. To do justice to hidden liquidity in the book (which results from iceberg orders in the system for which only the peak is displayed in the book), I first construct construct an order book of all visible limit orders and secondly a book which contains visible as well as hidden orders.

The second step is to reduce the data to a manageable and regularly spaced sample. The relevant frequency at which the order book should be sampled depends on the issue under study. I save snapshots of the order book every 5 minutes, 15 minutes and 60 minutes. This yields time series of the order book for all stocks at five different frequencies. Furthermore, I cut off the opening and closing auction because their trading mechanism is different from the normal continuous trading period. This reduces the data set to 64 trading days with the number of order book snapshots ranging from 8 (at the 60-minute frequency) to 102 (at the 5 minute frequency). Each time series is sampled for the visible order book and for the complete order book including hidden liquidity.

The third step is to match the trading activity to the reconstructed order books. To achieve this aim I first construct time series which count the number of each order type within the same time interval. Secondly I construct further time series which aggregate the volumes of all different order types within the same interval. I obtain two time series for each order type which give the number of orders and their cumulated volumes. Each element belongs to the order book of the corresponding interval. For longer intervals these volumes are larger by construction and smaller for shorter frequencies. Next I turn to some descriptive

Table 3.2: Summary Statistics of the Data Set – Aggregation

	Mean	Median	Max	Min
Market capitalization	18.20	10.39	61.29	2.95
Average daily trading volume	114.79	75.22	348.60	14.13
Absolute bid-ask spreads	0.04	0.03	0.09	0.01
Relative bid-ask spreads	0.09	0.10	0.15	0.05
Depth at the best bid and ask	180,487	148,461	979,214	65,773

Table 3.2 summarizes the main characteristics of the stocks and corresponding limit order books in the data set. Market capitalization is given in billions of Euros as of 01 January 2004. Trading volume is the average daily trading volume in millions of Euros between 01 January 2004 and 31 March 2004. The average bid-ask spread and depth at the best bid and ask are computed from the time series of limit order books. Absolute spreads and depth at the best price are in Euros and relative spreads in %.

statistics of the stocks in the data set, their limit order books and their order flow which characterize liquidity in the Xetra system.

3.2.2 Descriptive Statistics

Table 3.2 gives some summary statistics of the stocks in the data sample. The market capitalizations as of 01 January 2004 range from 2.95 billion to 61.29 billion Euros. Together, the market capitalization of all DAX 30 stocks makes up about 98% of the German market. The average daily trading volume varies from 14.13 million Euros to 348.60 million Euros. Trading activity is fairly high: even the least liquid stock is traded about 75 times a day on average. I use the 5-minute snapshots of the order book to compute a time series of the bid-ask spread and associated depth of all stocks. On average, the absolute spread between the best ask price and the best bid price is between 0.01 and 0.09 Euros or, in relative terms, 0.05% and 0.15%. The volume at the best bid and ask price lies between

Figure 3.1: Variation in the Liquidity of the Sample Stocks

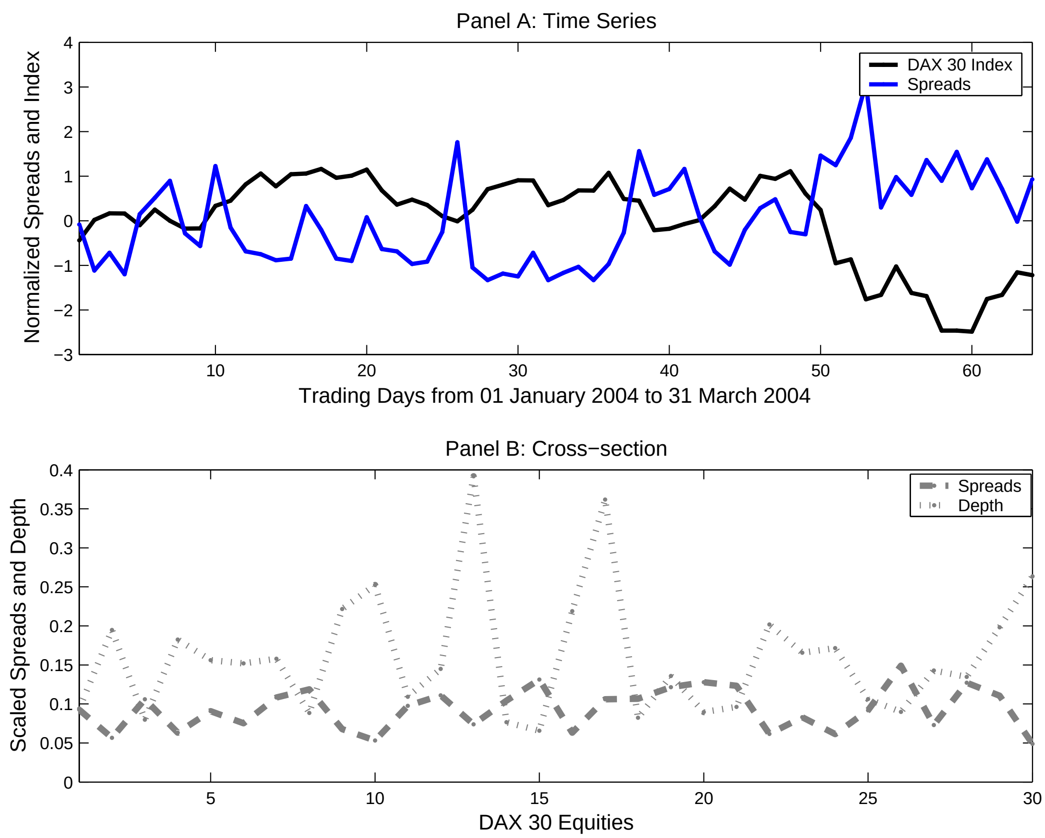

Figure 3.1 shows the variation in the liquidity of the sample stocks. Panel A is a time series graph of the market average of the bid-ask spread and the market trend. Panel B illustrates spreads and depth in the cross-section.

65,773 and 979,214 Euros. In all, these figures show that the DAX 30 is a very liquid segment that offers high volumes at fairly low transaction costs.[18]

Figure 3.1 illustrates the variation in the liquidity of the sample stocks, using the bid-ask spread as the measure of liquidity. Panel A plots the time series of the DAX 30 index performance and the time series of the average bid-ask spread across stocks. The panel makes very clear that there is considerable time

[18] Table A.1 in the appendix gives these figures individually for each stock of the data set.

variation in the spreads. Spreads tend to be low when the market rises and high when the market falls. However, in relative terms the variation of the market spread is far stronger than the variation in the index return. Panel B is a plot of the cross-sectional properties of liquidity. The differences in the spread are not specially striking across stocks. The difference in depth, however, is fairly strong: the deepest stock offers about seven times as much volume as the least liquid stock. Furthermore, Panel B shows that stocks with low spreads also tend to have high depth. The figure reinforces that liquidity should not be treated as a constant cost component; rather, it bears considerable risk along time.

Table 3.3 compares the submission behavior for the various order types. On average, 887,705 orders were submitted for a stock during the three month sample period. About 97.3 % of all submissions were limit orders, 2.1 % market orders, 0.6 % iceberg orders and only 0.05 % market-to-limit orders. These percentages show that the vast majority of order flow constitutes liquidity provision. Of the submitted limit orders, 23 % got executed and 77 % cancelled. Evidently, the largest proportion of liquidity supply gets cancelled and is not consumed. With regard to order size, the average market order is 23,959 Euros, while the average market-to-limit order is 19,522 Euros. If we compare these figures to the depth of the order book we can conclude that, on average, even the least liquid stock (with an average depth of 65,733 Euros) is liquid enough to absorb three normal-sized trades within its best price range. In contrast, iceberg orders exhibit much larger volumes (an average of 630,000 Euros). In relation to the normal order book

Table 3.3: Average Order Submissions

	Bid Side (Buys)	Ask Side (Sales)	All Orders (in sum)
Panel A: Number of Submissions			
Market Orders	9,512	9,563	19,075
Limit Orders	433,346	448,878	882,224
Market-to-Limit Orders	204	195	399
Iceberg Orders	2,568	2,514	5,082
Panel B: Average Size			
Market Orders	23,888	24,261	23,959
Limit Orders	34,918	34,382	34,639
Market-to-Limit Orders	21,056	17,807	19,522
Iceberg Orders	630,990	629,804	630,655
Panel C: Cancellations and Executions			
Limit Order Executions	103,044	100,777	203,821
Limit Order Cancellations	330,303	348,100	678,403

Table 3.3 gives an overview over the average number of order submissions and average order size (in Euros) in my data sample. The averages are computed over all 30 stocks in the blue-chip segment of FSE. Additionally, limit orders are split up into executions and cancellations. Column 1 lists the bid side of the book, column 2 the ask side and the third column gives figures for all orders.

depth, this volume is so large that anyone who submitted such a large limit order would run the risk that the order would be interpreted as an informed trade. This in turn would most probably lead to adverse price movements. An iceberg order is way to hide large volumes. It must also be said that very often only the tip of the iceberg gets executed and the rest is cancelled.

While the previous figures address average liquidity consumption, the data also allows to compute how far individual market orders walk up the book: 84% of all transactions take place within the best prices, while the remaining 16%

Figure 3.2: Histogram of Market Order Ticks

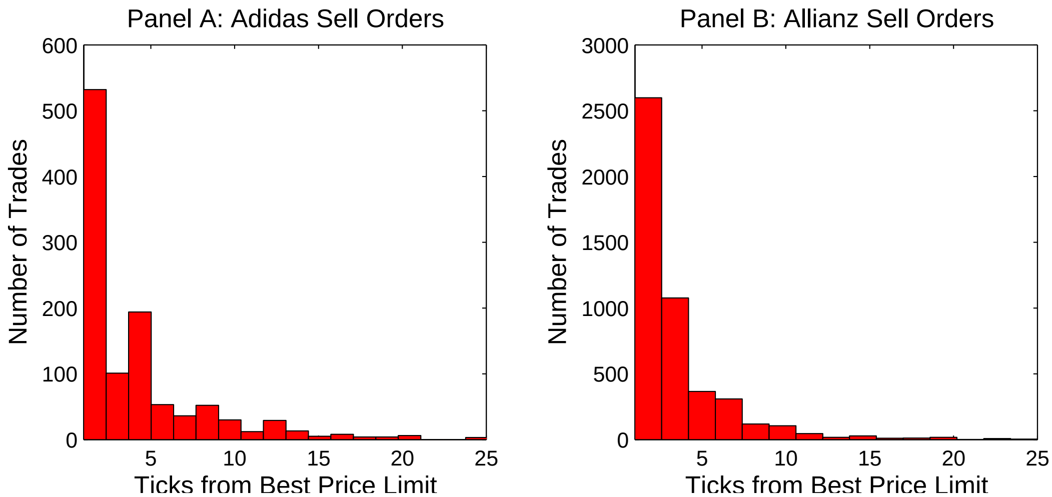

Figure 3.2 shows two histograms of the market order impact for two stocks, Adidas and Allianz. The x-axis gives the number of ticks that the market order walks up the book beyond the best price. The y-axis counts the number of events.

consume more liquidity than available at the best price. Figure 3.2 shows a histogram of the number of ticks that market orders walk up the book if they are too large to get executed at the best price. In this case, the largest part of all market orders (in the figure Adidas and Allianz) got settled within the first few ticks. However, some blocks clear the complete depth up to 25 ticks or more. These figures vary strongly from stock to stock. However, in all they show that the market as a whole is very liquid, that it mostly accommodates even large trades within the best spread and that even if the volume is too large, most of the matching takes place only very few ticks from the best price. Likewise, nearly all limit order activity takes place in this region as well: limit orders that are further from the best price because of price movements get cancelled or are

Figure 3.3: Time-of-day Effects

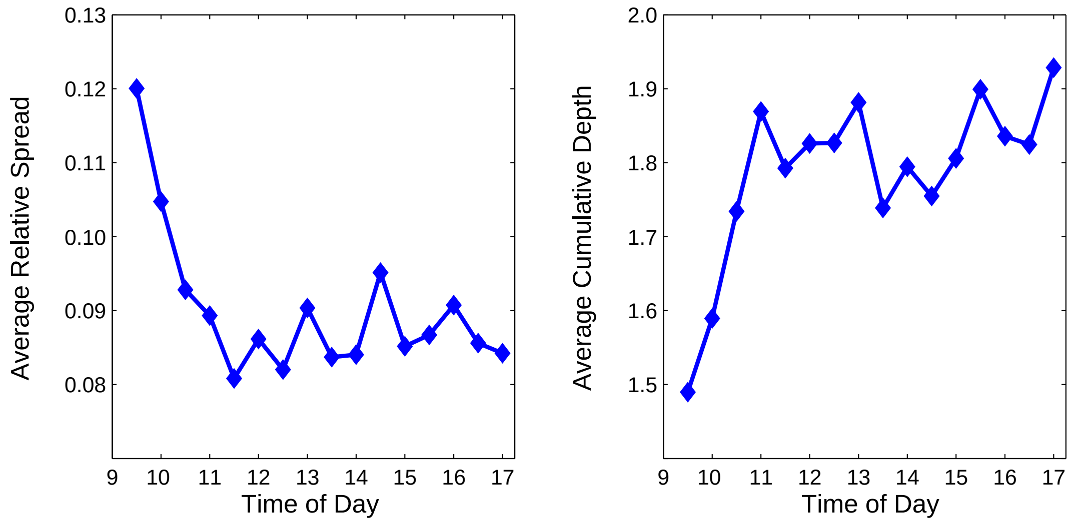

Figure 3.3 shows typical time-of-day effects for spreads and depth. The x-axis gives the time of day and the y-axis the average relative spread (in %) or the average depth (in 100,000 of Euros). Spreads exhibit the typical decreasing pattern, depth displays the typical increasing pattern.

revised to have new and more competitive price limits. Nearly all new liquidity flows into the book within the first ticks from the best price. Limit orders that are far away from the best price are mostly stale orders that have been in the order book for a longer time; new limit orders are only seldom submitted far from the best price.

Finally, I illustrate time-of-day effects in Figure 3.3. Previous literature such as Wood, McInish and Ord (1985), Jain and Joh (1988), Foster and Viswanathan (1990) or McInish and Wood (1992) documents that liquidity shows strong seasonal patterns, in particular on an intraday basis. I therefore plot average spreads against their specific time of day and likewise for depth. Fig. 3.3 shows that, consistent with previous empirical evidence, spreads are higher and depth is lower at the opening of the trading day. Spreads exhibit a pronounced L-shape over the time of the day, while depth exhibits the corresponding upside-down L-pattern. The figures for the market order flow are qualitatively identical. In other words, the variation of liquidity contains a deterministic time-of-day component.

Chapter 4

Resiliency of the Limit Order Book

This chapter investigates the resiliency of an electronic limit order book where there is centralized aggregation of liquidity and depth. I define resiliency as the speed with which the temporary erosion of liquidity is corrected through the inflow of new limit orders. I find strong evidence that the resiliency of each stock is consistently high. It is strongest around the best price and gets steadily weaker further from the best price. Furthermore, resiliency depends on microstructural determinants like order arrival rates, trader patience, trading volume, uncertainty and informed trading. Cross-sectionally, it has a high association with large market-capitalized and high-beta stocks. It is not strongly correlated with either spread or depth which reinforces its importance as an independent dimension of liquidity.

4.1 Introduction to Resiliency

It is widely recognized that market liquidity cannot be captured by a single measure. The seminal literature on liquidity (Garbade (1982), Kyle (1985), and Harris (1990)) identifies three main dimensions of liquidity: spread, depth and resiliency. Spread is the price dimension and represents the transaction costs faced by public traders, and is often measured by the quoted bid-ask spread or the trade-based effective spread. Depth is the quantity dimension and reflects the market's ability to absorb and execute large orders with minimal price impact, and is often measured by the quoted depth or by Kyle's Lambda. Finally, resiliency is the time dimension. In the context of a limit-order-book market, following Garbade (1982), resiliency is the speed with which the temporary erosion of the limit order book that is caused by a large uninformative order-flow shock is corrected through the flow of new orders into the market. It relates to a liquidity-induced reduction of the spread and replenishment of depth, not to information-induced price changes.

This chapter is on resiliency, the time dimension of liquidity. Resiliency addresses a question that is very important for market participants, stock exchanges and regulators, particularly in the context of order-matching market systems. Public traders in such markets potentially face significant price risk and execution risk when they wait for the price and depth to bounce back to normal levels after a large trade. For example, this happens to traders of large blocks who break up their trades into smaller blocks for a better execution price, and have to wait for non-trivial time periods before they can execute successive blocks. Arbitrageurs who typically work on large-volume-small-margin strategies also face similar risks; and the ability of arbitrageurs to arbitrage away small price discrepancies is essential for fair pricing and market integrity. With the market for supplying liquidity becoming increasingly competitive, and often transcending national boundaries, stock exchanges should arguably have a strong interest in understanding the replenishment mechanism of the order book in order to be able to attract and retain liquidity. Likewise, it is important for regulators to understand the resiliency dimension of liquidity in-depth, in order to factor an analysis of resiliency into their monitoring of market quality and stability.[19]

Spreads have been heavily researched: the literature is far too extensive to adequately summarise here.[20] Depth has also been reasonably well-researched.[21] Surprisingly, even though resiliency provides a key insight into the nature of liq-

[19] In a static context, Glosten (1994) and Seppi (1997) compare limit order markets and hybrid systems. In a dynamic context, Huang and Stoll (1996), Parlour (1998), Foucault (1999), Parlour and Seppi (2003) and Foucault, Kadan and Kandel (2005) study the equilibrium of the limit order book and order submission strategies, but they direct little attention to the stability of limit order book markets.

[20] One strand of the literature decomposes the spread into three components: one component reflecting the inventory holding risk of liquidity suppliers, another component reflecting the adverse-selection losses that liquidity suppliers make to more informed investors, and the last component reflecting order-processing costs. See, for example, Huang and Stoll (1997), Stoll (1989), and Glosten and Milgrom (1985). Another strand of the literature focusses on the individual trade-based effective spread, and its decomposition into the adverse selection spread and the realised spread. See, for example, Huang and Stoll (1996), Bessembinder (1997) and Naik and Yadav (2003).

[21] In particular, Hasbrouck (1991) and Kempf and Korn (1997) have analysed the effect of transactions on market prices. Additionally, for example, Glosten and Harris (1988) and Brennan and Subrahmanyam (1996) investigated the relationship between stock returns and measures of depth, similar to Kyle's Lambda.

uidity supply in the market, we know relatively little about the empirical properties of resiliency.[22] In particular, we do not know to what extent micro-structural factors like trading activity, uncertainty and asymmetric information, contrarian trading and momentum trading affect resiliency, we do not know what stock-specific factors determine resiliency in the cross-section, and we do not know if resiliency is related to the other dimensions of liquidity in some way or provides independent new information. This chapter aims to fill this major gap in the literature.

In this chapter, I investigate resiliency in an electronic limit order book market. There are several reasons for choosing to investigate resiliency in an order book setting rather than a dealer market setting.

1. With the enormous proliferation and growth in electronic order matching systems, stock exchanges around the world are increasingly organised as electronic order-driven markets. Except for the New York Stock Exchange, the NASDAQ and the London Stock Exchange, major stock markets conduct trading almost exclusively through open electronic limit order books.

2. An electronic limit order market is more crucially dependent on the exis-

[22] Coppejans, Domowitz, Madhavan (2003) do analyze the time variation of order book depth, though they do not consider changes in the liquidity flow over time. Gomber, Schweickert and Theissen (2004) measure the time it takes for a liquidity shock to dissipate to half its size. A related strand of literature investigates the submission behavior and order aggressiveness of market participants. Biais, Hillion and Spatt (1995) empirically analyze market order flows and order aggressiveness. They document substantial serial correlation of market orders. Ranaldo (2002), Ahn, Bae and Chan (2001), Bae, Jang and Park (2003) or Grammig, Heinen and Rengifo (2004) take such analyses further. In contradistinction, the focus in this chapter is on the continuous refreshment mechanism of the order book.

tence of adequate resiliency relative to a dealer or a hybrid market structure. Unlike a dealer market, a limit order book market depends only on limit order submissions for new liquidity. This raises the issue of how a limit order book market can ensure that enough new liquidity is submitted to the book as liquidity gets consumed. Understanding the replenishment mechanism requires a dynamic view of the limit order book.

3. Limit order books potentially allow a cleaner estimate of resiliency. In an order book context, the resiliency is, quite literally, the rate of mean reversion in the spread and the depth of the order book, with adequate controls for the information content in trades. This is relatively straightforward to observe, since the spread and depth of the order book can be measured with precision. Furthermore, this approach does not require knowledge of the "true price". While there exist ways in which such estimations of resiliency based on pricing errors can be made,[23] the associated estimates are considerably more noisy; we do not encounter this problem with limit order book resiliency.

4. Agents who wish to submit a new limit order can do so at a price/tick of their choice. This reveals an important facet of order-book based resiliency: the replenishment of the order book can take place at different points of the price-quantity schedule. If limit orders are submitted far from the former best price, implicit transaction costs stay high. However if the book is

[23] See Holthausen, Leftwich and Mayers (1990) and Dong, Kempf and Yadav (2005).

> refilled close to the former best price, transaction costs get reduced again very quickly. Clearly, the inflow of new liquidity close to the best price is more valuable to investors than at prices far from the best price. Therefore, resiliency can only be addressed adequately if the spread and depth at different ticks are considered. This is a nuance that makes a limit order book setting mandatory for the analysis of resiliency.

I investigate resiliency using limit order data from the electronic trading system Xetra at the German stock exchange in Frankfurt. I choose the German market over the markets in the US and the UK because, unlike these other markets, the German market provides a purely order-driven setting without any dealers and without significant lateral linkages to external liquidity suppliers or liquidity supply systems. The behavior of limit order traders is hence not influenced by resiliency supply from external sources. The trading platform of the Frankfurt Stock Exchange faces no competition from ECNs and hardly any competition from regional exchanges, and virtually all available liquidity is aggregated in the centralized limit order book. Limit order traders in the German market face virtually no competition from an upstairs market as in the UK or the US. Grammig and Theissen (2005) report for Germany that, in 2002, only 1.5% of trades (constituting only 0.25% of market value) went through the upstairs market. The limit order book of the Frankfurt Stock Exchange is affected by few external factors, and offers a clinically uncontaminated view of the behavior of limit order traders. Hence, it is well-suited for an investigation of resiliency.

This chapter empirically investigates, for the first time, the main features of resiliency as a dimension of liquidity in an electronic limit order market. Resiliency, to reiterate, addresses the following question: when trades, especially those resulting from relatively large and uninformative orders, consume liquidity by eroding the limit order book, how fast is liquidity replaced through the competitive actions of market traders. Resiliency results from the interaction of liquidity flowing into the market and liquidity being taken out. The inflow comes from the submission of new limit orders, while the outflow results either from the cancellation of limit orders or the execution of limit orders against newly submitted market orders. Together, inflow and outflow determine the evolution of the price-quantity schedule.

Specifically, I first address how I can formally measure resiliency. I set up a mean reversion model of liquidity to capture the dynamics of the spread and depth over time, and examine the relation between current and past liquidity flows. I examine ask-side and bid-side resiliency separately, and also analyze a range of different data frequencies. Second, I analyse the micro-structural time-series factors that affect resiliency, in particular, information asymmetry, uncertainty and trading activity. Third, I analyse the variation in resiliency across stocks. And finally, I examine the relationship between resiliency and the other two liquidity dimensions: spread and depth. My empirical investigation is based on three-months data on the thirty stocks that consitute the DAX.

I find strong evidence that the liquidity dynamics of the order book follow a stable replenishment process: the resiliency for each stock is consistently high and stable across different horizons. Empty limit order books are refilled quickly, while full books attract less liquidity. Most clearly I observe resiliency in the behaviour of liquidity suppliers around the first few ticks of the book, both for the reduction of the spread as well as for the provision of new depth: resiliency is strongest around the best price and gets steadily weaker the further I move away from the best price in the book. Clearly, trades that are executed against the book take away liquidity at the first few ticks, and traders who actively monitor the book jump in straight away to exploit these profit opportunities in the book.

I also find that resiliency is dependent in a robust manner on microstructural determinants. As predicted by Foucault, Kadan and Kandel (2005), the resiliency of the spread increases in the proportion of patient traders and decreases with the order arrival rate and at the end of the trading day. Furthermore, it decreases with the amount of trading volume, while volatility affects buy and sell side resiliency asymmetrically. The effects of these determinants on the resiliency of depth are contrarian. The most probable explanation is that, in the time series, spread resiliency and depth resiliency are displaced effects. When spread resiliency is high, price-improving limit orders erode the depth at the best ticks and therefore reduce depth resiliency. In the cross section, the results show very consistently that firms with high spread resiliency also have high depth resiliency. In particular, resiliency has a high association with large market-capitalized and high-beta stocks. I also find that resiliency is not significantly correlated with either spread or depth. This reinforces, ex post, the importance of resiliency as an independent dimension of liquidity. It cannot be seen or assumed as a replica of the price or quantity dimension. This time dimension of liquidity provides significant new information.

The chapter is organized as follows: Section 4.2 gives a brief outline of the liquidity proxies that are used in this chapter. Section 4.3 sets up a simple framework of resiliency which outlines the concept of resiliency and the main hypotheses. In Section 4.4 I present the empirical evidence of resiliency. In Section 4.5 I explore the interaction of resiliency with microstructural determinants in the time series. In Section 4.6 I analyze cross-sectional differences across stocks. Section 4.7 shows the link between the resiliency dimension of liquidity and other liquidity measures. Section 4.8 concludes.

4.2 Construction of Liquidity Measures

In the following section I construct liquidity measures on the basis of each stock's limit order book. As discussed above, liquidity is a property of an asset that ensures that the asset can be traded at any time without high price impacts, yet it is not obvious what proxies to use. In a limit order book market, the supply of liquidity results from the submission of limit orders, while the consumption can result from market orders being executed against the book or limit orders being cancelled. Models such as Glosten (1994) or Sandås (2001) characterize the limit order book by its price schedule. In line with their approaches, I choose two liquidity measures that capture the information of the price impact function: the depth and spread of the book. The limit order data I use for the measures are order book snapshots for 5-minute intervals.

