

Order Book Characteristics and the Volume-Volatility Relation: Empirical Evidence from a Limit Order Market*

Randi Næs

*Norges Bank, Bankplassen 2, 0107 Oslo, Norway and
Norwegian School of Economics and Business Administration (NHH)*

and

Johannes A. Skjeltorp †

*Norges Bank, Bankplassen 2, 0107 Oslo, Norway and
Norwegian School of Management (BI)*

May 24, 2005

Abstract

Using unique data, we address the issue of price formation in a limit order market. A standard volume-volatility relation is documented with the number of trades acting as the important component of volume. The main contribution of the paper is to identify strong evidence that volume, volatility, and the volume-volatility relation are negatively related to the order book slope. These results are robust to the inclusion of several liquidity measures. A significant empirical relationship between the order book slope and the coefficient of variation in earnings forecasts by financial analysts suggests that the slope is proxying for disagreement among investors. Hence, our results support models where investor heterogeneity intensifies the volume-volatility relation.

Keywords: Market Microstructure, Equity Trading, Asymmetric Information

JEL Codes: G10, G20

*The views expressed are those of the authors and should not be interpreted as reflecting those of Norges Bank. We are grateful to Bernt Arne Ødegaard, Dagfinn Rime, Thierry Foucault, Richard Lyons, Soren Hvidkjaer, Loran Chollete, and participants at the AFFI 2003 Conference in Lyon, the EFA 2004 Conference in Maastricht, and the Skinance 2005 Conference in Hemsedal for providing useful comments and suggestions. We would also like to thank Sverre Lilleng and Thomas Borchgrevink at the Oslo Stock Exchange for providing us with the transactions data and for being very helpful in the data preparation process. All remaining errors or omissions are ours.

†Corresponding author. E-mail:johannes-a.skjeltorp@norges-bank.no Phone:(+47)22316740
Fax:(+47)22424062

Introduction

In this paper, we study the informativeness of the order book in an electronic limit order market for equities. The importance of electronic trading systems is growing. Outside the US, electronic order matching is today the typical way to arrange equity trading.¹ Through the development of electronic communication networks (ECNs), traditional US marketplaces like the NYSE and the NASDAQ are also being challenged by electronic trading methods.² Parallel to the increased popularity of limit order markets, there is a growing academic interest in the properties of this market structure. One important property is a high degree of transparency. A limit order book is more informative about aggregate liquidity supply and trading interests than quoted volumes and prices from a dealer. The transparency of a limit order book is illustrated in Figure 1. The figure depicts snapshots of the total and visible order books of the two largest and most liquid securities in our data.³ Figure (a) shows the order book of a traditional blue chip firm, Norsk Hydro (NHY), and Figure (b) shows the order book of a relatively young IT firm, Opticom (OPC). The figure reveals a substantial difference in the liquidity provision for the two stocks. While on average about 50 percent of the order volume for NHY has limit prices which lie within 5 ticks from the quoted spread, the similar percentage for OPC is only about 10 percent.

Overall, we find that there are systematic variations in the order book slope across securities and over time. The main contribution of the paper is to identify strong evidence that the order book slope is negatively related to both trading volume and price volatility, as well as to the correlation between volume and volatility (“the volume-volatility relation”). This result is robust to the inclusion of several liquidity measures.

The volume-volatility relation is a well documented empirical fact found for most types of financial contracts, including stocks, Treasury bills, currencies and various futures contracts. The main theoretical explanation for the relation is that the arrival of new information makes prices adjust to new equilibria over time. A complementary explanation is that prices also change because investors have dispersed beliefs about asset values. As both explanations are hard to test, our result that the order book slope contains information about the phenomenon is interesting. To investigate further the relationship between the order book slope and the volume-volatility relation, we study factors that interact with the order book slope. A significant

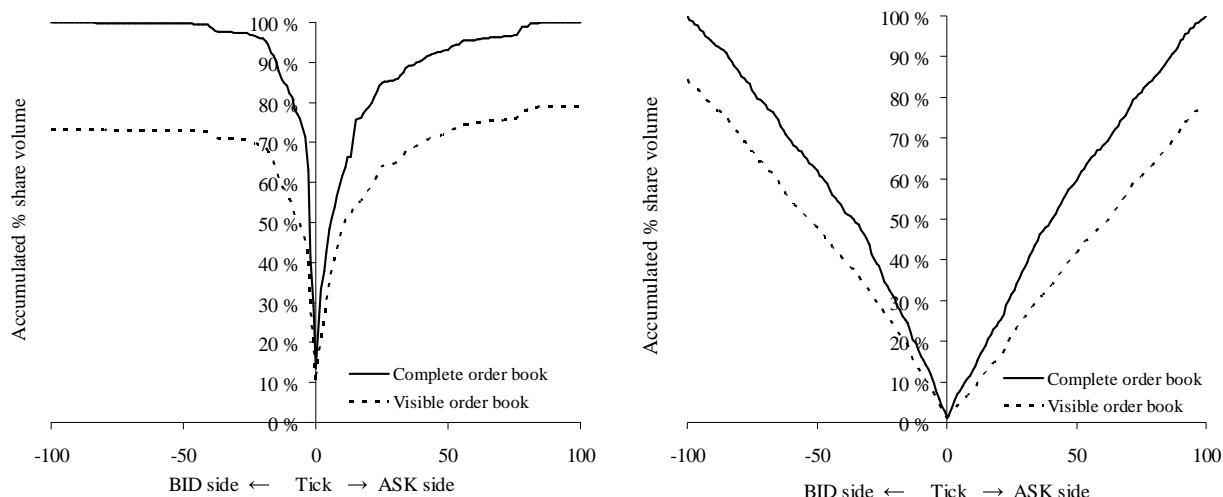
¹Limit order-driven stock exchanges are found in, for example, Toronto, Tel Aviv, Paris, Frankfurt, Stockholm, and Oslo.

²According to Bloomfield et al. (2003), ECN’s “such as Island, Instinet, and Archipelago use an electronic order book structure to trade as much as 45 % of the volume on NASDAQ”.

³The total order book includes the hidden order volume.

Figure 1: Transparency in a limit order market

The figure shows the total and visible order books of two large Norwegian firms, averaged over the last five trading days of May 2001. The upper picture is the order book of Norsk Hydro (NHY), a traditional blue chip firm, while the lower picture is the order book of Opticom, a relatively new IT firm. The vertical axis shows the cumulative percentage share volume in the book, and the horizontal axis shows ticks away from the quotes. Zero represents the best quote on each side of the market.



negative relationship is found between the order book slope and the coefficient of variation in analysts' earnings forecasts, i.e. the greater the disagreement among analysts, the more gentle the average slope of the order book. This result suggests that the order book slope proxies for dispersed beliefs about asset values. Applied to the two securities in Figure 1, this would imply that investors disagree more about the value of OPC than NHY. Looking at the main characteristics of the two firms, this makes sense