I define the depth as the number of shares in the order book. Let n_A^k and n_B^k denote the number of shares at tick k. At the best price in the book $k = 0$, one tick away from the best price k is 1 and so forth. A or B indicate the ask and bid sides of the order book. For each stock i I compute the cumulative depth in the limit order book at any tick from the best price. For example, at tick 3 I obtain $DEP_{A,i}^k = \Sigma_{k=1}^3 n_{A,i}^k$ and $DEP_{B,i}^k = \Sigma_{k=1}^3 n_{B,i}^k$. By construction, each tick comprises the depth of the previous tick. Because traders can choose where to submit limit orders in the book, the different tick regions reflect regions of different liquidity supply behavior. To keep the analysis tractable, I consider depth up to a maximum of ten ticks from the midquote. More than 99% of all trades are executed within this range of the order book, so it seems reasonable to restrict ourselves to the active part of the book.[24]

Beside the depth of the limit order book I also compute the spread that any trader has to pay in excess of the midquote. Let l be an index of the number of limit prices in the limit order book. At the best price $l = 0$, at the next best price $l = 1$ and so forth. Furthermore, let us denote any price in the book by p and the midquote by MQ. For each stock I compute half-spread measures for

[24] I do not include the hidden part of iceberg orders in aggregate order book depth. Therefore I use the exact limit order book that market participants observe and use to condition their behavior on. Hidden liquidity is, by definition, not visible and does not belong to the decision set of investors.

different prices in the limit order book. For example, $p_A^0 - MQ$ is the half-spread that an investor pays at the best ask price. $p_A^3 - MQ$ is the half-spread that is incurred for the last unit of trade if the trade walks up three more steps in the limit order book. The corresponding half-spread on the bid side is $MQ - p_B^3$. While depth measures the volume of liquidity, the half-spread measure captures the quality of liquidity in terms of its price difference to the perceived fair value (which is assumed to be the midprice).

4.3 Framework and Hypotheses

Resiliency refers to the dynamic dimension of liquidity which measures how prices and quantities in the order book evolve over time. In many studies, the time dimension of liquidity is ignored. While it is straightforward to determine the costs that one single trade would incur at a single point in time, it is difficult to determine the costs of a sequence of trades. Firstly, any early transaction will have an impact on the market price and liquidity which will then affect future lots of the same trade. Furthermore, the submission of orders transmits a signal to the market and can change the course of market events in an unpredictable way. To capture the time dimension, I set up a model of the observed order book dynamics and then include the interaction with market events.

Departing from the Garbade (1982) definition, I expect that a resilient limit order book will get refilled as soon as its liquidity has been consumed. In other words, I investigate the relationship between the past level of liquidity and current liquidity flow. Foucault, Kadan and Kandel (2005) derive an equilibrium range of spreads with a fixed upper and lower boundary. They define resiliency as the probability that spreads revert back to their competitive level (see Chapter 2). Consistent with this notion, I set up a mean reversion model of liquidity. Let L_t denote liquidity a time t and let ΔL_t be the increment of liquidity from $t-1$ to t, $\Delta L_t := L_t - L_{t-1}$. I model liquidity as a stochastic process that consists of a mean reversion component and a stochastic increment. For the empirical implementation, let α and φ be the coefficients to be estimated. This yields the specification

$$\Delta L_t = \alpha - \varphi L_{t-1} + \varepsilon_t \tag{4.1}$$

where φ is an estimate of κ. It measures the intensity of mean reversion which depends on the level of liquidity, L. The higher φ is, the stronger the pull-back effect of liquidity to its long-run mean is and thus the higher resiliency is.

One implicit assumption of Equation 4.1 is that Δt is equally spaced. As I use order book snapshots at a fixed time interval, this assumption is not critical. However, validity of the equal spacing for the cross-section of stocks is not that clear. Comparing the resiliency of small and large stocks against the same time frame might only measure differences in size instead of resiliency. This is an argument in favor of choosing the time interval in relation to the trading intensity of a stock. On the other hand, any trader who needs to liquidate a large position is bound to the same time horizon for all stocks. From a trader's point of view, using the hourly clock instead of the trading clock is the more appropriate approach. I follow the latter avenue and use the hourly clock keeping mind that it need not be the only valid specification.

A second issue to discuss is the choice of the best-suited frequency. In electronic markets, traders can monitor the limit order book more or less continuously. In times of heavy trading activity, several limit orders are processed within a second. Market orders take place about every 30 seconds for active stocks and up to every five or six minutes for less active stocks – for example every 30.60 seconds for Deutsche Telekom and every 6 minutes and 31.80 seconds for Fresenius. If I choose a frequency that is too long I will not be measuring the liquidity adjustment of the book to a trade because very many trades will have taken place. If we choose a frequency that is too short we will have too many intervals in which no trades take place and liquidity adjustments are not order-induced. Since I assume that not all market participants react instantaneously I allow some time until the discovery of the cleared order book and compute resiliency at a 5-minute, 15-minute and 60-minute frequency.

Market participants who trade assets to do arbitrage or rebalance portfolios will need to know how resiliency interacts with market events and microstructural factors to assess favorable and unfavorable moments for trading. I depart from the limit order book model in Foucault, Kadan and Kandel (2005) which provides the following hypotheses (see Chapter 2):

- Resiliency increases (φ rises) with the patience of traders.

- Resiliency decreases (φ falls) with the order arrival rate.
- Resiliency decreases (φ falls) with at the end of the trading day.

In addition, the demand curve literature[25] makes predictions about prices in the presence of trading volume and information. The information hypothesis states that an adverse impact of informative trades on the price will not be reversed. This implies that less new limit orders get submitted in information-intensive periods. The price pressure hypothesis states that trading volume only has a temporary effect on the price until liquidity suppliers reverse the price impact. This implies that trade-intensive periods stimulate the submission of new limit orders and thus increase resiliency. Finally, Foucault (1999) and Handa, Schwartz and Tiwari (2003) predict an asymmetric effect of bad news on limit order submissions: buy limit orders are submitted more cautiously and sell limit orders more aggressively.[26] Let us summarize the implications for resiliency:

- Resiliency decreases (φ falls) in the presence of informed trading.
- Resiliency increases (φ rises) with trading volume.
- Bad news has asymmetric effects on ask and bid side resiliency.

[25] The demand curve literature discusses whether the demand curve of stock prices is flat or sloped. For more details and evidence see Harris and Gurel (1986), Shleifer (1986) Dhillon and Johnson (1991), Beneish and Whaley (1996) or Kaul, Mehrotra and Morck (2000).

[26] Holthausen, Leftwich and Mayers (1990) and Saar (2001) make a similar point by arguing that the ask and bid side of the limit order book behave asymmetrically. Brokers are particularly unwilling to take short positions to accommodate large block purchases, because they might be forced to buy the assets at unfavorable conditions later. In the case of a limit order book, the argument implies that traders on the bid side might be more hesitant to post limit orders.

In the following sections I estimate the mean reversion model of resiliency and analyze the impact of microstructural factors along time. As traders will choose more resilient stocks over less resilient stocks all things equal, I will also assess cross-sectional differences in φ.

4.4 Dynamics of the Limit Order Book

The previous section set up an econometric model for the estimation of resiliency. In the following section I estimate the resiliency of the depth and spread for each stock in the data sample. I then vary the tick at which liquidity is measured and use different frequencies to examine how robust resiliency is across ticks and time horizons.

4.4.1 Base Estimation of Resiliency

Let us begin the empirical implementation of Equation 4.1 by estimating the spread and depth reversal. I adopt a cross-sectional SUR estimation in which all stocks are pooled. The advantage of this approach is that it offers a lot of freedom to constrain parameters to be stock-specific or to vary across stocks.[27] The cross-sectional SUR approach yields joint GLS estimates which are corrected for contemporaneous correlation and heteroscedasticity. I start with a relatively long lag length and shorten the model by the usual t-statistics of the lag coefficients. I repeat the procedure until the lag length is significantly different from zero and

[27] To test the stability of the results, I also performed all following computations on a stock-by-stock basis. Qualitatively the results are the same, however the approach has the strong disadvantage that it is more difficult to assess the significance of coefficients across equations.

all further autocorrelation of the error term is eliminated. The constant and the mean reversion parameter are firm-specific while the parameters of the lagged values are constrained to be equal across stocks. I run the following regression:

$$\Delta L_{i,t} = \alpha_i - \varphi_i L_{i,t-1} + \Sigma_{k=1}^{n} \gamma_k \Delta L_{i,t-k} + \varepsilon_{i,t} \tag{4.2}$$

In the equation, L denotes liquidity. Furthermore, let us define ΔL as the liquidity change in the current period, $\Delta L_t := L_t - L_{t-1}$. In the estimation, I substitute L by the spread and depth. The parameter φ measures the mean reversion of liquidity while the γ parameters are lag coefficients. Lags are included up to a length of 20; beyond that they are not significant anymore and the usual diagnostic checks show no evidence of serial correlation. ε is a normally distributed white noise error term. If φ in the above model equals zero, the equation is entirely in first differences and will have a unit root. The t-statistic will not follow the usual t-distribution anymore; instead I test for the presence of a unit root by means of the augmented Dickey-Fuller critical t-values. As the model only contains an intercept and no time trend, the correct value is the so-called τ_μ statistic which is 3.43 at the 1% significance level and 2.86 at the 5% level.

Table 4.1 gives the results for the resiliency of the limit order book at the third tick for depth (i.e. 0.03 Euros from the best price) and at the third price (i.e. the third-best price after the best price) for the half-spread. The mean reversion parameter is significantly positive for every single stock both for the depth and

Table 4.1: Resiliency of Order Book Liquidity

	Depth				Spread			
	Ask		Bid		Ask		Bid	
	φ	t_φ	φ	t_φ	φ	t_φ	φ	t_φ
Adidas	0.26	17.40	0.24	15.73	0.29	18.58	0.28	17.78
Allianz	0.42	26.10	0.42	25.50	0.38	27.05	0.37	25.97
Altana	0.28	18.46	0.29	19.25	0.30	18.88	0.27	17.66
BASF	0.37	23.67	0.39	23.99	0.37	24.68	0.37	24.47
Bayer	0.25	17.12	0.33	21.17	0.40	27.74	0.38	26.31
BMW	0.14	10.92	0.30	18.95	0.34	22.96	0.35	22.57
Commerzbank	0.18	12.80	0.21	14.47	0.38	24.20	0.36	22.25
Continental	0.20	13.66	0.23	15.00	0.22	14.49	0.29	18.28
DaimlerChrysler	0.37	24.15	0.24	16.73	0.42	40.58	0.40	38.47
Deutsche Bank	0.39	24.34	0.40	24.19	0.38	26.34	0.36	24.67
Deutsche Boerse	0.23	15.38	0.33	19.19	0.26	16.49	0.27	17.34
Deutsche Post	0.13	10.23	0.27	17.75	0.29	18.20	0.29	18.61
Deutsche Telekom	0.11	9.63	0.13	10.48	0.44	36.24	0.42	30.88
E.ON	0.43	27.17	0.30	19.68	0.40	32.97	0.38	29.01
Fresenius	0.17	12.21	0.15	10.24	0.23	15.19	0.21	14.14
Henkel	0.28	18.34	0.26	17.17	0.28	18.06	0.28	17.97
Infineon Technologies	0.16	13.45	0.16	12.20	0.41	26.98	0.39	25.35
Linde Lufthansa	0.30	19.56	0.23	15.58	0.27	17.00	0.27	17.33
Lufthansa	0.20	13.29	0.25	16.45	0.42	40.11	0.39	36.88
MAN	0.25	16.41	0.24	16.54	0.25	16.17	0.29	18.16
Metro	0.30	19.02	0.25	16.95	0.33	20.99	0.29	18.05
Münchener Rück	0.40	25.35	0.39	23.77	0.35	24.03	0.32	22.09
RWE	0.35	22.32	0.33	20.87	0.39	24.82	0.37	23.49
SAP	0.38	23.98	0.43	26.03	0.35	23.08	0.36	23.01
Schering	0.35	22.46	0.27	17.99	0.39	29.09	0.39	26.72
TUI	0.20	13.73	0.24	15.95	0.30	18.62	0.29	18.76
Volkswagen	0.35	22.47	0.33	23.53	0.40	29.14	0.37	26.44
ThyssenKrupp	0.23	15.79	0.27	17.81	0.35	21.48	0.34	21.43
HypoVereinsbank	0.16	13.57	0.14	11.36	0.42	40.51	0.40	38.33
Siemens	0.37	23.87	0.39	24.35	0.40	28.46	0.39	27.02
Aggregate	0.22	41.57	0.24	41.02	0.34	46.91	0.33	46.34

Table 4.1 reports the regression results for the resiliency of the limit order book. The columns give the resiliency parameter φ_i with its corresponding t-values for the ask side and the bid side of the order book. The results are reported both for the depth and spread of the limit order book. The final row gives the resiliency parameter if constrained to be equal across stocks.

the spread of the order book. The last row of the table summarizes these results by repeating the estimation and constraining the mean reversion parameter to be the same across stocks. For depth, the constrained parameter φ is 0.22 with a t-statistic of 41.57 on the ask side and 0.24 with a t-statistic of 41.02 on the bid side. For the half-spread, the constrained parameter φ is 0.34 with a t-statistic of 46.91 on the ask side and 0.33 with a t-statistic of 46.34 on the bid side. In other words, deviations from the average level of liquidity are reversed by about 20 to 30% in the next time interval. As discussed above, positive estimates of φ imply that the inflow of new liquidity to the order book is strongest if the past level of liquidity was low. Liquidity flows to the book faster if the order book has been cleared, which is evidence of resiliency. Comparing the ask side and the bid side results I observe that there is little difference between the coefficients on the buy side and sell side (0.22 vs 0.24 and 0.34 vs 0.33). The R^2 values are around 0.35 for depth and 0.41 for the spread. All Durbin-Watson statistics are very close to 2, which implies that there is no autocorrelation in the error term.

4.4.2 Order Book Tick and Time Horizon

Table 4.2 reports the behavior of the mean reversion parameter if I vary the number of ticks for which I determine the depth and the half-spreads. The tick in the left column indicates the maximum tick or step in the limit order book for which I measure liquidity. The table displays estimates constrained

Table 4.2: Resiliency at Different Ticks

	Depth				Spread			
	Ask		Bid		Ask		Bid	
Tick/ Step	φ	t_φ	φ	t_φ	φ	t_φ	φ	t_φ
0	0.40	44.43	0.41	45.03	0.47	50.31	0.47	50.31
1	0.29	43.35	0.30	41.50	0.43	49.01	0.42	48.09
2	0.22	41.82	0.24	39.94	0.38	47.31	0.38	47.44
3	0.22	41.57	0.24	41.02	0.34	46.91	0.33	46.34
4	0.15	39.05	0.17	38.81	0.30	45.01	0.30	45.33
5	0.13	38.88	0.15	37.99	0.27	44.89	0.26	45.02
10	0.08	36.57	0.08	33.70	0.16	40.23	0.17	40.99

Table 4.2 reports the regression results of limit order book resiliency for the depth and half-spread. In the case of depth, the measures are computed for different tick sizes and, in the case of the half spread, for different steps in the limit order book.

over all stocks.[28] The estimates in Table 4.2 highlight that the limit order book is strongly resilient irrespective of the tick which is considered. However, the strength of mean reversion gets less with increasing tick size. For the depth of the limit order book, the coefficient φ drops from around 0.40 (on the ask side) at the best price to 0.08 ten ticks away from the best price. For the spread, φ is about 0.47 (on the ask side) and drops to 0.16 ten steps from the best price. The bid side behaves in the same way. In all, resiliency at limit prices close to the prevailing best price is strongest. Most of the action in the limit order book takes place in a very limited range. When the order book is empty, new limit orders get submitted to gain price priority and have a high probability of getting

[28] Again the computations were also performed with non-constrained coefficients and yielded qualitatively identical results.

Table 4.3: Resiliency at Different Frequencies

	Depth		Spread	
Frequency	Ask	Bid	Ask	Bid
5 minute intervals	0.223	0.241	0.349	0.338
	41.573	41.017	46.910	46.344
15 minute intervals	0.391	0.432	0.538	0.559
	41.516	42.175	45.662	46.600
60 minute intervals	0.712	0.777	0.871	0.873
	52.489	52.778	53.347	56.803

Table 4.3 reports the regression results of order book resiliency for liquidity at different frequencies. The table reports the ask and bid side results for both the depth and the half-spreads in the limit order book. The depth of the limit order book is computed for the first three ticks and the spread is computed for the third step of the book.

executed. Thus the replenishment mechanism is strongest around the best prices and we see the strongest evidence of resiliency in that region.

To complete the estimation of the basic model let us focus on the impact of the trading frequency on the resiliency estimates. Table 4.3 displays the resiliency estimates obtained at different frequencies. For the sake of brevity, I only report the constrained estimates across all stocks. Again, the depth in the limit order book is computed up to the third tick, while the spread refers to half-spreads at the third step in the limit order book. The table shows that all results remain strongly significant over the different frequencies. It is evident that for all measures the resiliency parameters become stronger for longer intervals. While the level variable stays the same, the difference from $t-1$ to t can become much larger if the trading interval is longer. Thus, the process exhibits stronger mean reversion. We see that mean reversion is significant at various frequencies as documented by the high t-statistics; it is a robust phenomenon in high-frequency data as well as longer intervals. From an economic perspective, it seems plausible that the amount of new liquidity which is submittted after an eroded order book increases with the time horizon.

4.5 Interaction of Resiliency and Microstructural Factors

In this section I investigate how resiliency interacts with microstructural factors such as information asymmetry, trading activity and uncertainty. The first subsection explains the construction of the microstructural proxies, while the second subsection gives details on the estimation procedure and presents the results.

4.5.1 Construction of Microstructure Proxies

In section 4.3 I discussed that resiliency should be associated with the order arrival rate, the patience of traders, the information intensity of trading, the trading volume and bad news. Furthermore, I expect a different level of resiliency at the end of the trading day. Let us now turn to the construction of these determinants.

I measure the information intensity of the trading period by proxying the probability with which trades were submitted by informed traders. A prominent summary measure of informed trading is the PIN measure whose empirical implementation is developed in Easely, Kiefer, O'Hara and Paperman (1996), Easely, Kiefer and O'Hara (1997) and Easley, Hvidkjaer and O'Hara (2002). PIN departs from the order imbalance in a stock to derive a probability measure. Since it has to be computed over various days and I need a measure for each period, I proxy PIN by the imbalance between buy and sell market orders. This proxy assumes that the order imbalance reflects information directly which is the closest I can get to PIN for high-frequency levels. To compute the measure, I determine the number of market orders from $t-1$ to t on the ask side and the bid side of the order book, $nMO_{A,t}$ and $nMO_{B,t}$. I can then calculate the surplus of market orders for each order book side which corresponds to a simple measure of order imbalance OI:

$$OI_{A,t} = nMO_{A,t} - nMO_{B,t} \tag{4.3}$$

$$OI_{B,t} = nMO_{B,t} - nMO_{A,t} \tag{4.4}$$

$OI_{A,t}$ and $OI_{B,t}$ are identical by construction except for their sign. The reason for this convention is that suppliers of sell limit orders should react to a sell surplus in the same way as suppliers of buy limit orders do to a buy surplus. As the absolute value of the surplus has no real economic meaning, I finally transform order imbalance into a 0-1-variable where the highest 10% of imbalances receive a 1 and all other values are set to zero. This dummy variable isolates the effect of periods in which trading is information-intensive.

The next determinant that I turn to is trading activity. I proxy trading activity by the market order volume submitted in a given interval. Let $MO_{A,t}$ denote the volume of all market orders on the ask side between $t-1$ and t. Likewise, let $MO_{B,t}$ denote market order volume on the bid side of the order book between $t-1$ and t. I then obtain trading volume as the sum of all buy and sell market orders:

$$TRV_t = MO_{A,t} + MO_{B,t} \tag{4.5}$$

As trading volume is clustered very strongly in some periods I take the log values of volume as the proxy for trading volume.

I proxy bad news by the unexpected component of volatility. To determine unexpected volatility I employ the Bollerslev (1986) GARCH approach. I choose a GARCH(1,1) specification and model the variance of stock returns conditional on past squared residuals. At this point however I have the problem that conditional volatility will no doubt be strongly correlated with trading volume. As I want to include both variables in the following estimations, I include trading volume in the estimation of conditional volatility as in Lamoureux and Lastrapes (1990). I thus succeed in disentangling volume and volatility. The equations of the model are:

$$r_t = \mu_{t-1} + \varepsilon_t \tag{4.6}$$

$$\varepsilon_t \mid (TRV_t, \varepsilon_{t-1}, \varepsilon_{t-2}, ...) \sim N(0, h_t) \tag{4.7}$$

$$h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1} + \alpha_2 h_{t-1} + \alpha_3 TRV_t \tag{4.8}$$

where r_t is the rate of return, μ_{t-1} is the mean of r_t conditional on past information and trading volume, α_0 to α_3 are the coefficients of the conditional volatility equation and h_t is the conditional volatility. The residuals from this model, $\varepsilon_{t.}$, yield the unexpected component of volatility, $UNXV$, which, by construction, will be orthogonal to TRV.

Furthermore, I construct the variables which are associated with resiliency in the Foucault, Kadan and Kandel (2005) model. Firstly, I compute the proportion of patient traders in the population in each interval as the number of limit orders (nLO), corrected for the number of cancellations (nCA) in relation to the overall number of orders in the interval:

$$PAT_{A,t} = (nLO_{A,t} - nCA_{A,t})/(nLO_{A,t} + nMO_{A,t} - nCA_{A,t}) \quad (4.9)$$

$$PAT_{B,t} = (nLO_{B,t} - nCA_{B,t})/(nLO_{B,t} + nMO_{B,t} - nCA_{B,t}) \quad (4.10)$$

The arrival rate is simply the sum of market and limit orders in an interval, again corrected for the number of cancellations:[29]

$$AR_{A,t} = nLO_{A,t} + nMO_{A,t} - nCA_{A,t} \quad (4.11)$$

$$AR_{B,t} = nLO_{B,t} + nMO_{B,t} - nCA_{B,t} \quad (4.12)$$

As with trading volume, the arrival rate has very strong peaks at times. Therefore I use logs again. Finally, I construct a 0-1-variable END_t which gets assigned a value of 1 if the observation is from the last 45 minutes of the day and zero otherwise. In all, I now have six variables (OI, TRV, $UNXV$, PAT, AT and END) which I use to condition resiliency in a time series framework.

[29] I follow Foucault, Kadan and Kandel (2005) in their definition of order arrivals. Strictly speaking the term arrival rate is misleading, because order arrivals are defined in terms of their absolute numbers (and not as a rate).

4.5.2 Impact of Microstructure Proxies on Resiliency

In the assessment of the interaction effects of microstructural factors and resiliency, I re-estimate the basic resiliency model, yet I include the microstructural determinants as conditioning variables. In particular, the resiliency parameters φ are now functions of the conditioning variables. The regression model is the following:

$$\Delta L_{i,t} = \alpha_i - \varphi_i L_{i,t-1} + \Sigma_{k=1}^{n} \gamma_k \Delta L_{i,t-k} + \varepsilon_{i,t} \tag{4.13}$$

$$\begin{aligned} \varphi_i = {} & \beta_{0,i} + \beta_1 OI_{i,t} + \beta_2 END_{i,t} + \beta_3 PAT_{i,t} + \beta_4 AR_{i,t} \\ & + \beta_5 TRV_{i,t} + \beta_6 UNXV_{i,t} \end{aligned} \tag{4.14}$$

In the above equations, φ_i in the top line is substituted by the functional relationship specified in the bottom two lines. OI is the dummy variable of information-intensive periods, AR is the order arrival rate, PAT is the proportion of patient traders, END is a dummy variable for the end of the trading day, TRV corresponds to trading volume and $UNXV$ is the proxy of bad news. The indices in the equations show that the conditioning variables are time-varying and stock-specific. In the SUR estimation, I let the base level of resiliency be stock-specific ($\beta_{0,i}$), however I constrain the impact of the conditioning variables to be the same for all stocks, therefore β_1 to β_6 carry no firm indices. This notion is appealing because if there is an economically meaningful link between resiliency and the microstructural determinants, it should be present for all stocks.

Table 4.4 shows the results for the resiliency of the limit order book. The

Table 4.4: Time Series Impact on Depth Resiliency

	Depth				Spread			
	Ask		Bid		Ask		Bid	
Variables:	β	t_β	β	t_β	β	t_β	β	t_β
Constant	-0.83	-39.83	-0.82	-37.91	-0.02	-43.57	-0.02	-44.01
OI	-0.01	-3.02	-0.01	-1.26	-0.00	-1.01	-0.01	-4.12
AR	0.03	12.70	0.02	8.99	-0.00	-3.86	-0.00	-0.04
PAT	-0.02	-1.36	0.03	2.26	0.02	3.19	0.01	1.89
END	0.01	1.24	0.02	3.79	-0.01	-2.05	-0.01	-4.35
TRV	0.01	4.90	0.01	6.42	-0.01	-13.83	-0.01	-10.26
UNXV	-1.88	-8.09	1.09	6.03	3.68	26.10	-3.83	-22.83

Table 4.4 reports the results of the SUR estimation for resiliency coefficients that have been conditioned on time series factors. The mean reversion parameter φ is a function of six time series factors (as specified in Equation 4.14): information intensity (OI), order arrival rate (AR), proportion of patient traders (PAT), time of the trading day (END), trading intensity (TRV) and bad news ($UNXV$).

table gives resiliency both for the depth as well as the half-spreads. A positive β coefficient means that resiliency is positively associated with the conditioning variable and a negative coefficient implies negative association. I can draw the following conclusions:

- With regard to order imbalance, OI, all coefficients are negative. For the resiliency of depth, the ask side coefficient of order imbalance is significant and for the resiliency of the half-spread, the bid side coefficient is significant. Negative coefficients imply that resiliency is lower in information-intensive periods and higher in the presence of non-informative trades. These results are evidence of the fact that resiliency is lower when the informativeness of trades is high. Economically this implies that limit order traders are not prepared to replace liquidity in the presence of informed trading; if they did they would lose money against better informed investors.

- We observe that the effects of the order arrival rate (AR), the proportion of patient traders (PAT) and the end of the trading day (END) on the resiliency of the spread are as predicted in Foucault, Kadan and Kandel (2005). The resiliency of the spread is significantly higher if the proportion of patient traders is high (both on the ask and bid side). In economic terms this means that, to a large extent, patient traders are responsible for the supply of new liquidity and therefore the replenishment of the limit order book. In contrary, resiliency is negatively associated with the order arrival rate (ask side coefficient is significant) and the end of the trading day (bid and ask coefficients are significant). This implies that a high number of new orders leads to less resiliency because new orders consume the supplied liquidity and therefore impede that the book gets refilled. The most plausible explanation for weak resiliency at the end of the trading day is that there is no incentive for limit order traders to submit new liquidity as it is less likely to get consumed before the close; therefore the book does not refresh as fast anymore.

- Furthermore, the effects of the arrival rate and the end of the day on the resiliency of the depth are the opposite to the effects on spread resiliency: depth resiliency is positively associated with the order arrival rate and the end of the trading day. The results for trader patience are mixed. This indicates that spread resiliency and depth resiliency are not synchronous. Rather, spread improvement and depth replenishment do not seem to take place in the same time periods. This result is plausible because high spread resiliency erodes depth by definition: a new best quote improves the spread, but it reduces depth as depth now only comprises one single order. Therefore, depth recovers later than the spread.

- Similar to the order arrival rate, trading volume also has a negative effect on the resiliency of the spread. Both the coefficients on the ask side and on the bid side are highly significant. Subsequently, the improvement of the spread is lower in environments with a high amount of trading. Again, the impact on the resiliency of the depth is opposite: when trading is high, limit order traders are less likely to improve the spread and more likely to increase depth. This is additional evidence that there is lagged relationship between spread and depth improvement.

- The effect of unexpected volatility on resiliency is asymmetric. For the resiliency of depth, the ask side estimate is significantly negative and the bid side is significantly positive. This implies that, in the presence of high unexpected volatility, resiliency is weaker on the ask side and stronger on the bid side. The evidence suggests an unwillingness of limit order traders to buy when volatility is high and a stronger willingness to sell. Again, the resiliency of the half-spread is the other way round.

All together, informed trading has a negative effect on resiliency as we expected in section 4.3. Likewise, the results for the resiliency of the spread confirm the hypotheses in Foucault, Kadan and Kandel (2005). Another striking point are the opposite effects of the time series variables on the resiliency of the depth. The most plausible explanation is that, in the time series, the resiliency of the depth and of the spread are not synchronous: fast spread improvements reduce the depth at the best tick and, likewise, large depth improvements require stable spreads. Therefore the effects are opposite.[30]

4.6 Resiliency in the Cross-Section

If resiliency varies over time and across stocks, the choice of resilient stocks will ceteris paribus lead to more successful trade execution. It will be more important, the smaller the per-unit profits and the higher the turnover volume of trading strategies. In this section I focus on the cross-sectional perspective of resiliency and examine what characteristics stocks with high resiliency share.

An important property of a stock is its risk. I measure the overall risk of a stock i by the volatility of its return, VOL_i. Furthermore, I use the beta factor of stock, BF_i, for its exposure to systematic market risk. It is computed over

[30] With the help of the conditioning variables I construct time series of resiliency and conduct Granger causality tests. I observe Granger causality both from spread resiliency to depth resiliency and vice versa. The study of the asynchronous effects of spread and depth resiliency would probably require an examination of spread improvements and depth reactions on a much finer frequency to see which effect leads the other. I leave this to future research.

a time series of stock returns that takes the DAX30 as the market portfolio. This data is provided publicly by the German Stock Exchange. Because the previous section showed that resiliency is negatively associated with informed trading, I also include a factor that measures such information asymmetries. In the cross section it seems plausible that informed trading has an impact on the liquidity supply if the losses of liquidity suppliers to informed traders are high. Therefore I take the stock return in each interval, sign it by the direction of order imbalance and take the sum of all signed returns. Under the assumption that order imbalance reflects information, this measure computes the overall profits of an informed traders, IP_i. Finally, I follow Banz (1981) and Fama and French (1992) in adding the firm size as a cross-sectional stock characteristic. Like in the literature, I use the log of market capitalization, MC_i, to measure size.

In a first step, Panel A of Table 4.5 shows the correlation of the factors among each other. Overall risk (volatility) and systematic risk (beta factor) are associated fairly strongly (0.465) which is not surprising. Otherwise, the correlations are low as we would expect, since their is not obvious relationship between these factors. Panel B shows the correlations of the depth and spread resiliency with these factors. The results are averages of the ask side and the bid side. Market capitalization and the beta factor are both positively correlated with resiliency. In economic terms this implies that investors are prepared to provide new liquidity fast if the stock is very large (and therefore usually well-known and heavily traded). Likewise, investors provide new liquidity quickly if

Table 4.5: Correlation of Cross-Sectional Factors and Resiliency

Panel A: Correlation Matrix of Cross-Sectional Factors					
		BF	IMB	MC	VOL
Beta Factor	BF	1.000	0.086	0.250	0.465
Informed Trader Profits	IMB		1.000	0.194	-0.126
Market Capitalization	MC			1.000	-0.161
Return Volatility	VOL				1.000

Panel B: Correlation Matrix of Resiliency and Cross-Sectional Factors					
		BF	IMB	MC	VOL
Resiliency of Depth	φ_D	0.085	-0.161	0.268	-0.198
Resiliency of the Spread	φ_S	0.356	-0.221	0.210	-0.223

Table 4.5 reports correlation of cross-sectional factors and resiliency measures. Panel A shows the top half of the correlation matrix of the cross-sectional factors amongst each other. Panel B shows the correlation of the resiliency measures with these cross-sectional factors. The critical value at the 5%-quantile is $\mp$ 0.153.

they know that the stock is a high-beta stock. In contrast, volatility and informed trader profits are negatively correlated with resiliency. Economically speaking, volatility and informed trader profits are viewed as risky properties of a stock which impede investors from readily supplying more liquidity; hence resiliency is lower and the correlation between these factors and resiliency negative. The correlation structure is consistent for spread and depth resiliency.

Let us now turn to the cross-sectional estimation. I re-estimate Equation 6.1 with informed trader profits, market capitalization, beta and return volatility as conditioning variables. The estimation procedure is the same as in the time series section: again I model the resiliency parameter φ as function of the conditioning

Table 4.6: Cross-Sectional Impact on Mean Reversion

		Depth		Spread	
Conditioning variables		Ask	Bid	Ask	Bid
Base Level of Resiliency	δ_0	0.34	0.27	0.27	0.29
		17.30	13.87	15.43	15.60
Informed Trader Profits	δ_1	-1.84	-1.36	-1.20	-1.07
		-8.98	-6.53	-8.17	-6.88
Beta Factor	δ_2	0.03	0.04	0.05	0.06
		2.72	4.02	7.54	8.68
Market Capitalization (logs)	δ_3	0.00	0.01	0.00	0.00
		0.48	3.90	1.36	1.07
Return Volatility	δ_4	-9.48	-6.59	0.63	-1.39
		-7.57	-5.16	0.58	-1.23

Table 4.6 shows the results for the estimation of Equation 4.16. I examine the cross-section impact of informed trader profits, the beta factor, market capitalization and return volatility on resiliency. The results are displayed for the ask and bid side of spread and depth resiliency. The first row is the estimate of the coefficient and the second row always gives the corresponding t-statistic.

variables. The model is:

$$\Delta L_{i,t} = \alpha_i - \varphi_i L_{i,t-1} + \Sigma_{k=1}^{n} \gamma_k \Delta L_{i,t-k} + \varepsilon_{i,t} \tag{4.15}$$

$$\varphi_i = \delta_0 + \delta_1 IP_i + \delta_2 BF_i + \delta_3 MC_i + \delta_4 VOL_i \tag{4.16}$$

In the above equations, φ_i in the top line is substituted by the functional relationship specified in the bottom two lines. Therefore, the mean reversion parameter becomes a function of certain variables. The subscripts in the equations show that the conditioning variables are not time-varying yet stock-specific. The parameters of these conditioning variables are assumed to be the same over all stocks, therefore δ_0 to δ_5 carry no firm indices. I allow the intercept of the resiliency model to be stock-specific, α_i.

Table 4.6 shows the results for the estimation of Equation 4.16. Firstly, δ_0 reflects the base level of resiliency if all conditioning variables are zero (which I call the “base level” of resiliency). Clearly, the base level of resiliency remains strongly significant for depth and spread resiliency. The coefficients are slightly lower than without conditioning variables. The coefficients of the cross-sectional factors confirm the relationships that we observed in the correlation structure. Informed trader profits and resiliency are negatively associated (with highly significant coefficients for both spread and depth resiliency): stocks for which informed traders make higher profits have a less resilient liquidity supply. This is plausible as liquidity suppliers fear providing liquidity to better informed traders who will then make a profit from the transaction. If they therefore anticipate insiders making profits they reduce their liquidity supply. In contrast, the beta factor has strongly positive relationship with resiliency. Stocks that have a high beta factor also have a high resiliency. This reflects that investors care about beta and see potential profits in providing liquidity to high-beta stocks. The results for market capitalization are not quite as strong: all coefficients are positive, however only one estimate is significant. This is evidence that large stocks tend to be more resilient, however the evidence is fairly weak. Finally, return volatility and resiliency have a negative relationship: more volatile stocks also have a less resilient liquidity supply. A plausible explanation is that volatility reflects risk and uncertainty; in risky and uncertain environments liquidity suppliers fear making losses as they cannot be sure about the true asset value. Therefore they reduce their liquidity supply which leads to lower resiliency. This result is very strong for depth resiliency (on the bid and ask side). However, volatility has no significant impact on spread resiliency.

An interesting point to note is how consistent the cross-sectional effects are for spread and depth resiliency. The relationship with informed trader profits, beta and market capitalization is exactly the same. Volatility only has an effect on depth resiliency, while it does not affect spread resiliency. If I compare these results to the time series results, I see that, on a more microstructural scale, spread and depth resiliency are not synchronous. This suggests a lead-lag relationship on a very high frequency. In the cross section however, stocks that have a high depth resiliency also tend to have a high spread resiliency; spread and depth resiliency are linked to the same firm characteristics in the cross section.

4.7 Relationship with other Liquidity Measures

Evidently, the limit order book of stocks shows a strong tendency to refill once it has been cleared. The mechanism is particularly strong for larger stocks with a high exposure to risk and for stocks with low informational asymmetries. A final question that I pose is whether this difference in liquidity is already taken into account for by other dimensions. To do this, I examine the correlation of the resiliency measures with the bid-ask spread and depth at the spread.

At first I re-estimate the resiliency model in Equation 6.1 on a daily basis and collect a time series of resiliency estimates for the spread and depth. Then I collect

Table 4.7: Correlation of Resiliency Measures

	SPR	DEP	RES_D	RES_S
Bid-ask spread (SPR)	1.000	-0.370	-0.224	-0.167
Depth (DEP)		1.000	0.030	0.124
Resiliency of the depth (φ_D)			1.000	0.463
Resiliency of the spread (φ_S)				1.000

Table 4.7 reports the correlation of the different dimensions of liquidity with each other. In particular I correlate the bid-ask spread and depth at the spread with resiliency. The table shows the results for the resiliency of the spread and depth. The diagonal consists of ones and only the top half of the table is reported for the sake of brevity. The critical value at the 5%-quantile is $\mp$ 0.153.

time series of the average daily bid-ask spread and the average daily depth at the spread.[31] Table 4.7 shows the correlation structure of spread and depth resiliency with the bid-ask spread and depth. We see that the bid-ask spread and depth are negatively correlated (-0.370). This implies that spreads are low when depth is high and vice versa. This result is as we would have expected. The bid-ask spread is negatively associated with resiliency (-0.224 for depth resiliency and -0.167 for spread resiliency). Depth has a positive relationship both with depth resiliency (0.030) and spread resiliency (0.124). Again, these results are not surprising: high resiliency reflects the fact that a stock is liquid with regard to the time dimension of liquidity. Therefore, high liquidity in the time dimension coincides high liquidity in the spread dimension (a negative correlation with the spread) and high liquidity in the volume dimension (a positive correlation with depth). Finally, spread and depth resiliency are positively correlated (0.463). Taking the

[31] I collect these time series for all 30 stocks and for each variable stack the time series into one single vector. Correlations are computed on the basis of these stacked vectors.

results of the cross-sectional estimation into account, a positive correlation for daily estimates is not surprising.

In all, the correlation of the bid-ask spread, depth and resiliency are as expected. The association of resiliency with the bid-ask spread is slightly higher and the association with depth is slightly lower. All coefficients are not particularly high which indicates that resiliency does not simply duplicate the other liquidity dimensions. It contributes to the understanding of liquidity as a separate dimension, the time dimension of liquidity.

4.8 Conclusion

This chapter empirically investigates, for the first time, the main features of resiliency as a dimension of liquidity in an electronic limit order market. Resiliency addresses the following question: when trades, especially those resulting from relatively large and uninformative orders, consume liquidity by eroding the limit order book, how fast is the spread reduced again and how fast is the depth of the limit order book replenished through the competitive actions of market traders. Resiliency results from the interaction of liquidity flowing into the market and liquidity being taken out. The inflow comes from the submission of new limit orders, while the outflow results either from the cancellation of limit orders or the execution of limit orders against newly submitted market orders. Together, inflow and outflow determine the evolution of the price-quantity schedule.

Specifically, I first address how we can formally define and measure resiliency.

I accordingly set up a mean reversion model of liquidity to capture the dynamics of the price and quantity schedule over time, and examine the relation between current and past liquidity flows. I examine ask-side and bid-side resiliency separately, and also analyze a range of different data frequencies. Secondly, I analyze the microstructural time-series factors that affect resiliency. In particular, these factors include informed trading, the order arrival rate, the patience of traders, the time of the day, trading activity and unexpected volatility. Thirdly, I analyze the variation in resiliency across the cross section of stocks. And finally, I examine the relationship between resiliency and the other two liquidity dimensions, the spread and depth.

I find strong evidence that the liquidity dynamics of the order book follows a stable replenishment process: the resiliency for each stock over a five-minute horizon is consistently high and stable across different frequencies. Empty order books are refilled quickly, while full books attract less new liquidity. The overall order flow reflects this resiliency, but, far more strongly, the resiliency can be seen in the behaviour of liquidity suppliers around the first few ticks of the book. Both for the depth and for the spread, resiliency is very strong around the best price and gets steadily weaker the further we move away from the best price in the book. Clearly, trades that are executed against the book take away liquidity at the first few ticks, and traders who actively monitor the book jump in straight away to exploit these profit opportunities in the book.

I also find that, in its time series behavior, resiliency is dependent in a stable and robust manner on microstructural determinants. In information-intense periods, resiliency is lower than in liquidity-intense periods. This seems plausible as liquidity suppliers fear that they will lose to their better-informed counterparties when trades are particularly informative. As predicted by Foucault, Kadan and Kandel (2005), the resiliency of the spread is associated positively with the proportion of patient traders in the population and negatively with the order arrival rate and the time of the day. Like the arrival rate, trading volume also has a negative impact on resiliency. The effect of unexpected volatility is asymmetric: resiliency on the bid side increases, while it falls on the ask side. Interestingly, these effects on the resiliency of the spread are the opposite for depth resiliency. The most probable explanation is that, in the time series, spread and depth resiliency are asynchronous. When spread resiliency is high, the spread gets improved rapidly which, in turn, erodes depth. I would therefore expect spread resiliency to come before depth resiliency and, hence, an opposite effects of time series factors on these two dimensions.

My results also show that, ceteris paribus, investors will choose more resilient stocks over less resilient stocks. Resiliency has a high association with large market-capitalized and high-beta stocks. On the other hand, stocks which offer high profits to informed traders tend to be less liquid. I also find that resiliency is not significantly correlated with either spread or depth. This reinforces, ex post, the importance of resiliency as an independent dimension of liquidity. It cannot be seen or assumed as a replica of the price or quantity dimension. This time dimension of liquidity provides significant new information.

Chapter 5

Commonality Across Limit Order Books

This chapter investigates the commonality of liquidity in the limit order book for an electronic limit order book market. I use order book data from the electronic trading facility Xetra of the German Stock Exchange. I construct liquidity measures of a stock's limit order book and estimate the common movement of these liquidity measures. I find strong evidence that there is commonality in liquidity. It is much stronger for liquidity measures of the limit order book than for simple proxies like spreads. Secondly, it shows time variation both on an intradaily basis and over longer horizons. The common movement of liquidity implies that trading in limit order book markets is subject to systematic liquidity risk.

5.1 Introduction

Liquidity risk is a major concern to investors, because it implies that they might have to trade when markets are especially illiquid, and trading in illiquid markets is very costly. If the liquidity of different stocks also moves together, liquidity risk will be an even greater concern to investors:[32] common liquidity movements imply that liquidity risk is market-wide and, apart from idiosyncratic shocks, cannot be diversified. Investors will have to bear the systematic component of liquidity risk and will therefore ask for compensation.

Empirical asset pricing papers such as Amihud (2002), Pástor and Stambaugh (2003) or Gibson and Mougeot (2004) show that investors receive a notable compensation for bearing liquidity risk. In contrast, empirical microstructure papers such as Chordia, Roll and Subrahmanyam (2000), Hasbrouck and Seppi (2001), Halka and Huberman (2001) and Brockman and Chung (2002) suggest that liquidity risk is almost entirely firm- specific. It implies that liquidity risk is diversifiable and should not be priced.

One possible explanation for the weak evidence of commonality in the microstructure literature is that earlier studies used poor proxies for liquidity. For example, Chordia, Roll and Subrahmanyam (2000), Hasbrouck and Seppi (2001), Halka and Huberman (2001) and Brockman and Chung (2002) all measure liquidity by looking at quotes and quantities at best prices. However, if investors want to trade large positions, their orders will walk up the book and therefore they will not only care about liquidity at best prices, but also about liquidity beyond best prices.

---

[32] The theoretical literature develops several kinds of mechanisms through which the liquidity supply of different stocks is linked. In these models, the correlation of liquidity preferences, intermediary behavior or informational shocks create contagion effects in the liquidity of different stocks. See Allen and Gale (2000), Kyle and Xiong (2001), Gromb and Vayanos (2002), Fernando (2003), Watanabe (2003) and Brunnermeier and Pedersen (2005).

A second explanation is that best quotes are particularly noisy and therefore not well-suited for the study of commonality. Since liquidity suppliers compete fiercely for new price priority, the bid-ask spread and depth at best prices are subject to strong idiosyncratic variation. This hypothesis is consistent with Domowitz, Hansch and Wang (2005) who show that order type correlation has far more explanatory power for liquidity commonality inside the limit order book than at the best prices.

A third explanation is time variation. Chordia, Roll and Subrahmanyam (2001) show that liquidity is time-varying and that it is particularly low in falling markets. If commonality exhibits a similar time variation, it might be low on average but higher in falling markets. The empirical evidence on this issue is ambiguous. For the US market, Coughenour and Saad (2004) find that in falling markets specialist behavior tends to be more strongly correlated across the stocks that they manage than in rising markets, while Domowitz, Hansch and Wang (2005) observe no systematic differences for the Australian market.

In this chapter I focus on (i) the level of commonality in liquidity beyond best prices and (ii) the time variation of commonality in a pure limit order market. The basic methodology follows the market model used in Chordia, Roll and Subrahmanyam (2000) and extends it to liquidity measures beyond the bid-ask spread and depth. Furthermore, I examine the link between commonality in the limit order book and movements of the market return.

This study is most closely related to the work of Bauer (2004) and Domowitz, Hansch and Wang (2005). Bauer (2004) performs a Principal Component Analysis (PCA) for liquidity measures of the limit order book and relates commonality to underlying financial variables. Domowitz, Hansch and Wang (2005) investigate the impact of order-type and order-flow correlations on return and liquidity commonality. I differ from these papers in three ways: firstly, I focus on how commonality in liquidity depends on how deep we look into the limit order book (as opposed to its relation with underlying financial variables or order-type and order-flow commonality). Secondly, from a methodological point of view, I apply the Chordia, Roll and Subrahmanyam (2000) market model instead of using a correlation or common factor approach. This approach has the advantage that it allows me to control for external factors and that the common factor can directly be interpreted as market liquidity. Thirdly, I use data from the German stock market, one of the world's largest and most important markets, instead of the Swiss and Australian market.[33]

Based on the Xetra limit order book, I measure liquidity both at best prices (using the bid-ask spread and its depth) and beyond best prices (using depth deeper in the limit order book and the slope of the price-quantity schedule). My study yields the following main results: (i) I find evidence of significant common variation in liquidity throughout the order book. The more the liquidity

[33] The growing availability of limit order book data has produced many further studies on such markets. With regard to commonality, Chordia, Sarkar and Subrahmanyam (2005) study common liquidity movements across asset classes. Benston, Irvine and Kandel (2000) and Cao, Hansch and Wang (2004) study the information content of the limit order book. Coppejans, Domowitz, Madhavan (2003) focus on dynamic issues of the limit order book. Beltran, Giot and Grammig (2005) relate commonalities across price-volume pairs of the limit order book to underlying microstructural factors.

measures are extended beyond best prices, the stronger commonality is: at best prices, common variation in depth only accounts for roughly 2% of all liquidity variation, while it increases to a maximum level of about 20% deeper in the book. (ii) Commonality exhibits strong time variation associated both with the time of the day and the movement of the market. Most notably, liquidity commonality increases strongly with the absolute value of negative returns. This implies that diversifying liquidity risk becomes more difficult in falling markets when diversification is particularly important.

The remainder of the chapter is organized as follows: Section 5.2 investigates the commonality of liquidity at the best prices in the book. In Section 5.3 I construct liquidity measures that incorporate the liquidity in the order book. I use a market model and principal component analysis to explore the commonality in order book liquidity. Section 5.4 investigates the influence of stock industry, time of day and market momentum. Section 5.5 concludes.

5.2 Commonality at Best Prices

Previous studies have implicitly assumed that best limit prices alone are sufficient to capture the liquidity of an asset. To relate my results to the literature, I analyze the commonality of liquidity for the bid-ask spread and for the depth of the order book at the best bid and ask prices.

In Chapter 3 I showed that, consistent with the literature, liquidity exhibits strong time-of-day effects. Like Hasbrouck and Seppi (2001) I therefore standardize liquidity by time-specific means and standard deviations. Thus I focus on the unexpected component of liquidity variation which corresponds to the liquidity risk that investors bear. I illustrate the standardization procedure for the bid-ask spread: let p_A denote the best ask price, p_B the best bid price and MQ the midquote. I compute bid-ask spreads relative to the midquote, $RS = (p_A - p_B)/MQ$. Then let h denote the time of the day and d a specific trading day (with $d = 1, ..., 64$). For every stock i, I take subsamples that include observations at one time of the day but over all days. From these subsamples I calculate the time-specific mean of the spread, $\mu_{i,h} = \Sigma_{d=1}^{n} RS_{i,h,d}/64$, and its standard deviation, $\sigma_{i,h} = \sqrt{\Sigma_{d=1}^{64}(RS_{i,h,d} - \mu_{i,h})^2/64}$. For all stocks I then demean and standardize each observation of the spread according to its time of the day. Let RS^* denote spreads adjusted for trends:

$$RS^*_{i,h,d} = \frac{RS_{i,h,d} - \mu_{i,h}}{\sigma_{i,h}} \tag{5.1}$$

The same procedure is applied to depth at the best prices. Let n_A be the number of shares quoted at the best ask and n_B at the best bid. I compute depth as $DEP_{A,B} = n_A \cdot p_A + n_B \cdot p_B$ and then correct it for time-of-day trends to obtain $DEP^*_{A,B}$. The detrended time series of the spread and depth are the input for the following analysis of commonality at best prices.

The standard econometric approach to commonality is the market model in Chordia, Roll and Subrahmanyam (2000) which is estimated by time series regressions. The market model relates the liquidity of a single stock to the liquidity

Table 5.1: Market Model for Spreads and Depth – Individual Stocks

	Spreads			Depth		
Stock	β	t_β	R^2	β	t_β	R^2
Adidas-Salomon	0.6990	6.7864	0.0584	0.0753	1.1615	0.0109
Allianz	0.3734	6.0024	0.0579	0.2386	1.7794	0.0071
Altana	0.9078	6.5176	0.0512	0.0953	1.8156	0.0125
BASF	0.4568	6.8301	0.0720	0.4295	3.5770	0.0227
BMW	0.7256	7.8778	0.0822	0.4742	4.8448	0.0290
Bayer	0.5154	6.0242	0.0569	0.2042	2.0234	0.0275
Commerzbank	1.0406	9.9498	0.1098	0.2087	1.9895	0.0141
Continental	0.8969	6.2345	0.0497	0.1703	3.0950	0.0282
DaimlerChrysler	0.6308	8.6423	0.0885	0.3181	2.3183	0.0120
Deutsche Bank	0.3601	6.3999	0.0487	0.3441	1.9998	0.0299
Deutsche Boerse	0.6747	5.7529	0.0455	0.1925	2.7269	0.0201
Deutsche Post	0.7599	6.1584	0.0610	0.1286	1.3530	0.0104
Deutsche Telekom	0.3184	6.8685	0.0500	0.7755	1.1039	0.0420
E.ON	0.9983	8.7036	0.1132	0.2041	4.1170	0.0344
Fresenius	0.7957	5.2823	0.0357	0.2129	4.2120	0.0318
Henkel	0.4142	6.3119	0.0731	0.6281	5.1925	0.0409
HypoVereinsbank	0.1291	2.0315	0.0123	0.0179	0.1488	0.0087
Infineon Technologies	0.5589	6.5592	0.0468	1.0661	4.2320	0.0456
Linde	0.7843	6.4774	0.0574	0.2169	3.6965	0.0206
Lufthansa	0.9678	8.0727	0.0739	0.2359	2.6843	0.0195
MAN	1.1994	7.9897	0.0786	0.0931	1.5158	0.0197
Metro	0.8874	6.4716	0.0535	0.1176	1.6893	0.0193
Muenchener Rueck	0.5454	7.8595	0.0955	0.1777	1.0794	0.0055
RWE	0.5170	5.8814	0.0510	0.2081	1.9924	0.0066
SAP	0.3906	5.4891	0.0373	0.3358	3.0865	0.0239
Schering	0.5941	6.0253	0.0462	0.4035	5.4434	0.0552
Siemens	0.4222	6.5354	0.0715	0.1327	1.1256	0.0168
ThyssenKrupp	0.2495	3.8854	0.0273	0.0532	0.4538	0.0056
TUI	1.0952	6.4770	0.0465	0.2431	3.6895	0.0250
Volkswagen	0.5061	6.4617	0.0590	0.1066	1.2058	0.0148
Averages	0.6472	6.5520	0.0604	0.2703	2.5118	0.0220

Table 5.1 reports the parameter estimates of the liquidity market model for spreads and depths for each stock (Equation 5.2). It gives the coefficients of market liquidity, their corresponding t-values and the R^2 values of the regression. The last row gives the averages across stocks.

of the market in the same way as the CAPM does for returns. It employs a regression framework to measure the sensitivity of stock i's liquidity, L_i, to market liquidity, L_M.[34] Market liquidity is computed as the average liquidity across all stocks, $\Sigma_{i=1}^{29} L_{i,t}/29$, where t is a time subscript.[35] The estimation includes lead and lag market liquidity ($L_{M,t+1}$ and $L_{M,t-1}$), contemporaneous, lead and lag market returns ($r_{M,t}$, $r_{M,t+1}$ and $r_{M,t-1}$) as well as individual stock return volatility $VOL_{i,t}$ (proxied by the squared return) as additional regressors. With ε as an error term, I obtain the following specification:

$$\begin{aligned} L_{i,t} = \alpha + \beta_i^1 L_{M,t} + \beta_i^2 L_{M,t+1} + \beta_i^3 L_{M,t-1} + \delta_i^1 r_{M,t} + \delta_i^2 r_{M,t+1} \\ + \delta_i^3 r_{M,t-1} + \eta_i VOL_{i,t} + \varepsilon_{i,t} \end{aligned} \tag{5.2}$$

Tables 5.1 and 5.2 summarize the estimation results of Equation (5.2) for liquidity at best limit prices. Table 5.1 gives a detailed picture of all 30 β^1 coefficients. β^1 measures how strongly individual stock liquidity is determined by the market level of liquidity. For the relative spread the level of the coefficients varies between 0.1291 and 1.1994 with an average of 0.6472, while the t-statistic reflects significance for all 30 stocks. The R^2 values vary between 1% and 11% with an average of 6.04%. For depth at the best bid and ask, β^1 varies between 0.02 and 1.06 with an average of 0.27 across stocks. The parameter estimates are significant in 70% of all regressions, while the average R^2 value is about 2%. The

[34] The Chordia, Roll and Subrahmanyam (2000) model is in first differences. I follow Hasbrouck and Seppi (2001) and use standardized variables instead.

[35] For each stock i's regression, stock i is dropped in the calculation of market liquidity, because it would lead to additional correlation otherwise. Further robustness checks show that the results are unchanged for value-weighted market liquidity instead of the arithmetic average.

Table 5.2: Market Model for Spreads and Depth at the Best Limit Prices

		Spreads: Averages		Depth: Averages	
		Parameters	t-statistics	Parameters	t-statistics
Intercept	α	0.1309	0.2765	0.1186	3.6111
Market Liquidity	β^1	0.6472	6.5520	0.2703	2.5118
	β^2	0.0383	0.5094	0.1412	0.6699
	β^3	0.0508	0.6317	0.1268	1.0269
Market Return	δ^1	-0.0132	-0.2541	0.0053	0.0508
	δ^2	0.0064	0.0176	0.0797	0.4039
	δ^3	-0.0144	-0.0105	-0.1220	-0.2933
Return volatility	η	0.1115	0.6725	0.0049	0.0899
Adjusted R^2		0.0604		0.0220	

Table 5.2 reports the parameter estimates of the liquidity market model (Equation 5.2) for spreads and depths. It gives the mean parameter estimates across all 30 stocks and the corresponding average t-values for spreads and depth. The last row reports the average adjusted R^2 values of the regressions.

results show that the spread market model of liquidity is stronger than for depth. Furthermore, the adjusted R^2 measure of the regressions can be interpreted as a measure of commonality, since it explains the percentage of individual liquidity variation that is explained by market liquidity.[36] Thus, commonality in spreads is about 6% while it is lower at 2% for depth. Table 5.2 shows the average parameter estimates and t-values for the additional regressors. We see that their averages are clearly insignificant; unlike market liquidity they do not explain any movement in individual stock liquidity.[37]

[36] R^2 captures the effects of all explanatory variables rather than market liquidity alone. However, leaving out the insignificant regressors does not lower the adjusted R^2 values very much. Therefore its magnitude is nearly completely driven by the market liquidity coefficient. This procedure allows direct comparison with the R^2 values in Chordia, Roll and Subrahmanyam (2000).

[37] I average the estimated parameters and their corresponding t-values across all stocks as in

The market model results that I obtain are very much along the lines of Chordia, Roll and Subrahmanyam (2000) and Brockman and Chung (2002). In these studies, only the market liquidity parameter is significant, while the additional regressors are not. They report adjusted R^2 values of anywhere between 1.70% and 2.78% for spreads and 1.00% and 2.08% for depth. I obtain a value of 2% for the depth and 6% for spreads while all control variables are insignificant. Evidently, the spread market model is slightly stronger for my data, yet 6% is not that convincingly high, either.

5.3 Commonality Beyond Best Prices

In this section I extend commonality to the order book. First, I construct measures of order book liquidity and secondly compute the extent of commonality they show. Finally, I employ principal components analysis (PCA) to check that the results are stable with respect to the methodology. I compare the level of common liquidity movement to the levels evidenced by the bid-ask spread and the depth at the best bid and ask price.

5.3.1 Construction of Liquidity Measures

In limit order book markets, all orders are executed against the limit orders in the order book. If an order is very large, this implies that it will hit unexecuted limit orders which have different price limits. The larger the order, the more price limits will be hit and the further a market order walks up the limit order

Chordia, Roll and Subrahmanyam (2000) and Brockman and Chung (2002). A cross-sectional t-statistic implicitly assumes that the estimation errors in β_i are independent across stocks.

book. Furthermore, if many transactions take place before any limit order trader has time to submit new limit orders, limit orders that were deep in the order book suddenly become relevant. Evidently, the spread alone is not sufficient to characterize the liquidity of a limit order book market; I need a measure of the order book to assess a stock's liquidity.

The bid-ask spread gives the price discrepancy between the best prices in the book, while depth at the spread is the corresponding volume in the order book. A natural extension is to move away from the best prices and to consider prices p and volume x deeper in the order book. Because of asymmetries of the bid and ask side I construct separate measures for each order book side. In a first step I compute the price difference between all limit order prices in the book and the midquote at that point in time. In a second step I transform these price differences into price impacts relative to the midquote and link them to the cumulative volume in the order book, $PI_i^A(x) = (p_i(x) - MQ_i)/MQ_i$ and $PI_i^B(x) = (MQ_i - p_i(x))/MQ_i$. Finally, I compute the volumes that correspond to price impacts of 0.5%, 1%, 1.5% and 2%. I choose 2% as the cut-off value for this extended depth measure because in my data set market orders that are executed against the limit order book seldom incur higher price impacts.[38]

An advantage of this extended depth measure is that it is non-parametric and is not based on any restrictive assumptions. A disadvantage, however, is that it

[38] The measure is similar to the cost of round trip in Benston, Irvine and Kandel (2000), the XLM measure in Gomber, Schweickert and Theissen (2004) or hypothetical price impacts as in Kumar (2003). Unlike cost-of-round-trip measures, my measure captures asymmetries of the bid side and the ask side.

only characterizes one point on the price-quantity schedule at a time. I therefore introduce the slope of the price-quantity schedule as a second measure that summarizes all price-volume combinations in the limit order book simultaneously. Knowledge of the slope of the price-quantity schedule enables a trader to compute price impacts for any order size. With regard to the specification, the empirical literature comes to very mixed conclusions: Biais, Hillion and Spatt (1995) document linearity, while Coppejans, Domowitz, Madhavan (2003) or Cao, Hansch and Wang (2004) find evidence of some non-linearities.

Let l denote the pairs of price-volume combinations of an order book, let t denote the individual points in time for which I have order book snapshots and let x denote the volume in the order book. If ε is the error term, I obtain the following equation for a linear model price-quantity schedule:

$$PI_{i,t,l} = \lambda_{i,t} \cdot x_{i,t,l} + \varepsilon_{i,t,l} \tag{5.3}$$

As with the extended depth measures, I cut off the price-quantity schedule for price impacts higher than 2%. The indices indicate that I estimate the model in each point in time and for each stock. Furthermore, I do the estimation separately for the ask side and bid side. I also estimate Equation 5.3 with an additional quadratic term, $\rho_{i,t}x_{i,t,l}^2$. Negative estimates of $\rho_{i,t}$ imply a concave order book function and positive estimates imply a convex relationship. I find neither a significantly high number of positive nor negative estimates and therefore I choose the linear model. I obtain a time series of $\lambda_{i,t}^A$ and $\lambda_{i,t}^B$ for each stock. The model does not include a constant to ensure that the estimated graph starts at the origin. Subsequently, the state of the order book is summarized by one parameter. Since the price impact function is upward-sloping by construction, it is not surprising that the fits turn out to be very good and that the parameters estimates are highly significant.

Table A.2 in the appendix gives some descriptive statistics of my proxies. The average volume associated with 2% price impacts is 4,219,230 Euros on the ask side and 4,032,470 Euros on the bid side, yet there is considerable variation in these figures. The average slope estimate is 0.8992 on the ask side and 0.9557 on the bid side. Again there is considerable cross-sectional variation across stocks: As the model of the order book slope has no constant, the slope is a direct measure of the level of liquidity: for example, in the case of Fresenius an investor who wishes to buy a position of 500,000 Euros will incur a price impact of 1.1258%. In comparison, the half-spread is 0.0655%. Evidently, spread measures are bad proxies for large volumes.

5.3.2 Market Model Results

In this section I investigate market-wide liquidity movements of the entire limit order book. I use the same methodology as before, yet this time I substitute $L_{i,t}$ and $L_{M,t}$ in Equation 5.2 by my new measures of order book liquidity. Again I follow the standardization procedure of Equation 5.1 to eliminate trends from the data.

Table 5.3: Market Model for the Extended Depth Measure – Individual Stocks

	Ask Side			Bid Side		
Stock	β	t_β	R^2	β	t_β	R^2
Adidas-Salomon	0.7766	3.6945	0.0579	0.9571	5.0424	0.1020
Allianz	1.0389	4.9800	0.1042	0.8625	5.0533	0.2375
Altana	0.6372	3.0701	0.0918	0.8732	4.4712	0.0410
BASF	0.8268	4.1595	0.1397	1.0128	5.5817	0.1301
BMW	0.5456	2.7432	0.1508	0.6165	3.2135	0.0588
Bayer	0.8077	5.2372	0.3513	0.5407	2.7524	0.0205
Commerzbank	0.7146	3.8567	0.1980	0.6950	3.8881	0.1745
Continental	0.9604	4.7992	0.1835	0.5895	3.0350	0.0426
DaimlerChrysler	0.8281	5.1474	0.2641	1.0605	5.9898	0.2049
Deutsche Bank	0.7954	4.0889	0.1865	0.5277	2.8419	0.0753
Deutsche Boerse	0.6987	3.2872	0.0702	0.7246	3.7491	0.0341
Deutsche Post	0.8217	4.7963	0.3750	0.6051	3.3161	0.0704
Deutsche Telekom	0.5922	2.9398	0.1522	0.5497	3.0045	0.1287
E.ON	0.2957	1.4514	0.0636	0.3482	1.8335	0.0388
Fresenius	0.9057	5.0962	0.1631	0.9372	4.8645	0.0396
Henkel	0.8850	4.1493	0.0705	0.8781	4.6440	0.0517
HypoVereinsbank	0.0117	0.0557	0.0190	0.1885	0.9859	0.0017
Infineon Technologies	0.3439	1.6709	0.0830	0.1019	0.5343	0.0159
Linde	1.0261	4.9520	0.1348	0.7600	3.8896	0.0393
Lufthansa	0.5429	2.7324	0.1600	0.5930	3.2281	0.0770
MAN	0.7349	3.4949	0.0731	0.9213	4.7641	0.0706
Metro	0.9738	4.9160	0.1847	0.6173	3.1637	0.0164
Muenchener Rueck	0.8404	4.6200	0.3087	0.8559	4.5013	0.0777
RWE	0.6269	2.9312	0.0564	0.6129	3.1520	0.0496
SAP	0.9122	4.6560	0.2048	0.7529	3.9907	0.0819
Schering	0.5815	2.8550	0.1180	0.4207	2.1680	0.0328
Siemens	0.8535	4.3360	0.2043	1.0537	5.8729	0.1852
ThyssenKrupp	0.6324	3.5943	0.3301	0.4752	2.3938	0.0097
TUI	0.6339	3.1542	0.0716	0.6700	3.7105	0.1365
Volkswagen	0.8125	4.0291	0.1541	0.9084	5.3488	0.2565
Average	0.7200	3.7172	0.1609	0.6811	3.6532	0.0827

Table 5.3 reports the parameter estimates of the liquidity market model for the depth of the order book at 2% price impact for each stock. It gives the coefficients of market liquidity, their corresponding t-values and the R^2 values of the regression for the ask and bid side. The last row gives the averages across stocks.

Table 5.4: Market Model for Depth at 2.0 % Price Impact

		Ask Side: Averages		Bid Side: Averages	
		Parameters	t-statistics	Parameters	t-statistics
Intercept	α	-0.0084	-0.3581	-0.0034	-0.1151
Market Liquidity	β^1	0.7219	3.7165	0.6903	3.6995
	β^2	0.0634	0.4276	-0.0048	0.0121
	β^3	0.0725	0.4719	0.0009	0.0515
Market Return	δ^1	0.0794	0.5497	-0.0230	-0.1714
	δ^2	0.0323	0.2202	-0.0729	-0.5156
	δ^3	-0.0240	-0.1726	-0.0177	-0.1301
Return volatility	η	0.0492	1.1742	0.0760	1.5074
Adjusted R^2		0.1609		0.0827	

Table 5.4 reports the parameter estimates of the liquidity market model (Equation 5.2) for the extended depth measure at a price impact of 2.0%. It gives the mean parameter estimates across all 30 stocks, the average t-values and the percentage of significant coefficients at the individual stock level (in %) for the ask side and the bid side of the order book. The last row reports the average R^2 values of the regressions.

Tables 5.3 and 5.4 give the results for order book depth at 2% price impact. In Table 5.3 I list all individual β^1 coefficients with their corresponding t-statistics. Both on the ask side and on the bid side, β^1 is significant for 28 out 30 stocks (93.33%). Table 5.4 shows the average values. The beta coefficient for contemporaneous market liquidity is the only coefficient with a significant t-value; as before, the t-values of all other regressors indicate that they are not significant. Compared to the results of liquidity at the best limit price alone, however, the adjusted R^2 values are much higher. They climb to an average of 16.09% on the ask side and 8.27% on the bid side. In comparison, the corresponding R^2 value of depth at best prices was between 1.0% and 2.1%. In other words, commonality increases strongly if I consider the aggregate liquidity supplied to the limit order book at price impacts of 2%.

Next I proceed to the slope of the price-quantity schedule of the limit order book which is also estimated for price impacts up to 2%. Tables 5.5 and 5.6 present the results. Table 5.5 shows that the β^1 coefficient varies between 0.1139 and 1.0966 on the ask side of the limit order book and 0.2043 and 0.9949 on the bid side. 29 out of 30 stocks have significant coefficients on the ask side and 28 out of 30 on the bid side. On average, the market liquidity coefficients are 0.7302 on the ask side and 0.6645 on the bid side; both are highly significant. All other coefficients are close to zero in a range between -0.1063 and 0.0704 with small and insignificant t-values. Commonality on the ask side is 17.19% and on the bid side 10.90%. Evidently, the results for slope parameters of the price-quantity schedule are very similar to those for aggregate depth. The explanatory power of the slope model as opposed to the extended depth measure is minimally more powerful (17.19% versus 16.09% on the ask side and 10.90% versus 8.27% on the bid side). Compared to commonality at best prices, the extent of systematic liquidity risk is once again considerably higher. These results underline that the high extent of commonality for liquidity in the limit order book remains a robust result irrespective of the measure.

Obviously systematic movements in liquidity are quite different for the liquidity flow at best prices and the liquidity flow up until ticks further beyond best prices. As I have only focused on liquidity at 2% price impacts so far, I now

Table 5.5: Market Model for the Slope of the Price-Quantity Schedule – Individual Stocks

Stock	Ask Side β	Ask Side t_β	Ask Side R^2	Bid Side β	Bid Side t_β	Bid Side R^2
Adidas-Salomon	0.6706	4.2422	0.0812	0.8405	5.2077	0.1162
Allianz	1.0966	7.1134	0.1493	0.6913	4.8021	0.2748
Altana	0.8954	5.6654	0.0926	0.8583	5.1036	0.0495
BASF	0.8401	5.4983	0.1448	0.9949	6.4469	0.1969
BMW	0.4935	3.3006	0.1516	0.7096	4.3513	0.0897
Bayer	0.8649	6.1124	0.2614	0.6234	3.8197	0.0813
Commerzbank	0.9036	6.3157	0.2441	0.6266	4.4515	0.2953
Continental	0.8532	5.5842	0.1459	0.6429	3.7752	0.0211
DaimlerChrysler	0.6369	4.4276	0.2296	0.6894	4.4588	0.1678
Deutsche Bank	0.5673	3.7430	0.1380	0.5583	3.4095	0.0675
Deutsche Boerse	0.7413	4.7605	0.1053	0.7071	4.1456	0.0201
Deutsche Post	0.7483	5.7654	0.3705	0.5078	3.1506	0.0956
Deutsche Telekom	0.5207	3.6392	0.2253	0.4125	2.5542	0.0775
E.ON	0.2965	1.8984	0.0758	0.2043	1.2425	0.0384
Fresenius	0.7951	5.4021	0.1983	0.6289	3.7117	0.0283
Henkel	0.8211	5.1907	0.0945	0.7626	4.5393	0.0468
HypoVereinsbank	0.1139	0.7140	0.0173	0.3489	2.0703	0.0092
Infineon Technologies	0.3975	2.7986	0.2269	0.2417	1.4622	0.0308
Linde	0.7843	4.9910	0.0978	0.8434	4.9817	0.0434
Lufthansa	0.6055	3.9815	0.1325	0.5983	3.6188	0.0557
MAN	0.8529	5.4123	0.1103	0.8673	5.4119	0.1434
Metro	0.9083	6.0797	0.1832	0.6825	4.0259	0.0237
Muenchener Rueck	0.9139	6.8574	0.3416	0.9219	6.1417	0.2288
RWE	0.5247	3.3518	0.0853	0.4731	2.8092	0.0171
SAP	0.7750	5.1553	0.1674	0.7169	4.4542	0.1158
Schering	0.8488	5.7700	0.2030	0.7115	4.3647	0.0874
Siemens	0.8625	5.9839	0.2408	0.9600	6.6583	0.2860
ThyssenKrupp	0.8374	6.1046	0.3025	0.6241	3.6965	0.0342
TUI	0.8308	5.4249	0.1379	0.5796	3.7072	0.1507
Volkswagen	0.9060	6.1559	0.2029	0.9083	6.7677	0.3775
Average	0.7302	4.9147	0.1719	0.6645	4.1780	0.1090

Table 5.5 reports the parameter estimates of the liquidity market model for the slope of the price-quantity schedule for each stock. It gives the coefficients of market liquidity, their corresponding t-values and the R^2 values of the regression for the ask and bid side. The last row gives the averages across stocks.

Table 5.6: Market Model for the Slope of the Price Impact Function

		Ask Side: Averages		Bid Side: Averages	
		Parameters	t-statistics	Parameters	t-statistics
Intercept	α	0.0017	0.0600	0.0010	0.0335
Market Liquidity	β^1	0.7302	4.9147	0.6645	4.1780
	β^2	0.0600	0.4901	0.0568	0.4831
	β^3	0.0854	0.6713	0.0704	0.5787
Market Return	δ^1	-0.1063	-0.7185	0.0592	0.4333
	δ^2	-0.0359	-0.2520	0.0558	0.4003
	δ^3	0.0332	0.2250	-0.0109	-0.0677
Return volatility	η	-0.0041	-0.2888	-0.0003	-0.1862
Adjusted R^2		0.1719		0.1090	

Table 5.6 reports the parameter estimates of the liquidity market model (Equation 5.2) for the slope parameter of the price impact function. It gives the mean parameter estimates across all 30 stocks and the average t-values for both the bid side and the ask side of the order book. The last row reports the average adjusted R^2 values of the regressions.

turn my attention to the region between the best ask and ask prices that are 2% above the midquote and 2% below the best bid, respectively. Table 5.7 shows how commonality changes for different cut-off points of aggregate order book depth. On the ask side, β^1 starts out at 0.53 for 0.5% impacts and increases continually to 0.72 for order book depth at a 2% price impact. The corresponding average t-values are very clearly above the critical value of 1.65. Turning to the level of commonality, I observe that the adjusted R^2 value on the ask side is 9% for a price impact of 0.5%. It increases continually up to 16%. The pattern holds for the bid side of the order book as well where commonality increases to a value of 8%. In Figure 5.1 I illustrate the results graphically. In contrast to the previous liquidity measures, I have not cut liquidity off at 2% to show how commonality increases even further in the order book. It approaches levels beyond 20% on the ask side and beyond 10% on the bid side. While this kind of depth might not be required by traders in large-cap stocks, it might well get consumed in smaller and less liquid stocks (see Keim and Madhavan (1997)). The results clearly indicate that mismeasurement of liquidity is one reason for the underestimation of commonality.

5.3.3 Principal Components Results

In the following section I approach commonality from a different angle. While Chordia, Roll and Subrahmanyam (2000) directly assume that the market average of liquidity explains individual stock liquidity, other studies such as Hasbrouck and Seppi (2001) or Hansch (2001) use principal component analysis (PCA), a more statistical approach. Instead of imposing any pre-specified restrictions on the common liquidity factor, they use PCA to extract the factor with the highest explanatory power for individual liquidity variation. I compare the results of these two methodologies for my data and conclude to what extent they influence the findings.[39]

The main input of the PCA is the correlation matrix of my liquidity measures. Again I standardise all measures first as in Equation 5.1. PCA then extracts the

[39] Beltran, Giot and Grammig (2005) have also used principal components analysis in the context of limit order books. They apply PCA to price-quantity pairs to examine whether variation in such pairs can be attributed to one or more underlying factors. While they link variation within the order book of single stocks to microstructural factors, I focus on the variation of liquidity across stocks. Furthermore, I am less interested in the identification of such factors. Rather, I use PCA as a means to estimate the extent of covariation across a range of assets.

Table 5.7: Market Model for Increasing Depth of the Limit Order Book

	Ask Side			Bid Side		
Price impact	β_1	t-value	R^2	β_1	t-value	R^2
0.5%	0.53	3.19	0.09	0.46	2.74	0.06
1.0%	0.53	2.50	0.12	0.49	2.42	0.07
1.5%	0.67	3.33	0.14	0.62	3.22	0.08
2.0%	0.72	3.72	0.16	0.69	3.70	0.08

Table 5.7 reports parameter estimates of the liquidity market model (Equation 5.2) for increasing depth. The first column lists the price impact up to which depth is aggregated. The further columns give average parameter estimates, corresponding t-values and average R^2 values for the ask side and the bid side of the order book.

linear combination of individual liquidity measures with the highest explanatory power for the variability in the data. This linear combination – called the first principal component – is effectively a weighting vector of individual liquidity and is given by the first eigenvector of the correlation matrix. Its explanatory power is given by the corresponding eigenvalue. From the weighting vector I compute the value of the principal component for each point in time. I then regress this time series onto the time series of individual stock liquidity measures for each stock:

$$L_{i,t} = \xi_i + \psi_i \; PC_{i,t} + \varepsilon_{i,t}, \tag{5.4}$$

In the equation above, $PC_{i,t}$ denotes the realization of the first principal component in t. ξ_i and ψ_i are parameters and $\varepsilon_{i,t}$ is an error term. To generate a test statistic, I bootstrap new time series from the regression residuals $\varepsilon_{i,t}$ of all

Figure 5.1: Commonality for Increasing Depth of the Limit Order Book

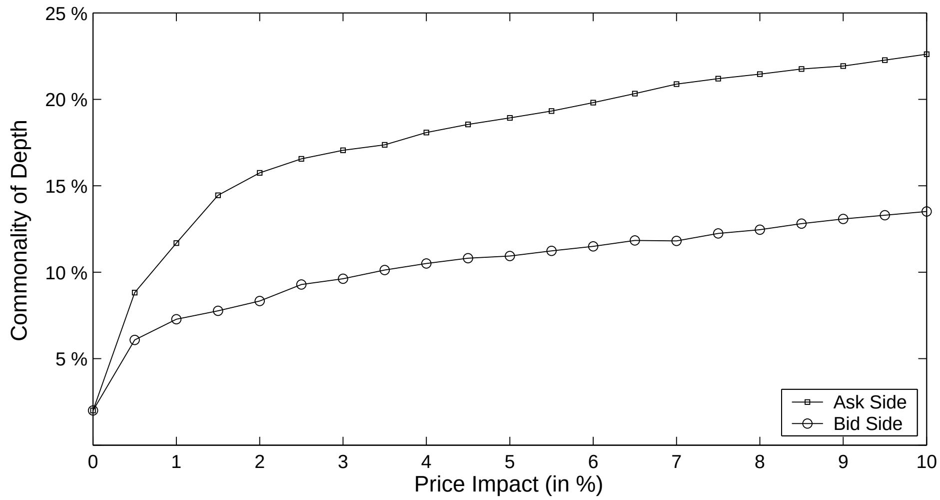

Figure 5.1 shows how the extent of commonality increases as more liquidity in the order book is considered. The x-axis gives the price impact up to which depth is aggregated and the y-axis gives the amount of market-wide movement. Commonality is measured as the R^2 value of the market model of liquidity.

stocks, compute the correlation matrix and perform PCA. I repeat this procedure 10,000 times until I obtain a smooth empirical distribution of the first eigenvalue and sample the 95%-quantile as our critical value.

Table 5.8 summarizes the PCA results for aggregate depth. On the ask side, the first principal component of aggregate depth at 0.5% price impact is 4.0. With a critical value of 2.1 it is clearly significant and accounts for 13.3% of overall variation in depth. If I successively increase aggregate depth up to 2%, the amount of common variation rises continually from 13.3% to 20.1%. All first principal components remain strongly significant.[40] On the bid side of the

[40] Although all additional regressors turned out to be insignificant in the previous market average approach, I eliminate their impact on our liquidity measures to doublecheck the significance of our results. In a first step, I regress the liquidity measures onto the same explanatory variables, then compute the correlation matrix of all 30 stocks from their residuals and finally

book, the pattern is identical. All first principal components for aggregate depth are significant. In level, commonality starts out at 10.8% and rises to 12.7% if I consider depth up to 2% price impacts. If I compare these results to the regression approach in the previous section, I see that they are very similar indeed. Firstly, commonality increases with order book depth. Secondly, commonality is stronger on the ask side of the book. In level, the PCA results are about 4% above the market index approach for ask side depth and 2% for the bid side. Somewhat higher PCA results are not surprising, since PCA is not restricted to a predetermined measure of market liquidity.

For the sake of completeness I report the PCA results for the spread and slope of the price-quantity schedule in Table A.3 in the appendix. Qualitatively they provide the same evidence: commonality around the spread is lower (7.74% for the bid-ask spread and 7.22% for depth at best prices), while commonality of order book liquidity is higher (24.36% for the ask side slope and 14.69% for the bid side slope). As with depth, PCA obviously reinforces the conclusion drawn from the Chordia, Roll and Subrahmanyam (2000) model.

5.4 Time Variation of Commonality

The previous section showed that the mismeasurement of the liquidity in the limit order book is one reason why the impact of commonality in liquidity has been underestimated in the past. A second reason is that commonality might be

repeat the PCA procedure. This leads to minimally lower levels of commonality, yet qualitatively identical results.

Table 5.8: PCA Results for the Extended Depth Measures

	Price Impacts of Depth			
PCA Output	0.5%	1.0%	1.5%	2.0%
Ask Side				
First eigenvalue	4.00	4.85	5.69	6.04
Critical value	2.13	2.53	2.84	2.92
Proportion of variation	13.32	16.17	18.97	20.13
Bid Side				
First eigenvalue	3.24	3.57	3.66	3.80
Critical value	2.01	2.19	2.48	2.54
Proportion of variation	10.80	11.92	12.20	12.66

Table 5.8 gives the results of PCA for the extended depth measure at different price impacts. In the first section, the table lists the first eigenvalue, its critical values at the 95% confidence level and the proportion of total variability explained by the first principal component (in %) for the ask side of the order book. The second section gives the same information for the bid side.

time-varying. In the following section I explore the time variation with regard to intradaily patterns and with regard to the momentum of the market.

5.4.1 Time of Day

It is widely recognized that liquidity exhibits strong time-of-day effects. In particular, liquidity is low in the morning at the opening of the market. This is illustrated very clearly in the L-shape over the bid-ask spread and the upside-down L-shape of depth. Furthermore, trading activity also declines at the end of the day. I now turn to the question whether commonality in liquidity exhibits similar time-of-day effects. The following analysis uses liquidity measures which are already free of trends, so remaining time-of-day effects only reflect trends in commonality, not in the level of liquidity.

Table 5.9: Impact of the Time of Day

	Ask						Bid					
	Open		Day		End		Open		Day		End	
Impact	β_1	R^2	β_1	R^2	β_1	R^2	β_1	R^2	β_1	R^2	β_1	R^2
0.5%	0.57	13	0.46	09	0.52	11	0.52	10	0.42	07	0.39	08
1.0%	0.51	15	0.41	12	0.54	14	0.46	13	0.37	08	0.54	09
1.5%	0.57	17	0.46	15	0.75	18	0.55	14	0.46	07	0.73	11
2.0%	0.56	19	0.47	16	0.82	19	0.60	15	0.51	07	0.81	12

Table 5.9 reports the average coefficients of market liquidity and the corresponding R^2 values of the regressions (in %) for the opening of the trading day ("open"), the midday trading period ("day") and the end of the trading day ("end"). The results are reported for the bid side and the ask side of the book.

To examine these effects I investigate commonality on an intradaily basis. In particular, I split our data into three subsamples: a morning sample (order book liquidity until 11 a.m.), a midday sample (from 11.30 a.m. until 3.30 p.m.) and an evening sample (from 4 p.m. to 5.30 p.m.). I then reestimate Equation 5.2 separately for all three subsamples.

Table 5.9 summarizes the results for the extended depth measure. Although time-of-day effects have been eliminated from individual liquidity levels, the systematic movement of liquidity across stocks shows clear intraday patterns. For example, the adjusted R^2 value for ask side depth at a price impact of 2% is 19% in the morning, drops to 16% in the course of the day and goes back up to 19% in the late trading period. On the bid side, the R^2 value is 15% in the morning, 7% during the trading day and 12% at the close of the day. There is a clear U-shape in the commonality of liquidity: commonality is higher in the morning, falls to lower levels during the day and rises again in the evening before the exchange closes. This pattern is visible for every depth measure from 0.5% to 2% price impacts. It is also robust with regard to the ask and bid side of the limit order book.[41]

5.4.2 Market Momentum

While time-of-day effects are an explanation of variation in commonality over a short time horizon, market momentum is a possible explanation for variation over a longer horizon. Numerous studies provide evidence that the correlation of stock returns is strongest in falling markets (see Conrad, Gultekin and Kaul (1991), Kroner and Ng (1998), Bekaert and Wu (2000), Longin and Solnik (2001) or Ang and Chen (2002)). In such environments, a flight to quality reduces the liquidity of equity markets in favor of safer investments. Therefore, in falling markets liquidity tends to get withdrawn from many stocks at the same time and that, in turn, induces a higher commonality.

Let w with $w = 1, 2, 3, ..., 53$ denote rolling ten-day intervals. For each such interval I calculate the portfolio return of our sample stocks, R_w, and estimate Equation 5.2. I take the R^2 of each interval as measure of commonality, C_w. The highest ten-day return is 1.48% and the lowest is -4.60%. In a first step, I present the results for these two subsamples in Table 5.10. If I compare the R^2 values for the upwards and downwards trending markets I see that there is very strong

[41] We obtain the same qualitative results if I use the slope of the price-quantity schedule as our measure of the liquidity in the limit order book. See Table A.4 in the appendix for the estimation results.

Table 5.10: Impact of Market Momemtum

	Ask				Bid			
	Up		Down		Up		Down	
Price impact	β^1	R^2	β^1	R^2	β^1	R^2	β^1	R^2
0.5%	0.51	0.10	0.63	0.15	0.31	0.06	0.53	0.15
1.0%	0.49	0.14	0.62	0.18	0.40	0.07	0.56	0.20
1.5%	0.61	0.14	0.74	0.21	0.55	0.09	0.72	0.18
2.0%	0.72	0.15	0.78	0.22	0.63	0.10	0.73	0.14

Table 5.10 reports the average parameter estimates of the market liquidity parameter and R^2 values for rising markets ("up") and falling markets ("down"). It lists the results for the ask side and the bid side of the order book and also differentiates with regard to the price impact up to which order book depth is aggregated.

evidence of a momentum effect: liquidity comoves far more strongly in falling markets than in rising markets. On the ask side, commonality is about one third higher in falling markets. At a price impact of 2% commonality is about 15% in the rising market and 22% in the falling market. On the bid side commonality in the rising market is 10% and in the falling market 14%.[42]

In a second step I relate the degree of commonality to the portfolio return in the same time window. For the depth measures from 0.5% price impact to 2% price impact all ask side and bid side correlation coefficients are below -0.4. Evidently, market momentum and commonality are negatively related. I add an error term ε and estimate the following regression:

$$C_w = \alpha + \beta \; R_w + \varepsilon_w \tag{5.5}$$

[42] The results for the slope of the price-quantity schedule are reported in Table A.5 in the appendix.

In the above equation, β measures the relation between commonality and market momentum. For the extended depth measure at 2% price impacts β is -1.125 with a Newey-West adjusted t-statistic of -2.110. On the bid side, β is -0.602 with a Newey-West corrected t-statistic of -3.853. Clearly, the coefficients on both sides of the book are significantly negative. The significance is just as strong for all other depth measures as well. This implies that commonality is stronger if momentum is negative and weaker if it is positive.[43]

5.5 Conclusion

In this chapter, I depart from the observation that market-wide liquidity apparently gets priced, yet there is only very weak evidence that stock liquidity movement does actually exhibit a market-wide component. One reason for this weak evidence might be the mismeasurement of order book liquidity. A further explanation is that commonality might be time-varying. I examine these hypotheses with the help of an extensive order book data set from the Frankfurt Stock Exchange (FSE). It enables us to measure common movements of the entire liquidity in the order book. I construct measures of order book liquidity and then estimate the sensitivity of firm liquidity to market liquidity. I alternatively use principal components analysis to doublecheck that our results are robust with respect to the methodology.

In all, I observe strong evidence of commonality in liquidity. In a reference

[43] The results are qualitatively virtually identical if I measure liquidity by means of the price-quantity schedule of the limit order book. See Table A.6 in the appendix for a summary of the results for both depth and the slope of the price-quantity schedule.

scenario, I analyze commonality for liquidity at the best price. Although I document significant commonality, it is fairly weak as in the literature (2% for depth and 6% for spreads). However, once I take the liquidity supply in the order book into account, commonality increases strongly. If depth measures at 2% price impacts are used, commonality rises up to 16%. A closer examination of the results also reveals that ask side commonality is stronger than on the bid side. Measures of the price-quantity schedule lead to similar results. Obviously, in a limit order market with a limit order book, the commonality of liquidity provision is drastically higher than the spread suggests.

In addition to the mismeasurement of order book liquidity I also examine the time variation of commonality. Firstly, commonality is far stronger at the opening of the trading day when liquidity is low and at the end of the day when the liquidity supply falls again. During the day when liquidity is high commonality is lower. Obviously there are intraday patterns of commonality which seem to be associated with the overall level of liquidity in the market. Likewise commonality is much stronger in falling markets than in rising markets. These results imply that commonality is also associated with the general momentum of the market. In all, this is strong evidence of time variation of commonality in liquidity.

One implication of commonality is that an asset's liquidity will affect the investor's risk of holding the asset. An investor who holds a portfolio of stocks will not be able to eliminate liquidity risk. Furthermore, commonality implies that assets will tend to be illiquid at the same time. This can potentially affect asset prices and lead to the instability of the market. Market-wide liquidity swings will tend to be stronger if a market is already trending downwards. Events such as the 1987 stock market crash show that market-wide liquidity outflows can be substantial and can have a destabilizing impact. Whether commonality really is responsible for market crises or whether effects remain confined to the asset class alone is an open question. Previous research presents first evidence that commonality also exists across asset classes.

In the light of our empirical findings, it is a natural question to ask where commonality comes from. One hypothesis is that commonalities arise from the correlated trading behaviour of market participants. The most plausible determinants seem to be correlated liquidity demands, informed trading or discretionary trading. The theoretical literature leaves a lot of room for models to be developed in this area. Even in the absence of theory, some of these hypotheses should still be accessible to closer empirical examination. The principal components in the PCA are a further source of information as to the identity of the economic factors at work behind the commonality of liquidity.

Chapter 6

Pricing Effects of Liquidity

This chapter investigates the impact of liquidity on asset returns in an electronic limit order book market. I decompose liquidity into three components: the level of liquidity, systematic liquidity risk and the resiliency of liquidity. These three factors are then incorporated into conventional asset pricing tests to examine their impact on stock returns. I find evidence that liquidity gets priced: the return of a stock is decreasing in the level of liquidity, increasing in the systematic component of liquidity risk and decreasing in the resiliency of its limit order book. This implies that liquidity is valued by investors and that it enters the pricing process of a stock.

6.1 Introduction

In general, liquidity is viewed as an important property of an asset that is traded in financial markets. It is widely established that it varies both across assets and also over time. This implies that the purchase price as well as the risk of holding two otherwise identical stocks can differ substantially due to differences in liquidity. Conventional wisdom holds that, ceteris paribus, an investor should require a higher expected return for a stock which is less liquid or whose liquidity supply is more risky. Traditionally, however, asset pricing has focused more strongly on fundamental factors. It has only very recently turned to microstructure issues. In this chapter I take a closer look at such features by decomposing liquidity into different dimensions and examining their impact on stock returns.

From a theoretical point of view the central question is whether the liquidity of a stock should have an effect on the required return at all. Amihud and Mendelson (1986) argue that traders will demand higher returns for holding stocks with larger spreads and, thus, in equilibrium liquidity should be priced. Some further approaches model the impact of different levels of trading costs on returns. These models include Constantinides (1986), Heaton and Lucas (1996), Vayanos (1998) or Gârleanu and Pedersen (2004). They argue that while liquidity should have an impact on prices, the level of the liquidity premium should be very low. Harris (2003) goes as far as saying that the resulting premium should be inconsequential.

At the same time, the empirical literature documents that illiquid assets earn significantly higher returns than liquid assets. The first empirical studies examine the relation between the level of liquidity and expected stock returns (see Amihud and Mendelson (1986), Amihud and Mendelson (1989), Eleswarapu (1997), Brennan and Subrahmanyam (1996), Brennan, Chordia and Subrahmanyam (1998) or Amihud (2002)). The next generation of studies focuses on the pricing of liquidity risk as opposed to the level alone. It includes Pástor and Stambaugh

(2003), Gibson and Mougeot (2004) and Acharya and Pedersen (2005). Together these studies provide positive evidence that both the level and risk of liquidity get priced.[44] However, they differ in the construction of liquidity factors and the strength of the pricing effects. Beyond that, they use low-frequency trade data instead of limit order book data. This leaves some doubt as to whether the observed pricing effects are really due to differences in liquidity in the microstructural sense.

In this chapter I investigate the pricing effects of liquidity risk in an electronic limit order book environment. Firstly, stock exchanges around the world conduct trading almost exclusively through open limit order books nowadays. Therefore limit order book liquidity is the kind of liquidity which is of interest to investors. This implies that pricing studies of liquidity should take place on the basis of limit order book data. Secondly, the limit order book is a very clean and immediate measure of liquidity which measures the costs of any potential trade. This data is necessary to estimate the resiliency of the liquidity supply of any stock. Resiliency is important to investors because it is a measure of trade execution risk: it ensures that trades can be placed rapidly and at reasonable prices. Finally, purely trade-based liquidity measures are biased because trades take place when liquidity is high. In contrast, the limit order book does not have this bias.

The chapter examines empirically, for the first time, the pricing effects that the

[44] There are in fact some studies that present contrasting evidence: for example, Eleswarapu and Reinganum (1993), Chen and Kan (1996) and Chalmers and Kadlec (1998) find that liquidity does not get priced.

liquidity has in electronic limit order book markets. I do not limit my approach to one all-embracing factor of liquidity such as the bid-ask spread or depth alone. Instead I split liquidity into a level component and two risk components. The level component reflects that investors have to make an up-front payment when they buy a stock and incur a further discount when they sell the stock again. The first risk component is resiliency. It measures the execution risk of individual trades. The second risk component of liquidity is commonality which takes into account that liquidity co-moves across stocks. Intuitively speaking, it should be valuable to possess stocks which are liquid when the market is illiquid. I expect a pricing impact of these components because, all other things held equal, a high level of liquidity, a low systematic exposure and a resilient limit order book are favorable for good trade execution. Surprisingly, asset pricing studies have not distinguished these components of liquidity. This chapter aims to fill this gap in the literature.[45]

While the measurement of liquidity levels is fairly straightforward, the correct specification of liquidity risk is more difficult. I use the same methodology as in the previous chapters: I set up a mean reversion model of the limit order book to estimate resiliency (see Chapter 4). This procedure yields measures of the riskiness of trade execution. Then I estimate the exposure of each stock's

[45] From this perspective, this chapter is most closely related to Acharya and Pedersen (2005). They set up a framework in which liquidity, stock returns and market returns are correlated and then distinguish between a level effect and a risk effect of liquidity. However, they do not consider resiliency as a dimension which is valuable to investors. Furthermore, their results rise and fall with the assumption that they have modeled the actual channels through which liquidity enters returns correctly. Thirdly, the estimation is done without order data.

liquidity to market movements of liquidity (see Chapter 5). This yields measures of each stock's systematic liquidity risk. With the help of the Fama and French (1992) asset pricing framework I examine the impact of these liquidity factors on stock returns. More specifically, I determine time-varying beta factors, book equity to market equity as well as firm size and then add liquidity into the pricing equations. Finally, I estimate the simultaneous impact on stock returns in this asset pricing scenario.

I find strong evidence that liquidity plays an important role for the pricing process of assets: both the level of liquidity and liquidity risk get incorporated into stock returns. The higher the level of stock liquidity is, the lower the return that an investor receives from buying that stock. Likewise, the lower the liquidity risk of the stock is, the higher the return paid on the investment. The main implication of these results is that, evidently, investors pay attention to the tradability of stocks. While fundamental factors have an impact on the value that an investor perceives as the fair price of the stock, the quality of the actual trading process itself also enters asset prices. I derive this conclusion on the basis of liquidity measures which are constructed directly from limit order book data. Therefore they are founded strongly in the fine microstructure of the market and unlikely to be proxying other factors.

The fact that the level of liquidity is priced reflects the initial investment and the expected final liquidation costs. Furthermore, liquidity is stochastic and varies along time. Therefore investors run the risk of having to liquidate positions at times when their stocks are particularly illiquid. I split liquidity risk into a systematic component and into resiliency. I find evidence that both components get priced as expected. In terms of magnitude, the pricing effects for the level of liquidity and systematic liquidity risk are very similar. A difference of one standard deviation across stocks will lead, all other things held equal, to an increase of about 10% return per annum (or 0.9% per month). This is roughly the same magnitude as for the remaining fundamental factors in the pricing framework. The liquidity premium for resiliency is about half the size.

The chapter is organized as follows: Section 6.2 presents the liquidity measures, briefly summarizes the microstructure models that are used for the estimation and examines the linkage between the various components of liquidity. In Section 6.3 I present the methodology of the asset pricing tests. Section 6.4 reports the empirical evidence of the pricing effects. First I give simple correlations and then report the test results. Section 6.5 analyses the impact of the liquidity measures to check the robustness of the results. Section 6.6 concludes.

6.2 Liquidity Measures and Pricing Factors

A frequent definition of liquidity is that it measures the speed and ease with which an investor can trade into or out of positions in a financial market. An asset is generally considered liquid if transactions take place at a reasonable price and without having to wait too long. The previous sections have already discussed the looseness of this concept. I now discuss how I measure the level and risk of liquidity for the asset pricing tests.

The most frequent proxies of liquidity are the bid-ask spread and the corresponding depth at the best bid and ask price. Therefore I use them as measures of the level of liquidity.[46] While the level of liquidity is a fairly easy input to the asset pricing tests, the risk dimension cannot be proxied that directly. Let us take a closer look at the microstructure models that I use to assess a stock's systematic component of liquidity risk and the resiliency of its limit order book. The estimation of resiliency is based on the model in Chapter 4 and the estimation of the systematic liquidity risk is based on the Chordia, Roll and Subrahmanyam (2000) model that was presented in Chapter 5. I briefly summarize these approaches:

1. **Resiliency:** The first component of liquidity risk that I model is the resiliency of the limit order book. We would expect resiliency to enter returns because higher resiliency ensures that liquidity shocks are overcome rapidly and that large trades can be executed at good prices. I set up a mean reversion model of the liquidity supply and then adopt a cross-sectional SUR for the empirical specification. The cross-sectional SUR approach yields joint GLS estimates which are corrected for contemporaneous correlation and heteroscedasticity. Let L_{t-1} denote the level of liquidity in the past period and let ΔL be the change in liquidity in the current period,

[46] To make sure that these measures are not too noisy, I will extend the spread and depth to ticks that are deeper in the book to ensure the robustness of the results.

$\Delta L_t := L_t - L_{t-1}$. I run the following regression:

$$\Delta L_{i,t} = \alpha_i - \varphi_i L_{i,t-1} + \Sigma_{k=1}^{n} \gamma_k \Delta L_{i,t-k} + \varepsilon_{i,t} \tag{6.1}$$

In the equation, the parameter φ measures the mean reversion of liquidity while the γ parameters are lag coefficients. ε is a normally distributed white noise error term. As in Chapter 4 the estimation of resiliency is based upon the bid-ask spread and depth around the spread. To reduce the noise in the data I also use depth up to the first three ticks in the book as well as bid and ask half-spreads for the third step of the limit order book.

2. **Commonality:** The second component of liquidity risk that I model is systematic component of liquidity risk. We would expect that the systematic component of liquidity enters returns because, ceteris paribus, stocks with higher systematic liquidity risk are riskier and should provide compensation for the additional risk that investors bear. The standard econometric approach is the market model in Chordia, Roll and Subrahmanyam (2000) which relates the liquidity of a single stock to the liquidity of the market. It employs a regression framework to measure the sensitivity of stock i's liquidity, L_i, to market liquidity, L_M:

$$\begin{aligned} L_{i,t} = \alpha + \beta_i^1 L_{M,t} + \beta_i^2 L_{M,t+1} + \beta_i^3 L_{M,t-1} + \delta_i^1 r_{M,t} + \delta_i^2 r_{M,t+1} \\ + \delta_i^3 r_{M,t-1} + \xi_i VOL_{i,t} + \varepsilon_{i,t} \end{aligned} \tag{6.2}$$

In the equation, the market return, $r_{M,t}$, and stock volatility, $VOL_{i,t}$, are used as control variables. ε is a normally distributed white noise error term. As in Chapter 5 the estimation of commonality is based on the bid-ask spread, depth at the spread as well as the depth at 2% price impact in the limit order book and the slope of the price-quantity schedule. These measures are standardised to correct for predictable trends in the data. The cut-off point of the limit order book is 2%.

In the context of the following pricing sections, the parameters of interest are φ_i in Equation 6.1 and β_i^1 in Equation 6.2, the coefficients of resiliency and systematic liquidity risk. These two measures characterize the liquidity risk of an investor when he trades a stock in a limit order book market. The resiliency measure is negative in its exposure to liquidity risk (the higher the coefficient, the lower the risk), while the commonality coefficient is positive in its exposure (the higher the coefficient, the higher the risk). Since investors should get compensated for holding risky and illiquid assets, let us depart with the hypothesis that in the asset pricing tests the sign of the resiliency measure should be negative and the sign of the commonality measure should be positive.

6.3 Methodology

In the following section I determine the Fama-French factor inputs for the asset pricing tests. I then present the conventional Fama and MacBeth (1973) test methodology and extend the approach to a joint GLS estimation to obtain higher statistical power. Finally I discuss potential errors and biases in the results.

6.3.1 Fama-French Factors

The limit order book data used for the previous studies of liquidity was presented in Chapter 3. For the asset pricing tests I need additional data on stock returns and firm characteristics. The data is available in Datastream. I collect return information and firm characteristics for all stocks in the DAX 100 index which were listed at the German stock exchange for the period from 1993 to 2004. Although the asset pricing tests will only use DAX 30 stocks, I use the DAX 100 stocks to construct the beta factors.[47] Return data is collected on a monthly basis which yields 144 monthly returns for each stock. Firm characteristics are sampled on a yearly basis.

The methodology that I will be using follows the Fama and French (1992) approach to allow for the comparability of my results to previous studies. Fama and French (1992) explored the determination of the cross-sectional variation in returns and found that the beta factor, firm size and the ratio of book equity to market equity all influenced stock returns. Therefore, I also include these factors in the asset pricing tests. A problematic feature to note is that the period that I use is fairly short (10 years), while asset pricing studies typically use very long time series. However, the longer the period that I choose, the less likely

[47] The reason for using more than 30 stocks will become apparent in the following discussion of the methodology. The central idea is that we use more stocks to be able to construct beta portfolios. If I only used thirty stocks I would only be able to construct either very few portfolios (for example only 3 ten-stock portfolios) or otherwise very small portfolios (for example 10 three-stock portfolios). However, I need a sufficient number of portfolios to be able to construct a sufficient time variation in beta and I need sufficiently large portfolios to reduce the estimation error in beta. The estimation procedure is explained in the following section.

liquidity features are to be valid for whole time horizon. I go back ten years because that spans the time for which trading has been organized in limit order book markets. I do not consider any longer time horizons because the market structure has changed too much in the last few years for the pricing of liquidity to yield meaningful results over such long periods.

Methodologically the most challenging task in the Fama and French (1992) framework is to determine the beta factor. Beta is considered to be time-varying, yet unfortunately it cannot be observed directly and has to be estimated. This poses a number of challenges which have been discussed extensively in the literature (see Black, Jensen and Scholes (1972), Fama and MacBeth (1973), Roll (1977), Fama and French (1992), Black (1993) or Kothari, Shanken and Sloan (1995)). I follow the Fama and French (1992) of using portfolio betas to proxy the time variation in beta. First, I estimate so-called pre-ranking portfolio betas for each stock by using two years of monthly returns, $R_{i,t}$. Let t be a time index in months, let T be the index for each test year and let i be an index for each stock. To obtain pre-ranking betas for each year, I estimate the following equation:

$$R_{i,t} = \alpha_i + \hat{\beta}_i (r_{M,t} - r_{f,t}) + \varepsilon_{it} \tag{6.3}$$

In the equation, ε is an i.i.d. error term. $R_{i,t}$ is the excess return of stock i in t over the risk-free rate, $R_{i,t} := r_{i,t} - r_{f,t}$. I use the German 1-month interbank rate, Fibor and Euribor,[48] as the return of the risk-free asset and the performance

[48] The sample period covers the introduction of the Euro in Germany which is why the Fibor rates (denominated in Deutschmark) were renamed into Euribor (denominated in Euros).

index DAX 30 as the proxy for the market return $r_{M,t}$. Equation 6.3 is estimated for each year T with the past 24 monthly returns so that I obtain one pre-ranking beta for each stock and for each year, $\hat{\beta}_{i,T}$. By construction, there is some overlap in the return data for the pre-ranking betas.

In the next step, the pre-ranking betas are sorted into deciles according to their beta factor each year. An important point to notice is that the composition of each beta decile changes from year to year. Monthly portfolio returns are calculated for each portfolio as the equally-weighted averages of individual stock returns, $r_{p,t}$, where p is an index for the portfolios. The monthly portfolio returns are then aggregated to yearly returns, $r_{p,T}$, and are regressed onto the yearly market returns, $r_{M,T}$:

$$R_{p,T} = \alpha_p + \hat{\beta}_p(r_{M,T} - r_{f,T}) + \varepsilon_{p,T} \tag{6.4}$$

This procedure yields one beta estimate over the whole period for each portfolio, $\hat{\beta}_p$. An advantage of this procedure is that the use of portfolios rather than individual securities diversifies most of the firm-specific return component and therefore enhances the precision of the estimates of beta (see Black, Jensen and Scholes (1972)). As I will be using individual stocks in the cross-sectional regressions, I take the portfolio beta factor of the portfolio to which stock belongs as its individual beta factor each year as in Fama and French (1992). Because stocks belong to different portfolios each year, their individual stock betas vary over time, $\hat{\beta}_{i,T}$.

The remaining factors in the Fama and French (1992) framework can be observed directly. The size factor is the market value of stock i at the end of year T. The book-to-market factor takes the ratio of a firm's book equity to its market capitalization. I take yearly time series from Datastream for each value at the end of year T. Like the literature, I set values of BE/ME below 0.005 to 0.005 and values above 0.995 to 0.995 and then take the log of the ratio, $BM_T := ln(BE_T/ME_T)$. With the beta estimates we obtain all inputs of the original Fama and French (1992): a market factor, a size factor and a book-to-market factor. We only use DAX30 stocks as these are the only stocks for which we can determine the liquidity factors. Finally, we follow Fama and French (1992) in deleting all financial companies because book equity to market equity has a different meaning for banks and financial corporations.[49]

6.3.2 Estimation Procedure

I now bring together the liquidity factors from section 6.2 and the Fama-French factors from section 6.3. The aim is to analyze the individual pricing effects of these risk components in a simultaneous model. The original model of Fama and French (1992) regresses the cross section of stock returns onto the stocks' firm characteristics:

$$R_i = \gamma_0 + \gamma_1\hat{\beta}_i + \gamma_2 SIZE_i + \gamma_3 BM_i + \varepsilon_i \tag{6.5}$$

[49] In my sample the financial firms are Allianz, Commerzbank, HypoVereinsbank, Münchener Rück and Deutsche Bank.

I now add the level, resiliency and commonality of liquidity ($LLEV$, $LRES$ and $LSYS$) to the model. As stated before, asset pricing tests require long time series. Typically, Equation 6.5 will be estimated for each year of the sample period, while we only have single observations of the liquidity factors. Therefore I assume that the liquidity factors remained constant over the estimation period and thus generate the same return contribution every year. In principle, there is no harm in this assumption, yet it will bias my study against finding significant results.[50] With the estimates of the liquidity level and liquidity risk, I obtain the following new pricing equation:

$$\begin{aligned} R_{i,T} \;&=\; \gamma_{0,T} + \gamma_{1,T}\hat{\beta}_{i,T} + \gamma_{2,T} SIZE_{i,T-1} + \gamma_{3,T} BM_{i,T-1} + \gamma_{4,T} LLEV_i \\ &\quad + \gamma_{5,T} LSYS_i + \gamma_{6,T} LRES_i + \varepsilon_{i,T} \end{aligned} \tag{6.6}$$

In Equation 6.6, $R_{i,T}$ is the yearly excess return of stock i in year T. The equation is estimated for each year of the data sample from 1995 to 2004 over the cross-section of all stocks i. $\gamma_{j,T}$ with $j = 0, ..., 6$ are the estimated coefficients that I obtain each year. $\varepsilon_{i,T}$ is the serially uncorrelated error term with a mean of zero. Finally I obtain a time series of 10 parameter estimates for each factor in Equation 6.6.

The standard procedure in Fama and MacBeth (1973) is to test whether the

[50] An alternative approach is to identify observable variables which determine liquidity. With the help of such variables, liquidity factors could be reconstructed under the assumption that the relationship of the observable variables and liquidity stays unchanged. This has the advantage that liquidity factors become time-varying. The previous chapters suggest possible factors. Also see Beltran, Giot and Grammig (2005) for an investigation of factors that help predict the liquidity of the limit order book.

average of the parameter estimates, $\bar{\gamma}_j = \frac{1}{T}\Sigma_{T=1}^{10}\hat{\gamma}_{j,T}$, are significantly above or below zero. Litzenberger and Ramaswamy (1979) derive a correction technique for time-varying volatility in which each coefficient is weighted by its precision (as measured by the standard error) when computing cross-sectional averages. In effect, this corresponds to a cross-sectional GLS estimation.[51] I use the Fama and MacBeth (1973) and Litzenberger and Ramaswamy (1979) as the starting point to obtain estimates that are comparable to the previous studies. However, as both the cross section of stocks as well as time horizon are relatively small, these standard tests will tend to have low statistical power. In a second step I therefore follow Amihud, Christensen and Mendelson (1992) who combine time series and cross-sectional observations in their estimation. This yields one single estimation equation:

$$\begin{aligned} R^s &= \gamma_0 + \gamma_1\hat{\beta}^s + \gamma_2 SIZE^s + \gamma_3 BM^s + \gamma_4 LLEV^s + \gamma_5 LSYS^s \\ &\quad + \gamma_6 LRES^s + \varepsilon^s \end{aligned} \tag{6.7}$$

In Equation 6.7, the superscript s indicates that all variables which carry the index have been stacked: R^s is an $i \times T$ vector which consists of i subvectors of length T. Each subvector is the return time series of one stock which is then followed by the time series of the next stock and so forth. Likewise, β^s, $SIZE^s$ and BM^s are also $i \times T$ vectors which stack the time series of the asset pricing factors. $LLEV^s$, $LSYS^s$ and $LRES^s$ are similar $i \times T$ vectors, yet each

[51] For further details see Litzenberger and Ramaswamy (1979).

subvector i consists of constants because we assume that liquidity is constant over the sample period. In this pooled approach, I use a standard t-test to test whether the estimates of γ are unequal to zero.[52] The main difference between Equations 6.6 and 6.7 is that Equation 6.6 is estimated each year, while the joint approach is estimated only once over the whole time horizon.

Let us briefly discuss the correct estimation technique for the pooled approach. GLS is the correct estimation method[53] where I proceed as follows: With the help of the original OLS residuals I estimate the variance-covariance matrix of stock returns for cross-sectional heteroscedasticity, $\hat{C}$. The next step is to determine the weighting matrix W such that $C^{-1} = W'W$. All variables are then pre-multiplied by the weighting matrix and Equation 6.7 is re-estimated for the transformed variables:

$$\begin{aligned} R^* &= \gamma_0 + \gamma_1 \hat{\beta}^* + \gamma_2 SIZE^* + \gamma_3 BM^* + \gamma_4 LLEV^* + \gamma_5 LSYS^* \\ &\quad + \gamma_6 LRES^* + \varepsilon^* \end{aligned} \tag{6.8}$$

In Equation 6.8, the superscript * indicates that the stacked variables have been replaced by their transformed counterparts. I can now use the standard OLS statistical inference to test for the significance of γ_j. Additionally, I use the

[52] This is the approach used in Amihud and Mendelson (1986), Amihud and Mendelson (1989) and Amihud, Christensen and Mendelson (1992).

[53] See, in the context of the CAPM, Brown and Weinstein (1983), Amihud and Mendelson (1986), Amihud and Mendelson (1989) and Shanken (1992).

Newey and West (1987) correction for possible remaining autocorrelation in the data.[54]

6.3.3 Potential Errors and Biases

Asset pricing tests in the tradition of Fama and French (1992) face the problem that their factors do not have a solid theoretical foundation. Therefore, the tests always run the danger that they are in fact using variables which have no causal relationship with returns. Instead, they might simply be capturing the measurement error in other factors such as beta. Likewise, the proposed factors could simply be proxies of more fundamental factors like uncertainty or information asymmetry which are not directly observable. For example, Miller and Scholes (1982) and Berk (1995) point out that many factors in asset pricing tests depend on the stock price. In turn, the stock price's inverse value is a good proxy of conditional beta. A similar point can be made for liquidity: Novy-Marx (2003) argues that a stock's exposure to risk leads to high return premia and low liquidity. Subsequently, it is not surprising that liquidity risk is associated with high returns, however that does not mean that high liquidity risk is the reason for high stock returns. Let us therefore take the time to check how reliable the factors in the following asset pricing tests are.

As the estimates of liquidity risk are not related to the stock price, they should definitely not be a proxy of beta. With regard to the second point it

[54] The estimation of the variance-covariance matrix, the GLS weighting method and the Newey and West (1987) correction are all implemented in EVIEWS.

is always difficult to rule out completely that a variable is just a proxy of a different unobservable pricing factor. However, I try to keep this danger as low as possible by controlling for the standard Fama-French factors. Secondly, the liquidity measures are connected very directly to each stock's limit order book, so that they should be reliable liquidity estimates. We can conclude that it is unlikely that we are confounding liquidity with any other factors. Therfore let us return to Equations 6.6 and 6.7. The main coefficients of interest are γ_4, γ_5 and γ_6 which we will now examine in more detail.

6.4 Asset Pricing Test Results

In the following section I present some basic properties of the pricing factors as well as their correlation structure. I then examine the impact that liquidity risk has in the Fama and French (1992) model and in the joint GLS framework.[55]

6.4.1 Correlation Structure of Pricing Factors

Let us first have a look at some summary statistics of the variables which are used in the asset pricing tests. Table 6.1 shows that over the whole period the mean excess return over the risk-free asset is 4.93%. There is a wide range between the lowest and highest annual return over the time period (-70.69% and 211.07%). The range of beta factors obtained from the portfolio approach reaches from 0.33

[55] The liquidity measures used are the bid-ask spread (for the level of liquidity), depth at 2% price impact in the limit order book (to estimate commonality) and depth up to the third tick (to estimate the resiliency of the limit order book). The results for all further measures are presented in the robustness section to underline that the following conclusions are irrespective of the liquidity proxies.

Table 6.1: Asset Pricing Inputs

	Mean	Median	S.D.	Min	Max
Excess Returns	0.049	0.045	0.240	-0.707	2.110
Beta Factor	0.871	0.920	0.278	0.330	1.277
Market Capitalization	2.745	2.208	0.962	0.190	4.883
Book Equity to Market Equity	-0.636	-0.501	0.498	-2.813	0.598
Level of Liquidity	0.084	0.084	0.024	0.045	0.136
Resiliency of Limit Order Book	0.536	0.545	0.179	0.225	0.861
Systematic Liquidity Risk	1.489	1.593	0.364	0.230	2.104

Table 6.1 gives the inputs of the asset pricing tests. Returns are the excess returns over the German interbank rate. The beta factors are the portfolio betas that are obtained from the Fama and French (1992) estimation procedure and assigned to individual stocks. Market capitalization of equity is the size factor at the end of the previous year (logarithmic values of equity in billions of Euros). Book equity to market equity is also given in log values. The level of liquidity is the average relative bid-ask spread (in %). The two components of liquidity risk are the estimates of resiliency and of commonality. All statistics are calculated over the last five years where there are no missing values.

to 1.28. Market capitalization and book equity to market equity are taken as logs, so that these values are more important in terms of their relative size compared to other stocks as opposed to their absolute levels. On average, the relative bid-ask spread is 0.084%. Liquidity risk in terms of execution risk (i.e. resiliency) is taken as the sum of the ask side and bid side coefficients. It ranges from 0.225 to 0.861. Low values imply strong resiliency and thus lower risk. Systematic liquidity risk, also computed as the sum of ask side and bid side coefficients, varies between 0.23 and 2.10. For systematic risk, high values imply strong systematic liquidity risk.

The first step that I take in investigating the return impact of the pricing factors is a canonical correlation analysis. Table 6.2 shows the correlation structure of all variables that enter the asset pricing regressions. The critical value

Table 6.2: Correlation Structure

	RET	BET	SIZE	BEME	LLEV	LRES	LSYS
Excess Returns	1.000	0.032	-0.165	-0.086	0.228	-0.110	0.104
Beta Factor		1.000	0.143	0.300	0.059	0.443	-0.041
Size Factor			1.000	-0.296	-0.510	-0.241	0.012
BE-ME Factor				1.000	0.146	-0.170	-0.024
Level of Liquidity					1.000	-0.049	0.020
Resiliency						1.000	-0.043
Commonality							1.000

Table 6.2 gives the correlation matrix of the asset pricing input factors. Returns are the excess returns over the risk-free rate. Beta factors are obtained from the Fama and French (1992) estimation procedure. Market capitalization of equity is the size factor (logarithmic values of equity in billions of Euros). Book equity to market equity is also given in log values. The liquidity level is measured by the bid-ask spread. The two components of liquidity risk are the estimates of resiliency and of commonality. All statistics are calculated over the cross section of sample stocks.

at the 5% level is 0.19 for my data sample. The strongest and clearly significant correlation is between the size factor and the bid-ask spread (-0.51): as we would expect large stocks tend to have low bid-ask spreads. Another high correlation is between the beta factor and the resiliency of the limit order book (0.44): stocks with high beta factors tend to have very resilient limit order books. All other correlation coefficients are in the range of -0.30 to 0.30. Next let us look at the correlations of the factors with the stock return. The correlation coefficient for the bid-ask spread and the return is 0.23 which is significant, while for resiliency and the return the correlation is -0.11 and for commonality and the return 0.10. The coefficients are not particularly large in value, yet they all have the expected signs. The correlation of the returns with the Fama-French factors is not much higher, either. Return and beta are positively correlated (0.03), while the coefficient for the size factor is negative (-0.16). These coefficients have the same sign as often reported in previous studies: high beta stocks and small-cap firms earn higher returns. The association between returns and book equity to market equity is negative (-0.09). This is opposed to the original positive relationship found in Fama and French (1992). However, Loughran (1997) reports that the BE/ME factor is an effect predominant in the technology section (for example NASDAQ stocks). Easley, Hvidkjaer and O'Hara (2002) also find an opposite sign of BE/ME for NYSE stocks.

A variable that has received no attention so far is trading volume (or turnover). Datar, Naik and Radcliffe (1998) find evidence of a negative relationship between turnover and returns. Chordia, Subrahmanyam and Anshuman (2001) also argue that both turnover as well as turnover volatility should affect stock prices. They present evidence that there is a negative relationship between these variables and returns. As turnover and firm size tend to be heavily correlated, I check their association and observe a correlation of 92.3%. Subsequently, firm size seems to be a very strong proxy for trading volume (or vice versa). Because of the high association, I do not include volume as a separate variable. I stick with firm size instead to maintain comparability with Fama and French (1992).

6.4.2 Pricing of Liquidity and Liquidity Risk

In this section I present the results of the asset pricing tests. Primarily I will be focusing on the liquidity factors in the pricing framework – the level of liquidity,

Table 6.3: Fama-MacBeth Asset Pricing Test

	Fama-MacBeth		LR-Correction	
	Coeff.	t-stat.	Coeff.	t-stat.
Constant	-0.172	-1.366	-0.028	-0.105
Beta Factor	-0.076	-0.661	-0.078	-0.728
Market Capitalization	0.050	1.358	0.034	0.880
Book Equity to Market Equity	-0.111	-2.071	-0.116	-1.710
Level of Liquidity	0.169	0.174	0.871	0.702
Resiliency of Limit Order Book	-0.042	-0.329	-0.020	-0.139
Systematic Liquidity Risk	0.100	1.488	0.069	1.195

Table 6.3 shows the results of the Fama-MacBeth asset pricing test. The coefficients are yearly averages and the t-statistics test whether the averages are equal to zero. The dependent variable is the yearly excess stock return over the German Interbank rate. The exogenous variables are the beta factor, market capitalization (logs of equity in billions of Euros) and book equity to market equity (logs). The level of liquidity is measured by the relative bid-ask spread while liquidity risk is captured by resiliency and commonality.

resiliency and systematic liquidity risk.[56] We start with the Fama and MacBeth (1973) framework and then compare the results to the pooled approach. After we have examined the pricing effects of liquidity, we will also have a look at the conventional Fama-French factors at the end of the section.

Table 6.3 displays the time series averages of the estimated coefficients in the cross-sectional Fama-MacBeth approach. Both the conventional Fama and MacBeth (1973) test as well as the Litzenberger and Ramaswamy (1979) correction yield very similar results. The coefficients of all three liquidity factors have the expected signs: firstly, the standard Fama and MacBeth (1973) coefficient of the bid-ask spread is 0.169 (with a t-statistic of 0.174) and the Litzenberger and Ramaswamy (1979) coefficient is 0.871 (with a t-statistic of 0.702). For re-

[56] For the sake of brevity I report the results for the bid-ask spread. The results for depth at the best spread are presented in the robustness check. However, the results remain unchanged.

Table 6.4: Pooled Asset Pricing Test

	Coefficient Estimates	Coefficient t-statistics	Impact on Returns %
Constant	-0.464	-2.533	
Beta Factor	0.001	0.010	0.273
Market Capitalization	0.055	1.656	5.240
Book Equity to Market Equity	-0.296	-5.126	-14.767
Level of Liquidity	0.013	2.246	9.412
Resiliency of Limit Order Book	-0.240	-1.429	-4.310
Systematic Liquidity Risk	0.301	3.607	10.973

Table 6.4 shows the results of the pooled asset pricing test. In the pooled approach the cross-section and time series are jointly estimated and tested against zero with standard t-tests. The dependent variable is the yearly excess stock return over the risk-free rate. The exogenous variables are the beta factor, market capitalization (logs of equity in billions of Euros) and book equity to market equity (logs). The level of liquidity is measured by the relative bid-ask spread while liquidity risk is captured by resiliency and commonality. The impact on returns gives the return impact (in %) for an increase of the exogenous variables by one standard deviation.

siliency, the standard coefficient is -0.042 and the corrected coefficient is -0.020 (with t-statistics of -0.329 and -0.139, respectively). Finally, the Fama and MacBeth (1973) estimate for commonality is 0.100 (with a t-statistic of 1.488) and the Litzenberger and Ramaswamy (1979) coefficient is 0.069 (with a t-statistic of 1.195). All together, high spreads, low resiliency and high commonality imply higher returns. However, the coefficients are not significant. As the yearly regressions have only 25 observations, this is not very surprising. We therefore turn to the joint estimation of the time series and cross section to examine whether lack of statistical power or lack of causality are responsible for the insignificant results.

The results for the pooled asset pricing tests are provided in Table 6.4. Again we focus on the parameter estimates of the liquidity factors. In the pooled approach we find that all estimates have the same sign as they had in the Fama and MacBeth (1973) and Litzenberger and Ramaswamy (1979) approaches. However, the pooling of the time series and cross-sectional observations has clearly increased the statistical power of the test: in the joint approach, the liquidity factors have a very pronounced effect on stock returns. Firstly, the coefficient of the bid-ask spread is 0.013 with a t-statistic of 2.246. As we would expect, stocks with higher bid-ask spreads (less liquid stocks) earn higher returns, all other things equal. Thus, investors get compensated for holding illiquid stocks. Secondly, systematic liquidity risk has a significantly positive impact on returns. The estimated coefficient is 0.301 and the corresponding t-statistic is 3.607: the higher the systematic component of liquidity risk is, the higher the return that is paid to the investor. Finally, the coefficient of resiliency is -0.240 with a t-statistic of -1.429. Subsequently, less resilient stocks earn higher returns. This result is significant at the 10% level, while the other liquidity coefficients are significant at the 5% level. In all, the coefficients of all liquidity factors have an important impact on stock returns. The Adjusted R^2 value of the model is 24.05% and the hypothesis that all parameters are significant is rejected even at the 1% level.

I now assess the economic importance of the pricing factors. The final column of Table 6.4 displays the magnitude of return differentials which are due to cross-sectional differences in the liquidity factors. An increase of one standard deviation of the bid-ask spread across stocks leads, all other things held equal, to an increase of the annual stock return of 9.41 percentage points. Likewise, the same increase of systematic liquidity risk leads to an increase of the return by 10.97 percentage points. With regard to resiliency, an increase of the resiliency factor by one standard deviation leads to a decrease of the return of 4.31 percentage points. Comparing these figures we see that the impact of the bid-ask spread and of systematic liquidity risk on returns are about twice as strong as the impact of resiliency. One reason why resiliency has a weaker impact might be that its risk is diversifiable across assets. In contrast, systematic liquidity risk cannot be diversified. Likewise, the bid-ask spread cannot be avoided, either.

Let us turn to the Fama-French factors in the asset pricing regressions. Their impact on returns varies greatly from factor to factor. The coefficient of the size factor is 0.055 with a t-statistic of 1.656. This result implies that larger firms earn higher returns. It is contrary to the well-known size effect (see Banz (1981)), however it is consistent with Easley, Hvidkjaer and O'Hara (2002). They also observe that the impact of market capitalization on returns is positive if the pricing equation includes a separate liquidity factor or a factor of informed trading. The most plausible reason is that with a separate liquidity factor in the equation, market capitalization does not proxy liquidity anymore and therefore need not have a negative coefficient. The coefficient of the beta factor is 0.001 and with a t-statistic of 0.010. Obviously, beta has little power to explain stock returns. This evidence is in line with prior research as in Fama and French (1992), Chalmers and Kadlec (1998) or Datar, Naik and Radcliffe (1998). The book-to-market factor also has a coefficient of -0.296 and a t-statistic of -5.128. The coefficient has a different sign compared to Fama and French (1992). However, as discussed before, a positive relationship of book equity to market equity with returns is primarily observed for NASDAQ stocks. As our sample does not include high-tech stocks, the opposite sign is therefore not so surprising (see Easley, Hvidkjaer and O'Hara (2002)). If we compare the magnitude of the Fama-French factors with the liquidity factors we see that beta has virtually no impact on returns. The impact of market capitalization is about as high as the impact of resiliency (5.24 percentage points). In contrast, the ratio of book equity to market equity has a stronger impact on returns than liquidity risk (14.76 percentage points).

In all, we observe that all dimensions of liquidity and liquidity risk get priced. The bid-ask spread and systematic liquidity risk have the strongest effects on stock returns. This is not surprising as systematic liquidity risk is non-diversifiable and the spread is a payment that is always incurred for any stock investment.

6.5 Robustness Checks

In this section I examine the robustness of the previous results. Firstly, I take a closer look at the influence which the choice of liquidity measures has on the estimation results. The measure of liquidity enters the estimation directly through the level of liquidity and indirectly through its impact on the estimates of resiliency and commonality.

Table 6.5: Asset Pricing Tests with Different Liquidity Measures

Pricing Factor	Measure	Coefficient	t-statistic	Impact
Level of Liquidity	SPR_i	0.013	2.246	9.412
	DEP_i	-0.931	-0.824	-4.516
	λ_i	0.204	1.712	7.008
Resiliency of Limit Order Book	$DEP^i_{k=3}$	-0.240	-1.429	-4.310
	$HS^i_{l=3}$	-1.024	-1.265	-5.126
Systematic Liquidity Risk	SPR_i	0.188	1.743	5.024
	$DEP^{ti}_{A,B}$	0.301	3.607	10.973
	λ_i	0.084	0.948	3.599

Table 6.5 shows the results of the pooled asset pricing test for different liquidity measures. The table displays the parameter estimates and t-statistics of the liquidity factors. The level of liquidity and systematic liquidity risk are estimated for the spread (SPR), depth (DEP) and the slope of the price-quantity schedule (λ). Resiliency is estimated for the depth of the limit order book at the third tick ($DEP^i_{i=3}$) and the area under the price-quantity schedule ($HS^i_{l=3}$) .

To assess the impact of the liquidity proxy I re-estimate Equation 6.8 for different measures of liquidity and for different estimates of resiliency and commonality. I use three different proxies for the level of liquidity: the bid-ask spread (SPR_i), the corresponding depth of the limit order book ($DEP^i_{A,B}$) and the slope of the price-quantity schedule (λ_i). With regard to resiliency, I estimate the mean reversion model in Equation 6.1 for depth at the third tick ($DEP^i_{k=3}$) and for the half-spread ($HS^i_{l=3}$) on the bid and ask side (see Chapter 4). Likewise, I estimate the systematic component of liquidity risk in Equation 6.2 on the basis of the spread, the depth of the limit order book (at 2% price impact) and the slope of the price-quantity schedule (see Chapter 5). I vary the liquidity measures for one factor at a time so that all other factors are always equal to the specification in the previous section. The results are provided in Table 6.5.

Table 6.5 shows that the only sign that changes is when I proxy the level of liquidity by depth instead of the bid-ask spread. However, this is consistent with the previous results, as a high bid-ask spread reflects low liquidity, while high depth reflects high liquidity. Therefore, in this case opposite signs actually underline that the effect of the level of liquidity on returns is the same for both liquidity measures. For all other proxies, the signs and directions of the impact on returns also remain the same. While the directions of the impact remain unchanged, the statistical power of most alternative liquidity measures is lower: all estimates were significant in the previous section, yet after the robustness checks only about two thirds of the coefficients remain significant. Likewise, the magnitude of the effects is smaller for the level of liquidity and for the systematic component of liquidity risk. The results for resiliency are unchanged.

In all, the results are qualitatively similar with regard to the signs of the coefficients, however the magnitude of the return effects differs for the various liquidity measures. The strongest effects can be observed if we measure liquidity by the bid-ask spread and by the depth of the limit order book. Maybe it is not surprising that the pricing effects of the slope are not very high because investors do not observe the slope of the price-quantity schedule directly. Furthermore, it is reasonable to assume that the slope measure is a particularly noisy pricing factor: the slope is an estimated variable which is used to obtain an estimate of commonality. This double estimation makes it particularly likely to be subject to an error-in-variables bias. The bias will tend to underestimate the coefficient in the pricing equation and thus biases against finding significant results.

6.6 Conclusion

This chapter investigates the impact of liquidity on stock returns. In the past, asset pricing has focused very strongly on more fundamental factors of asset prices and has only very recently turned to the characteristics of the actual trading process. In particular, aspects of liquidity have not been examined very comprehensively. I split liquidity into three factors: the level of liquidity, the systematic component of liquidity risk and the resiliency of the limit order book. We would expect that, ceteris paribus, an investor will demand higher returns for stocks which are less liquid, which have higher systematic liquidity risk or which have a less resilient liquidity supply. The chapter examines empirically, for the first time, the simultaneous pricing effects of these liquidity factors on the basis of tick data from a limit order book market.

We find strong evidence that liquidity plays an important role for the pricing process of assets: both the level of liquidity and liquidity risk are incorporated into stock returns in a significant way. The lower the level of liquidity is, the higher the return that investors receive from buying that stock. This effect is the strongest if we measure liquidity by the bid-ask spread or by the depth of the limit order book. Likewise, the higher the liquidity risk of a stock is, the higher is the return that an investor receives. This result can be observed both for the systematic component of liquidity risk as well as for the resiliency of the limit order book. This chapter presents evidence that investors pay attention to the tradability of stocks and that they require higher premia if they hold stocks which cannot be traded as easily or whose liquidity supply is more risky.

The results provided in this chapter are consistent with several studies in the literature that have provided evidence that the level of liquidity gets priced (see Amihud and Mendelson (1986), Amihud and Mendelson (1989) or Amihud (2002)). Likewise, the evidence is also consistent with the studies by Pástor and Stambaugh (2003) and Gibson and Mougeot (2004) who first put forward that liquidity risk gets priced. The only other study that examines the level and a risk component of liquidity simultaneously is Acharya and Pedersen (2005). In their framework, the impact of the level is more than twice as strong as the impact of liquidity risk. The results in this chapter confirm that both the riskiness and level of liquidity are priced. However, in my study the level of liquidity and its systematic risk have an equally strong impact on returns, while the impact of resiliency is much lower. One reason why liquidity risk has a stronger impact in my study is probably that the use of limit order book data allows a more precise estimate of liquidity risk.

In all, I provide positive evidence that liquidity is a stock characteristic that is valued by investors and thus enters stock returns. In this chapter, liquidity is split up into three components to determine the magnitude of the individual aspects of liquidity. The pricing effects for the level of liquidity and for the systematic component of liquidity risk are very similar: a difference of one standard deviation across stocks will lead, all other things held equal, to a return increase of about 10% return per annum. For resiliency, an increase of one standard deviation leads to a return increase of about 5%. These results are valid under the assumption that the liquidity factors are constant over time. When longer time series of the limit order book are available, future studies will be able to examine these effects for time-varying liquidity factors.

Chapter 7

Conclusion

In the previous chapters I examined the riskiness of liquidity in a limit order book market. Liquidity risk typically constitutes the danger that the liquidation of an asset can be very costly because trades incur high market impacts in illiquid markets. Two important aspects of such liquidity-induced price movements are, firstly, the replenishment process of the limit order book and, secondly, the spill-over process of liquidity from one stock to another. For investors, liquidity risk matters when they want to time their trades or when they split large volumes over time. In addition, liquidity spill-overs become relevant as soon as investors trade multiple securities. Against this background my thesis investigated the properties and magnitude of liquidity risk. Furthermore, it examined the pricing effects of liquidity risk on stock returns. Let us conclude by reviewing the most important results and implications for future research.

7.1 Main Results

Market structures have been changing and limit order book markets growing so fast that the theoretical literature has not kept pace with this rapid development. Therefore, all theoretical models of the limit order book are only partial views of limit order markets. Most notably, static models such as Glosten (1994) or Sandås (2001) derive the equilibrium state of the price-quantity schedule. Dynamic models such as Foucault, Kadan and Kandel (2005) characterize the dynamic equilibrium of very simple limit order book markets. The challenge of the empirical literature is to keep the theoretical models of liquidity in mind and, at the same time, to do justice to the complex nature of real-world markets. My research provides important insights for investors in limit order book markets and has implications for the further theoretical and empirical research.

I address liquidity risk empirically with the help of three months of limit order book data from a the Frankfurt Stock Exchange. On the basis of Foucault, Kadan and Kandel (2005) I set up an econometric model of resiliency which I estimate for the dynamics of the spread and depth. Then I build on Sandås (2001) to implement the price-quantity schedule of the limit order book empirically and estimate systematic liquidity movements across stocks. Finally, I use the Fama and French (1992) asset pricing framework to assess the impact of liquidity factors on stock returns.

With regard to resiliency (Chapter 4), I find strong evidence that the limit order book follows a stable replenishment process. If an order book is empty it attracts a lot of new liquidity which increases the book's depth. Likewise, if the current bid-ask spread is high, new limit orders get submitted that improve the price. These results reflect that the limit order book is resilient, because once liquidity has been consumed, it gets replaced by new liquidity. Along time, spread resiliency is negatively associated with the order arrival rate and positively associated with the proportion of patient traders in the trading population. This suggests that liquidity-consuming factors reduce resiliency and liquidity-providing factors increase resiliency. In general, spread resiliency and depth resiliency react contrarily because spread resiliency erodes depth at the best tick, so that the two effects are asynchronous. In the cross section, spread and depth resiliency are both positively associated with market capitalization and beta, yet negatively correlated with insider profits. Resiliency is not significantly correlated with either the bid-ask spread or depth. This reinforces, ex post, the importance of resiliency as an independent dimension of liquidity which provides significant new information.

With regard to commonality (Chapter 5), I find strong evidence of market-wide liquidity movements. At first, I analyze commonality for liquidity at the best prices which turns out to be fairly weak (2% for depth and 6% for spreads). However, once I take the liquidity supply in the order book into account, commonality increases strongly: for depth at 2% price impacts commonality rises to 16%. Measures of the price-quantity schedule lead to similar results. Obviously, in a limit order market with a limit order book, the commonality of liquidity provision is drastically higher than the spread suggests. In addition to the mismeasurement of order book liquidity I also examine the time variation of liquidity co-movement: commonality is far stronger at the opening of the trading day when liquidity is low and at the end of the day when the liquidity supply falls again. During the day when liquidity is high commonality is low. Furthermore, commonality is much stronger in falling markets than in rising markets. In all, this is strong evidence that in the past the mismeasurement of liquidity and the neglect of time variation has led to a serious underestimation of liquidity commonalities.

With regard to pricing effects (Chapter 6), I find strong evidence that liquidity plays an important role for the pricing process of assets: both the level of liquidity and liquidity risk are incorporated into stock returns in a significant way. The lower the level of liquidity is, the higher the return that investors receive from buying the stock. Likewise, the higher the liquidity risk of a stock is, the higher is the return that an investor receives. This result can be observed both for the systematic component of liquidity risk as well as for the resiliency of the limit order book. The pricing effects for the level of liquidity and for the systematic component of liquidity risk are very similar: a difference of one standard deviation across stocks will lead, all other things held equal, to a return increase of about 10% return per annum. An increase in resiliency leads to a return increase of about 5%. Evidently, investors require higher premia if they hold stocks which cannot be traded as easily or whose liquidity is more risky.

7.2 Further Research

The previous results show that there are considerable differences in liquidity both in the time series as well as in the cross section of stocks. The most important implication for investors is that liquidity risk can have a substantial impact on their trading performance. Therefore it is not surprising that liquidity risk is reflected in stock prices and returns. From an academic point of view, the results bring up further interesting issues which I leave to future research:

1. Theoretical models of resiliency have only addressed the dynamics of the spread so far, most notably Foucault, Kadan and Kandel (2005). However, the empirical results show that the replenishment mechanism of the order book is strong both for the bid-ask spread and for depth. These results imply the necessity to extend the theory to a more realistic limit order book setting. In particular, theoretical models should allow the queuing of limit orders at the same tick and should give up the assumption that limit orders are always price-improving. Although these elements will make the models more complex, they should at the same time offer richer and more realistic predictions about the behavior of limit order book markets.

2. While theory is in search of an adequate framework for resiliency, some open questions remain for further empirical study. In particular, the previous results show that depth and price resiliency are both strong, yet their relationship is not totally clear. In the current set-up, the fixed time hori-

zon aggregates limit and market orders within the interval and therefore makes it difficult to identify clear lead and lag structures of the spread and depth. A promising approach is to use the complete history of limit order books updated for each single order. This would allow to study how the spread and depth adapt to liquidity shocks (market orders that clear one tick, two ticks and so forth) and how spread and depth changes interact.

3. In the context of a limit order book structure, liquidity commonality has not been modeled formally in the literature. However, the most striking result in my empirical study is how strongly systematic movements of liquidity increase with the tick distance of liquidity from the best price. Domowitz, Hansch and Wang (2005) provide further empirical evidence that in limit order book markets such commonality arises from the correlation of order types instead of order flow. Therfore, it appears very promising to extend the current models of the limit order book (such as Sandås (2001)) to a multi-asset setting. The great challenge is to endogenize the choice of order type which, in turn, should produce correlation in the liquidity supply.

4. A further intriguing aspect of commonality is the link between systematic liquidity movements and the state of the macroeconomy. As financial crises are mostly accompanied by vast liquidity outflows, a plausible hypothesis is that macroeconomic variables such as the movement of the exchange rate have predictive power for systematic liquidity risk. From the opposite per-

spective, the effect of liquidity commonality on the stability of the market is unclear. Chordia, Sarkar and Subrahmanyam (2005) give first evidence that spill-overs between bond and stock markets are large. It remains to be examined whether high commonality in liquidity is evidence of an efficient and well-functioning liquidity supply or a signal of an approaching crisis.

5. The pricing results demonstrate that both the level of liquidity as well as its riskiness are factors which have a pronounced impact on stock returns. A central problem of such asset pricing tests is the length of the times series. A promising extension to my approach would be to make liquidity a function of observable economic variables. For example, Beltran, Giot and Grammig (2005) identify such factors that predict the liquidity in the limit order book. In a first step their coefficients would have to be estimated for the limit order book data. In a second step, liquidity factors could then be constructed going back in time. The advantages of such an approach are that constant coefficients are a less restrictive assumption than constant liquidity factors and that the longer time series increase the statistical power of the asset pricing test.

In all, the results in the previous chapters contribute to the understanding of liquidity risk in limit order book markets. The thesis underlines the importance of liquidity risk for investors' trading decisions. The main challenge of future research is to bring theory in line with the empirical results and to extend the empirical literature in order to obtain further insights into the role of liquidity risk in limit order book markets.

Appendix A

Additional Tables

Table A.1: Descriptive Statistics of Spread and Depth Measures at the Best Limit Prices in the Limit Order Book

Stock	Market Capital.	Trading Volume	Absolute Spread	Relative Spread	Aggregate Depth
Adidas	4.17	34.76	0.086	0.093	94,230
Allianz	38.50	295.99	0.056	0.057	194,850
Altana	6.73	33.47	0.052	0.106	79,820
Basf	25.39	135.67	0.027	0.063	182,560
Bmw	23.08	97.23	0.021	0.091	156,100
Bayer	17.05	94.15	0.026	0.075	152,060
Commerzbank	9.36	56.29	0.017	0.108	157,800
Continental	4.09	27.82	0.038	0.119	88,490
Daimler	37.36	201.79	0.025	0.068	221,550
Deutsche Bank	38.34	324.95	0.036	0.053	252,420
Deutsche Boerse	4.87	36.55	0.046	0.098	109,200
Deutsche Post	18.23	45.42	0.020	0.111	144,860
Deutsche Telekom	61.29	348.60	0.012	0.074	979,210
Eon	36.15	178.47	0.069	0.104	76,630
Fresenius	3.94	14.13	0.071	0.131	65,770
Henkel	3.70	19.61	0.033	0.063	218,860
Infineon	8.01	153.05	0.012	0.106	361,810
Linde	5.15	24.03	0.047	0.107	82,390
Lufthansa	4.96	45.19	0.017	0.121	135,730
Man	3.42	28.34	0.035	0.128	89,570
Metro	11.42	42.63	0.043	0.123	96,300
Muenchener Rueck	22.13	221.43	0.058	0.062	201,930
RWE	16.49	106.93	0.028	0.083	165,730
SAP	42.49	201.83	0.080	0.061	171,560
Schering	7.91	53.69	0.037	0.092	106,230
Tui	2.95	26.53	0.028	0.149	89,860
Volkswagen	14.19	114.85	0.028	0.073	142,820
HypoVereinsbank	9.19	101.00	0.020	0.111	198,140
Thyssen Krupp	8.19	40.37	0.020	0.110	134,610
Siemens	57.16	339.02	0.030	0.050	263,520
Averages	18.20	114.793	0.037	0.093	180,487

Table A.1 reports descriptive statistics for different liquidity measures of our data set. The second column is market capitalization of each stock (as of 01 January 2004), the third colum is average daily trading volume (in 1,000 Euros), the fourth column is the absolute spread in Euros, the fifth column is the relative spread in % and the final column is depth at the best limit prices in Euros.

Table A.2: Descriptive Statistics for the Depth of the Order Book at 2% Price Impact and for the Slope of the Order Book

	Depth		Slopes	
Stock	Ask Side	Bid Side	Ask Side	Bid Side
Adidas	2,437.00	2,044.00	0.8522	0.9701
Allianz	8,833.00	8,657.00	0.2263	0.2332
Altana	1,776.00	1,602.00	1.2289	1.2673
BASF	3,924.00	4,360.00	0.4721	0.4463
Bayer	3,273.00	2,922.00	0.6139	0.6919
BMW	3,075.00	3,186.00	0.6809	0.6602
Commerzbank	2,885.00	2,295.00	0.7663	0.9186
Continental	1,239.00	1,036.00	1.6666	1.9632
Daimler	5,789.00	5,547.00	0.3326	0.3441
Deutsche Boerse	11,145.00	10,916.00	0.1836	0.1866
Deutsche Bank	1,830.00	1,861.00	1.1555	1.1305
Deutsche Post	2,411.00	2,060.00	0.8619	1.0021
Deutsche Telekom	17,655.00	16,668.00	0.1084	0.1133
Eon	1,173.00	997.00	1.6551	1.7904
Fresenius	823.00	831.00	2.2516	2.2441
Henkel	5,459.00	5,303.00	0.3422	0.3631
Infineon	7,150.00	5,614.00	0.3391	0.395
Linde	1,079.00	1,123.00	1.8555	1.8497
Lufthansa	1,734.00	1,621.00	1.1827	1.3244
MAN	1,046.00	945.00	2.0172	2.1075
Metro	1,485.00	1,429.00	1.4324	1.4314
Muenchener Rueck	7,815.00	7,048.00	0.2832	0.3002
RWE	2,770.00	2,705.00	0.7139	0.7267
SAP	6,423.00	6,915.00	0.2982	0.284
Schwering	2,200.00	2,467.00	1.0523	0.9295
TUI	1,223.00	1,102.00	1.7401	2.0498
VW	3,213.00	3,887.00	0.6709	0.5903
ThyssenKrupp	1,762.00	1,318.00	1.2102	1.4617
HypoVereinsbank	3,943.00	3,408.00	0.6082	0.7217
Siemens	11,007.00	11,107.00	0.1747	0.1752
Average	4,219.23	4,032.47	0.8992	0.9557

Table A.2 gives descriptive statistics of the liquidity of the order book. The second and third column give the aggregate depth (in 1,000.00 Euros) of the order book up to a price impact of 2%. The fourth and fifth column give parameter estimates of a linear model of the order book function.

Table A.3: PCA Results for the Spread and the Slope Measures

	Best Prices		Slope Measure	
PCA Output	Spread	Depth	Ask	Bid
First eigenvalue	2.23	2.17	6.85	5.42
Critical value	1.54	1.39	4.25	3.50
Proportion of variation	7.74	7.22	24.36	14.69

Table A.3 gives the results of PCA for liquidity at best prices (bid-ask spread and depth at the best bid and ask price) as well as the slope of the price-quantity schedule. The table lists the first eigenvalue, its critical values at the 95% confidence level and the proportion of total variability explained by the first principal component (in %).

Table A.4: Impact of the Time of Day: Slope of the Price-Quantity Schedule

	Open		Day		End	
Slope Coefficient	β_1	R^2	β_1	R^2	β_1	R^2
Ask Side	0.60	15.64	0.40	8.45	0.48	12.90
Bid Side	0.49	17.89	0.43	11.74	0.50	15.63

Table A.4 reports the average coefficients of market liquidity and the corresponding R^2 values of the regressions (in %) for the opening of the trading day ("open"), the midday trading period ("day") and the end of the trading day ("end"). The results are reported for the bid side and the ask side of the book. The liquidity measure used is the price-quantity schedule of the order book.

Table A.5: Impact of Market Momemtum: Slope of the Price-Quantity Schedule

	Up		Down	
Slope Coefficient	β^1	R^2	β^1	R^2
Ask Side	0.541	0.110	0.694	0.185
Bid Side	0.506	0.141	0.728	0.178

Table A.5 reports the average parameter estimates of the market liquidity parameter and R^2 values for rising markets ("up") and falling markets ("down"). It lists the results for the ask side and the bid side. The liquidity measure used is the price-quantity schedule of the book.

Table A.6: Relationship between Commonality and Market Return

Liquidity Measures	α	$t(\alpha)$	β	$t(\beta)$	R^2
Depth (Ask Side)	0.303	1.129	-1.125	-2.110	0.087
Depth (Bid Side)	0.227	0.794	-0.602	-3.853	0.010
Slope (Ask Side)	0.187	0.669	-1.004	-2.365	0.085
Slope (Bid Side)	0.194	0.890	-0.708	-3.112	0.099

Table A.6 reports the parameter estimates for the regression of commonality onto the market return (in Equation 5.5). It lists the results for the ask side and the bid side of the order book. The liquidity measures used are the depth in the limit order book and the slope of the price-quantity schedule.

Appendix B

Principal Component Analysis

Principal Component Analysis (PCA) is a technique that goes back to Pearson (1901) and Hotelling (1933). It is used to analyze variation of multivariate data structures. PCA decomposes the total variability within a data structure into components which account for the total variation. A textbook application of PCA to finance is given in Watsham and Parramore (1997).

The starting point of PCA is the covariance matrix. PCA requires stationary and standardized data because it will otherwise give too much weight to variables with large variances. Therefore if the data are of different orders of magnitude it is usual to standardize the variables first. Let X denote the data matrix and let x_i denote column i of this matrix. Each observation of each column is then standardized by subtracting the column mean $\overline{x}_i$and dividing by its standard deviation σ_{x_i}. Let z_i denote the columns of the standardized data and let Z be the corresponding data matrix:

$$z_{i,t} = \frac{x_{i,t} - \overline{x}_i}{\sigma_{x_i}} \tag{B.1}$$

The next step is then to compute the variance-covariance matrix of the standardized data matrix which I call C:

$$C = cov(Z) \tag{B.2}$$

The variance-covariance matrix enables us to extract from this matrix the eigenvectors and their associated eigenvalues. The eigenvectors are linearly independent combinations of the variables of the variance-covariance matrix which account for the total variance of the multivariate data structure. The eigenvalues give the proportion of risk that each eigenvector accounts for.

Mathematically, each eigenvector j is a vector γ_j which has an associated scalar λ_j known as the corresponding eigenvalue. The relationship between eigenvectors and eigenvalues is such that when the variance-covariance matrix is multiplied against the eigenvector, it simply scales the eigenvector by the eigenvalue:

$$C\ \gamma_j = \lambda_j\ \gamma_j \tag{B.3}$$

Now let matrix Λ be a diagonal matrix which contains all eigenvalues of C sorted from highest to lowest. Let Γ be the corresponding matrix of eigenvectors in which the eigenvectors are sorted as columns from highest to lowest. The symmetry and non-singularity of the variance-covariance matrix C then ensures that the inverse of Γ exists and equals Γ. Finally, PCA has decomposed the complete covariance matrix into eigenvectors and eigenvalues that account for the total variation in the original multivariate data:

$$C = \Gamma \Lambda \Gamma \tag{B.4}$$

The eigenvectors in Γ are the principal components that contribute to the total variance. The eigenvalues in Λ give the risk accounted for by each of these principal components. Dividing the first eigenvalue by the dimension of the variance-covariance matrix yields the proportion of risk explained by the first principal component. I perform PCA for our various liquidity measures and devote special attention to the proportion of total variation that the obtained principal components explain. I perform this analysis for standardized liquidity measures for which we compute the variance-covariance matrix and then follow the procedure above.

Appendix C

Order Book Reconstruction

The following source code reconstructs the Xetra order book at any desired frequency. It was programmed in Gauss 6.0 and consists of one main text file as well as several additional functions. The main file structures the computation procedure and the functions apply the Xetra matching rules in detail. The included file `globvariables.src` contains global variable definitions, the `helpprocs.src` file contains auxiliary functions and the additional file `mainprocs.src` contains a list of functions used in the main body of the program. Text that is enclosed by a slash and an asterisk – /* ... */ – is a comment and not executed. For the sake of brevity, the loop is shortened to three stocks that are sufficient to illustrate the program structure.

```
/* GENERAL PATH AND FORMAT SETTINGS */

chdir /orderbooks/input;
format /m1 /rd 5,2;

/* FILES AND SOURCE CODED TO BE INCLUDED */

#include /book/iput/globvariables.src;
#include /book/iput/mainprocs.src;
```

```
#include /book/iput/helpprocs.src;

/* LOOP FOR ALL 30 DAX EQUITIES (SHORTENED TO 3 STOCKS HERE):
DATA IS LOADED; ENTRIES IS PRE-SPECIFICATION OF MAXIMUM NUMBER
OF DIFFERENT LIMIT PRICES (REDUCES COMPUTATION; ERROR RETURNED
IF SPECIFICATION TOO LOW; DATA UP TO 3-MIO-BY-50 MATRICES */

stock=1;
do while stock <=30;

/* ADIDAS */

if stock == 1;
let data_all="allADS";
let data_init="initADS";
entries=300;
endif;

/* ALLIANZ */

if stock == 2;
let data_all="allALV";
let data_init="initALV";
entries=700;
endif;

/* here come all other stocks */


/* SIEMENS */

if stock == 30;
let data_all="allSIE";
let data_init="initSIE";
entries=400;
endif;

/* SAMPLING FREQUENCY IN SECONDS; UP TO 15 MINUTES SUPPORTED BY 1
GB PROCESSOR, HIGHER FREQUENCIES REQUIRE MORE POWERFUL SERVER */

sample_time=60*5;

/* COLUMN DEFINITIONS FOR FOLLOWING MATRIX MANIPULATION */
```

```
ordertype=1;
orderprice=2;
ordervolume=3;
direction=4;
ordertime=5;
orderdate=6;
weekday=7;
rem_ordervolume=8;
tradingperiod=9;
hid_ori_ordervolume = 10;
hid_rem_ordervolume = 11;
htb_ordervolume =12;
chge_rem_ordervolume = 13;

/* DAILY OPENING AND CLOSING HOURS (IN SEC) AND CURRENT PERIOD */

begin_sample=8.50*3600;
end_sample=17.5*3600;
number_period=(end_sample-begin_sample)/sample_time+1;

/* GENERATION OF VARIABLE NAMES */

vnames=seqa(1,1,250);
datanams=seqa(1,1,8);
vnames=0$+"v"$+ftocv(vnames,3,0);
datanams=0$+"v"$+ftocv(datanams,3,0);

/* DATA IMPORT FROM RAW DATA (FUNCTION "data_in_matrix");
"data_init" IS DATA FOR RECONSTRUCTION OF INITIAL ORDER
BOOK, "data_all" IS DATA FOR CONTINUOUS UPDATING */

temp_data_all=data_in_matrix(data_all);
temp_data_all=temp_data_all[.,1:14];
data=data_in_matrix(data_init);
data=data[.,1:14];

/* CALCULATION OF START DAY, END DAY AND NUMBER OF ORDER BOOKS */

dayvar=minc(temp_data_all[.,orderdate]);
endday=maxc(temp_data_all[.,orderdate]);
number_dates=rows(unique(temp_data_all[.,orderdate],1));
n_total=number_dates*number_period;
```

```
/* INITIALIZATION OF PRICE AND QUANTITY MATRICES THAT CONTAIN
ORDER BOOK INFORMATION (FOR VISIBLE AND HIDDEN BOOKS) AND
INFORMATION REGARDING TIME, TRADING PERIOD ETC. */

pask_total=zeros(n_total,entries);
pbid_total=zeros(n_total,entries);
vask_total=zeros(n_total,entries);
vbid_total=zeros(n_total,entries);
hid_vask_total=zeros(n_total,entries);
hid_vbid_total=zeros(n_total,entries);
stime_total=zeros(n_total,1);
info_total=zeros(n_total,3);
period_total=zeros(n_total,1);

/* POINTER FOR PLACE OF ORDER BOOK TO BE UPDATED; TEMPORARY PRICE
AND QUANTITY VARIABLES THAT ARE OVERWRITTEN EACH DAY */

indic=1;
mnanz_initial=1; /* INITIAL ASK POINTER POSITION IS 1 */
mnbnz_initial=1; /* INITIAL BID POINTER POSITION IS 1 */
mpask_initial=zeros(1,entries);
mpbid_initial=zeros(1,entries);
mvask_initial=zeros(1,entries);
mvbid_initial=zeros(1,entries);
hid_mvask_initial=zeros(1,entries);
hid_mvbid_initial=zeros(1,entries);

/* CONSTRUCTION OF INITIAL ORDER BOOK BY THE PROCEDURE
"construct_one_day1" */

construct_one_day1(mnbnz_initial,mnanz_initial,mpask_initial,
mpbid_initial,mvask_initial,mvbid_initial,entries,ordertype,
orderprice,ordervolume,direction,ordertime,rem_ordervolume,
hid_mvbid_initial,hid_mvask_initial,hid_ori_ordervolume,
hid_rem_ordervolume,htb_ordervolume,chge_rem_ordervolume,
tradingperiod);

/* SAVE PROCEDURE OUTPUT FOR FURTHER USE */

mpask_initial=mpask[rows(mpask),.];
mvask_initial=mvask[rows(mvask),.];
hid_mvask_initial=hid_mvask[rows(mvask),.];
```

```
mnanz_initial=mnanz[rows(mnanz),.];

mpbid_initial=mpbid[rows(mpbid),2:entries]~0;
mvbid_initial=mvbid[rows(mvbid),2:entries]~0;
hid_mvbid_initial=hid_mvbid[rows(mvbid),2:entries]~0;
mnbnz_initial=mnbnz[rows(mnbnz),.]-1;

/* CLEAR MEMORY TO REDUCE COMPUTATION TIME */

clear temp_data_init,mpbid,mvbid,mpask,mvask,mnbnz,mnanz,mtime,
merror_in_record;

/* LOOP FOR MAIN COMPUTATION: RECONSTRUCTION OF ORDER BOOKS DONE
DAILY AND RESULTS ARE SAVED TO FINAL OUTPUT MATRICES */

do while dayvar le endday;

data = selif(temp_data_all,temp_data_all[.,orderdate].==dayvar);
if not ismiss(data);
weekday = minc(data[.,weekday]);

/* CONSTRUCTION OF DAILY BOOK BY PROCEDURE "construct_one_day2" */

construct_one_day2(mnbnz_initial,mnanz_initial,mpask_initial,
mpbid_initial,mvask_initial,mvbid_initial,entries,ordertype,
orderprice,ordervolume,direction,ordertime,rem_ordervolume,
hid_mvbid_initial, hid_mvask_initial, hid_ori_ordervolume,
hid_rem_ordervolume,htb_ordervolume,chge_rem_ordervolume,
tradingperiod);

/* SAVE PROCEDURE OUTPUT FOR FURTHER USE */

lastobs=rows(mvask);
mvask_initial=mvask[lastobs,.];
mpask_initial=mpask[lastobs,.];
mvbid_initial=mvbid[lastobs,.];
mpbid_initial=mpbid[lastobs,.];
mnbnz_initial=mnbnz[lastobs,.];
mnanz_initial=mnanz[lastobs,.];
hid_mvask_initial=hid_mvask[lastobs,.];
hid_mvbid_initial=hid_mvbid[lastobs,.];
sampling(sample_time,begin_sample,9*3600-1);
topic=rows(pask);
```

```
/* SAVE PARTIAL ORDER BOOKS TO FINAL OUTPUT MATRICES */

pask_total[indic:indic+topic-1,.]=pask;
pbid_total[indic:indic+topic-1,.]=pbid;
vask_total[indic:indic+topic-1,.]=vask;
vbid_total[indic:indic+topic-1,.]=vbid;
hid_vask_total[indic:indic+topic-1,.]=hid_vask;
hid_vbid_total[indic:indic+topic-1,.]=hid_vbid;
info_total[indic:indic+topic-1,.]=(dayvar*ones(rows(info),1))...
~info~(weekday*ones(rows(info),1));
stime_total[indic:indic+topic-1,.]=stime;
period_total[indic:indic+topic-1,.]=period;

indic=indic+topic;
else;
endif;

dayvar=dayvar+1;
endo;

/* SAVE OUTPUT MATRICES */

if stock == 1;
save path = /book/oput/adidas/pask_total=pask_total;
save path = /book/oput/adidas/pbid_total=pbid_total;
save path = /book/oput/adidas/vask_total=vask_total;
save path = /book/oput/adidas/vbid_total=vbid_total;
save path = /book/oput/adidas/hid_vask_total=hid_vask_total;
save path = /book/oput/adidas/hid_vbid_total=hid_vbid_total;
save path = /book/oput/adidas/info_total=info_total;
save path = /book/oput/adidas/stime_total=stime_total;
save path = /book/oput/adidas/period_total=period_total;
endif;

if stock == 2;
save path = /book/oput/allianz/pask_total=pask_total;
save path = /book/oput/allianz/pbid_total=pbid_total;
save path = /book/oput/allianz/vask_total=vask_total;
save path = /book/oput/allianz/vbid_total=vbid_total;
save path = /book/oput/allianz/hid_vask_total=hid_vask_total;
save path = /book/oput/allianz/hid_vbid_total=hid_vbid_total;
save path = /book/oput/allianz/info_total=info_total;
```

```
save path = /book/oput/allianz/stime_total=stime_total;
save path = /book/oput/allianz/period_total=period_total;
endif;

if stock == 30;
save path = /book/oput/siemens/pask_total=pask_total;
save path = /book/oput/siemens/pbid_total=pbid_total;
save path = /book/oput/siemens/vask_total=vask_total;
save path = /book/oput/siemens/vbid_total=vbid_total;
save path = /book/oput/siemens/hid_vask_total=hid_vask_total;
save path = /book/oput/siemens/hid_vbid_total=hid_vbid_total;
save path = /book/oput/siemens/info_total=info_total;
save path = /book/oput/siemens/stime_total=stime_total;
save path = /book/oput/siemens/period_total=period_total;
endif;

" "; "Reconstruction (30min) completed for stock";;
stock;

stock=stock+1;
endo;
```

Bibliography

Acharya, V. and L. Pedersen, 2005, Asset Pricing with Liquidity Risk, Journal of Financial Economics 77, 375-410.

Ahn, H., K. Bae and K. Chan, 2001, Limit Orders, Depth and Volatility: Evidence from the Stock Exchange in Hong Kong, Journal of Finance 56, 767-788.

Allen, F. and D. Gale, Financial Contagion, Journal of Political Economy 108, 1-33.

Amihud, Y., 2002, Illiqudity and Stock Returns, Journal of Financial Markets 5, 31-56.

Amihud, Y., B. Christensen and H. Mendelson, 1992, Further Evidence on the Risk-Return Relationship, Research Paper Series 1248, Stanford University, Version November 1992.

Amihud, Y. and H. Mendelson, 1986, Asset Pricing and the Bid-Ask Spread, Journal of Financial Economics 17, 311-353.

Amihud, Y. and H. Mendelson, 1989, The Effects of Beta, Bid-Ask Spread, Residual Risk and Size on Stock Returns, Journal of Finance 44, 479-486.

Ang, A. and J. Chen, 2002, Asymmetric Correlations of Equity Portfolios, Journal of Financial Economics 63, 443-494.

Bae, K., H. Jang and S. Park, 2003, Traders' Choice between Limit and Market Orders: Evidence from NYSE Stocks, Journal of Financial Markets 6, 517-538.

Banz, R., 1981, The Relationship between Return and Market Value of Common Stocks, Journal of Financial Economics 9, 3-18.

Barclay, M., W. Christie, J. Harris, E. Kandel and P. Schultz, 1999, Effects of Market Reform on the Trading Costs and Depths of Nasdaq Stocks, Journal of Finance 54, 1-34.

Bauer, W., 2004, Commonality in Liquidity in Pure Order-Driven Markets, Working Paper, March 2004.

Bekaert, G. and G. Wu, 2000, Asymmetric Risk and Volatility in Equity Markets, Review of Financial Studies 13, 1-42.

Beltran, H., P. Giot and J. Grammig, 2005, Commonalities in the Order Book, Working Paper, Version 26th January 2005.

Beltran, H., J. Grammig and A. Menkveld, 2005, Understanding the Limit Order Book: Conditioning on Trade Informativeness, Working Paper, Version February 2005.

Beneish, M. and R. Whaley, 1996, An Anatomy of the S&P Game: The Effects of Changing the Rules, Journal of Finance 51, 1909-1930.

Benston, G., P. Irvine and E. Kandel, 2000, Liquidity beyond the Inside Spread: Measuring and Using Information in the Limit Order Book, Working Paper, Version July 2000.

Berk, J., 1995, A Critique of Size-Related Anomalies, Review of Financial Studies 8, 275-286.

Bessembinder, H., 1997, The Degree of Price Resolution and Equity Trading Costs, Journal of Financial Economics 45, 9-34.

Bessembinder, H., 2003, Trade Execution Costs and Market Quality after Decimalization, Journal of Financial and Quantitative Analysis 38, 747-777.

Biais, B., P. Hillion and C. Spatt, 1995, An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse, Journal of Finance 50, 1655-1689.

Black, F., 1993, Beta and Return: Anouncements of "Death" Seem Premature, Journal of Portfolio Management 20, 8-18.

Black, F., M. Jensen and M. Scholes, 1972, The Capital Asset Pricing Model: Some Empirical Tests, in: M. Jensen (ed.), Studies in the Theory of Capital Markets, New York.

Bollerslev, T., 1986, Generalized Autoregressive conditional heteroscedasticity, Journal of Econometrics 31, 307-327.

Bowsher, C., 2004, Modelling the Dynamics of Cross-Sectional Price Functions: An Econometric Analysis of the Bid and Ask Curves of an Automated Exchange, Working Paper, Version September 2004.

Brennan, M., Chorida, T. and A. Subrahmanyam, 1998, Alternative Factor Specifications, Security Characteristics and the Cross-section of Expected Stock Returns, Journal of Financial Economics 49, 345-373.

Brennan, M. and A. Subrahmanyam, 1996, Market microstructure and Asset Pricing: On the Compensation for Illiquidity in Stock Returns, Journal of Financial Economics 41, 441-464.

Brockman, P. and D. Chung, 2002, Commonality in Liquidity: Evidence from an Order-Driven Market Structure, Journal of Financial Research 25, 521-539.

Brown, S. and M. Weinstein, 1983, A New Approach to Test Asset Pricing Models: The Bilinear Paradigm, Journal of Finance 38, 711-743.

Brunnermeier, M. and L. Pedersen, 2005, Predatory Trading, Journal of Finance, 60, 1825-1863.

Cao, C., O. Hansch and X. Wang, 2004, The Informational Content of an Open Limit Order Book, EFA 2004 Maastricht Meetings Paper 4311, Version 15 March 2004.

Chakravarty, S., 2001, Stealth-trading: Which traders' trades move stock prices?, Journal of Financial Economics 61, 289-307.

Chalmers, J. and G. Kadlec, 1998, An Empirical Examination of the Amortized Spread, Journal of Financial Economics 48, 159-188.

Chan, L. and J. Lakonishok, 1995, The Behavior of Stock Prices around Institutional Trades, Journal of Finance 50, 1147-1174.

Chen, N. and R. Kan, 1996, Expected Return and the Bid-Ask Spread in K. Saitou and K. Kubota (editors): Modern Portfolio Theory and Applications, Osaka 1996.

Chordia, T., R. Roll and A. Subrahmanyam, 2000, Commonality in Liquidity, Journal of Financial Economics 56, 3-28.

Chordia, T., R. Roll and A. Subrahmanyam, 2001, Market Liquidity and Trading Activity, Journal of Finance 56, 501-530.

Chordia, T., R. Roll and A. Subrahmanyam, 2005, Evidence on the Speed of Convergence to Market Efficiency, Journal of Financial Economics 76, 271-292.

Chordia, T., A. Sarkar and A. Subrahmanyam, 2005, An Empirical Analysis of Stock and Bond Market Liquidity, Review of Financial Studies 18, 85-129.

Chordia, T., A. Subrahmanyam and V. Anshuman, 2001, Trading Activity and Expected Stock Returns, Journal of Financial Economics 59, 3-32.

Cochrane, J., 2001, Asset Pricing, Princeton (New Jersey).

Conrad, J., M. Gultekin and G. Kaul, 1991, Asymmetric Predictability of Conditional Variances, Review of Financial Studies 4, 597-622.

Constantinides, G., 1986, Capital Market Equilibrium with Transaction Costs, Journal of Political Economy 94, 842-862.

Coppejans, M., I. Domowitz and A. Madhavan, 2003, Resiliency in an Automated Auction, Working Paper, Version July 2003.

Coughenour, J. and M. Saad, 2004, Common Market Makers and Commonality in Liquidity, Journal of Financial Economics 73, 37-70.

Datar, V., N. Naik and R. Radcliffe, 1998, Liquidity and Stock Returns: An Alternative Test, Journal of Financial Markets 1, 203-219.

Degryse, H., F. Jong, M. Ravenswaaij and G. Wuyts, 2005, Aggressive Orders and the Resiliency of a Limit Order Market, Review of Finance 9, 201-242.

Deutsche Boerse Group, 2004, Xetra Release 7.1, Market Model Stock Trading, Version Sep 2004.

Demsetz, H., 1968, The Cost of Transacting, Quarterly Journal of Economics 82, 33-53.

Diebold, F. and C. Li, 2002, Forecasting the Term Structure of Government Bond Yields, Working Paper, Version October 2003.

Dhillon, U. and H. Johnson, 1991, Changes in the Standard and Poor's 500 List, Journal of Business 64, 75-85.

Domowitz, I., O. Hansch and X. Wang, 2005, Liquidity Commonality and Return Co-Movement, Journal of Financial Markets 8, 351-376.

Dong, Y., A. Kempf and P. Yadav, A Study on the Resiliency of Prices, Working Paper, Centre for Financial Research Cologne, Version November 2005.

Easley, D., S. Hvidkjaer and M. O'Hara, 2002, Is Information Risk a Determinant of Asset Returns?, Journal of Finance 57, 2185-2221.

Easley, D., N. Kiefer and M. O'Hara, 1997, One Day in the Life of a very Common Stock, Review of Financial Studies 10, 805-835.

Easley, D., N. Kiefer, M. O'Hara and J. Paperman, 1996, Liquidity, Information and Infrequently Traded Stocks, Journal of Finance 51, 1405-1436.

Easley, D. and M. O'Hara, 1987, Price, Trade Size and Information in Securities Markets, Journal of Financial Economics 19, 69-90.

Easley, D. and M. O'Hara, 1992, Time and the Process of Security Price Adjustment, Journal of Finance 47, 577-605.

Elesvarapu, V., 1997, Cost of Transacting and Expected Returns in the NASDAQ Market, Journal of Finance 52, 2113-2127.

Elesvarapu, V. and M. Reinganum, 1993, The Seasonal Behavior of the Liquidity Premium in Asset Pricing, Journal of Financial Economics 34, 373-386.

Fama, E. and K. French, 1992, The Cross-Section of Expected Stock Returns, Journal of Finance 47, 427-465.

Fama, E. and J. MacBeth, 1973, Risk, Return and Equilibrium: Empirical Tests, Journal of Political Economy 81, 607-636.

Fernando, C., 2003, Commonality in Liquidity: Transmission of Liquidity Shocks across Investors and Securities, Journal of Financial Intermediation 12, 233-254.

Foster, F. and S. Viswanathan, 1990, A Theory of the Interday Variations in Volume, Variance, and Trading Costs in Securities Markets, Review of Financial Studies 3, 593-624.

Foucault, T., 1999, Order Flow Composition and Trading Costs in a Dynamic Limit Order Market, Journal of Financial Markets 2, 99-134.

Foucault, T., O. Kadan and E. Kandel, 2005, Limit Order Book as a Market for Liquidity, Review of Financial Studies 18, 1171-1217.

Frey, S. and J. Grammig, 2005, Liquidity Supply and Adverse Selection in a Pure Limit Order Book Market, Working Paper, Version February 2005.

Fujimoto, A. and M. Watanabe, 2005, Time-Varying Liquidity Risk and the Cross Section of Stock Returns, Working Paper, Version August 2005.

Garbade, K., 1982, Securities Markets, New York.

Gârleanu, N. and L. Pedersen, 2004, Adverse Selection and the Required Return, Review of Financial Studies 17, 643-665.

Garman, M., 1976, Market Microstructure, Journal of Financial Economics 3, 257-275.

George, T, G. Kaul and M. Nimelandran, 1991, Estimation of the Bid-Ask Spread and its Components: A New Approach, Review of Financial Studies 4, 623-656.

Gibson, R. and N. Mougeot, 2004, The Pricing of Systematic Liquidity Risk: Empirical Evidence from the US Stock Market, Journal of Banking and Finance 28, 157-178.

Glosten, L., 1994, Is the Electronic Limit Order Book Inevitable?, Journal of Finance 49, 1127-1161.

Glosten, L. and L. Harris, 1998, Estimating the Components of the Bid-Ask Spread, Journal of Financial Economics 21, 123-142.

Glosten, L. and P. Milgrom, 1985, Bid, Ask and Transaction Prices in Specialist Market with Heterogenously Informed Traders, Journal of Financial Economics 13, 71-100.

Goettler, R., C. Parlour and U. Rajan, 2004, Equilibrium in a Dynamic Limit Order Book Market, Journal of Finance 60, 2149-2192.

Gomber, P., U. Schweickert and E. Theissen, 2004, Zooming in on Liquidity, EFA 2004 Maastricht Meetings Paper No. 1805, Version February 2004.

Grammig, J., A. Heinen and E. Rengifo, 2004, Trading Activity and Liquidity Supply in a Pure Limit Order Book Market – An Empirical Analysis using a Multivariate Count Data Model, Working Paper, Version August 2004.

Grammig, J. and E. Theissen, 2005, Is Best Really Better?, Internalization in Xetra BEST, Working Paper, Version July 2005.

Griese, K. and A. Kempf, 2004, Predicting Liquidity from Order Book Data, Working Paper, Version October 2004.

Gromb, D. and D. Vayanos, 2002, Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs, Journal of Financial Economics 66, 361-407.

Handa, P. and R. Schwartz, 1996, Limit Order Trading, Journal of Finance 51, 1835-1861.

Handa, P., R. Schwartz and A. Tiwari, 2003, Quote Setting and Price Formation in an Order-Driven Market, Journal of Financial Markets 6, 461-489.

Halka, D. and G. Huberman, 2001, Systematic Liquidity, Journal of Financial Research 24, 161-178.

Hansch, O., 2001, Island Tides: Exploring ECN Liquidity, Working Paper, Version January 2003.

Hansch, O., N. Naik and S. Viswanathan, 1999, Preferencing, Internalization, Best Execution, and Dealer Profits, Journal of Finance 54, 1799-1828.

Harris, L., 1990, Liquidity, Trading Rules and Electronic Trading Systems, New York Salomon Center (ed.), Monograph in Finance and Economics, 1990-4, New York.

Harris, L., 1990, Trading and Exchanges, Oxford University Press, New York.

Harris, L. and E. Gurel, 1986, Price and Volume Effects Associated with Changes in the S&P 500 List: New Evidence for the Existence of Price Pressures, Journal of Finance 41, 815-829.

Harrison, J. and D. Kreps, 1979, Martingales and Arbitrage in Multiperiod Securities Markets, Journal of Economic Theory 20, 381-408.

Hasbrouck, J., 1991, The Summary Informativeness of Stock Trades: An Econometric Analysis, Review of Financial Studies 4, 571-595.

Hasbrouck, J., 2004, Empirical Market Microstructure: Economic and Statistical Perspectives on the Dynamics of Trade in Security Markets, Teaching Notes, Version January 2004.

Hasbrouck, J. and D. Seppi, 2001, Common Factors in Prices, Order Flows and Liquidity, Journal of Financial Economics 59, 383-411.

Heaton, J. and D. Lucas, 1996, Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing, Journal of Political Economy 104, 443-487.

Ho, T. and H. Stoll, 1981, Optimal Dealer Pricing under Transactions and Return Uncertainty, Journal of Financial Economics 9, 47-73.

Ho, T. and H. Stoll, 1983, The Dynamics of Dealer Markets under Competition, Journal of Finance 40, 21-42.

Holthausen, R., R. Leftwich and D. Mayers, 1987, The Effect of Large Block Transactions on Security Prices: A Cross-Sectional Analysis, Journal of Financial Economics 19, 237-267.

Holthausen, R., R. Leftwich and D. Mayers, 1990, Large-block Transactions, the Speed of Response, and Temporary and Permanent Stock-Price Effects, Journal of Financial Economics 26, 71-95.

Hotelling, H., 1933, Analysis of a Complex of Statistical Variables into Principal Components, Journal of Educational Psychology 24, 417-441 and 498-520.

Huang, R. and H. Stoll, 1996, Dealer versus Auction Markets: A Paired Comparison of Execution Costs on NASDAQ and the NYSE, Journal of Financial Economics 41, 343-357.

Huang, R. and H. Stoll, 1997, The Components of the Bid-Ask Spread: A General Approach, Review of Financial Studies 10, 995-1034.

Jain, P. and G. Joh, 1988, The Dependence between Hourly Price and Trading Volume, Journal of Financial and Quantitative Analysis 23, 269-283.

Jones, C., 2002, A Century of Market Liquidity and Trading Costs, Working Paper, Version May 2002.

Kaul, A., V. Mehrotra and R. Morck, 2000, Demand Curves for Stocks *Do* Slope Down: New Evidence from an Index Weights Adjustment, Journal of Finance 55, 893-829.

Kavajecz, K., 1999, A Specialist's Quoted Depth and the Limit Order Book, Journal of Finance 54, 747-771.

Keim, D. and A. Madhavan, 1995, Anatomy of the Trading Process: Empirical Evidence on the Behavior of Institutional Traders, Journal of Financial Economics 37, 371-398.

Keim, D. and A. Madhavan, 1997, Transaction Costs and Investment Style: an Inter-Exchange Analysis of Institutional Equity Trades, Journal of Financial Economics 46, 265-292.

Kempf, A. and O. Korn, 1997, Market Depth and Order Size, Journal of Financial Markets 2, 29-48.

Kothari, S., J. Shanken and R. Sloan, Another Look at the Cross-Section of Expected Stock Returns, Journal of Finance 50, 185-224.

Kroner, K. and V. Ng, 1998, Modeling Asymmetric Comovements of Asset Returns, Review of Financial Studies 11, 817-844.

Kumar, S., 2003, Liquidity: Persistence and Forecast in Open Electronic Limit Order Book Market, Working Paper, Version July 2003.

Kyle, A., 1985, Continuous Auctions and Insider Trading, Econometrica 53,1315-1336.

Kyle, A. and W. Xiong, 2001, Contagion as a Wealth Effect, Journal of Finance 56, 1401-1440.

Lamoureux, C. and W. Lastrapes, 1990, Heteroscedasticity in Stock Return Data: Volume versus GARCH Effects, Journal of Finance 45, 221-229.

Large, J., 2005, Measuring the Resiliency of an Electronic Limit Order Book, Working Paper, 26 August 2005 (forthcoming in Journal of Financial Markets).

Litzenberger, R. and K. Ramaswamy, 1979, The Effect of Personal Taxes and Dividends on Capital Asset Prices: Theory and Empirical Evidence, Journal of Financial Economics 7, 163-196.

Longin, F. and B. Solnik, 2001, Extreme Correlation of International Equity Markets, Journal of Finance 56, 649-676.

Loughran, T., 1997, Book-to-Market across Firm Size, Exchange and Seasonality: Is there an effect?, Journal of Financial and Quantitative Analysis 32, 249-268.

Lowenstein, R., 2000, When Genius Failed: The Rise and Fall of Long-Term Capital Management, New York.

Lütkepohl, H., 2005, New Introduction to Multiple Time Series Analysis, Berlin.

Madhavan, A. and M. Cheng, 1997, McInish, In Search of Liquidity: Block Trades in the Upstairs and Downstairs Markets, Review of Financial Studies 10, 175-203.

Madhavan, A. and G. Sofianos, 1997, An Empirical Analysis of NYSE Specialist Trading, Journal of Financial Economics 48, 189-210.

McInish, T. and R. Wood, 1992, An Analysis of Intraday Patterns in Bid/Ask Spreads for NYSE Stocks, Journal of Finance 47, 753-764.

Miller, M. and M. Scholes, 1982, Dividends and Taxes: Some Empirical Evidence, Journal of Political Economy 90, 1118-1141.

Naik, N. and P. Yadav, 2003, Trading Costs of Public Investors with Obligatory and Voluntary Market-Making: Evidence from Market Reforms, EFA 2003 Annual Conference Paper 408.

Newey, W. and K. West, 1987, A Simple Positive Semi-Definite, Heteroscedasticity and Autocorrelation Consistent Covariance Matrix, Econometrica 55, 703-708.

Novy-Marx, R., 2003, On the Excess Returns to Illiquidity, Working Paper, Version April 2004.

New York Stock Exchange, 2000, Fact Book Online, `http://www.nysedata.com/factbook/`.

O'Hara, M., 1995, Market Microstructure Theory, Cambridge (Massachusetts).

O'Hara, M., 2003, Presidential Address: Liquidity and Price Discovery, Journal of Finance 58, 1335-1354.

Parlour, C., 1998, Price Dynamics in Limit Order Markets, Review of Financial Studies 11, 789-816.

Parlour, C. and D. Seppi, 2003, Liquidity-Based Competition for Order Flow, Review of Financial Studies 16, 301-343.

Pástor, L., and R. Stambaugh, 2003, Liquidity Risk and Expected Stock Returns, Journal of Political Economy 111, 642-685.

Pearson, K., 1901, On Lines and Planes of Closest Fit to Systems of Points in Space, Philosophical Magazine 6, 559-572.

Ranaldo, A., 2002, Order Aggressiveness in Limit Order Book Markets, Journal of Financial Markets 7, 53-74.

Rock, K., 1996, The Specialist's Order Book and Price Anomalies, Working Paper, Version January 1996.

Roll, R., 1977, A Critique of the Asset Pricing Theory's Tests, Journal of Financial Economics 4, 1073-1103.

Roll, R., 1984, A Simple Implicit Measure of the Effective Bid-Ask Spread, Journal of Finance 51, 613-636.

Ross, S, 1976, The Arbitrage Theory of Capital Asset Pricing, Journal of Economic Theory 13, 341-360.

Ross, S, 1978, A Simple Approach to the Valuation of Risky Streams, Journal of Business 51, 453-475.

Saar, G., 2001, Price Impact Asymmetry of Block Trades: An Institutional Trading Explanation, Review of Financial Studies 14, 1153-1181.

Sandås, P., 2001, Adverse Selection and Competitive Market Making: Empirical Evidence from a Limit Order Market, Review of Financial Studies 14, 705-734.

Seppi, D., 1997, Liquidity Provision with Limit Orders and a Strategic Specialist, Review of Financial Studies 10, 103-150.

Shanken, J., 1992, On the Estimation of Beta-Pricing Models, Review of Financial Studies 5, 1-33.

Shleifer, A., 1986, Do Demand Curves for Stocks Slope Down?, Journal of Finance 41, 579-590.

Stoll, H., 1978, The Supply of Dealer Services in Securities Markets, Journal of Finance 33, 1133-1151.

Stoll, H., 1989, Inferring the Components of the Bid-Ask-Spread: Theory and Empirical Tests, Journal of Finance 44, 115-134.

Vayanos, D., 1998, Transaction Costs and Asset Prices: A Dynamic Equilibrium Model, Review of Financial Studies 11, 1-58.

Watanabe, M., 2003, A Model of Stochatic Liquidity, Working Paper, Version June 2003.

Watsham, T. and K. Parramore, 1997, Quantitative Methods in Finance, First Edition, Thomson Learning, London.

Wood, R., T. McInish and J. Ord, 1985, An Investigation of Transaction Data for NYSE Stocks, Journal of Finance 25, 723-739.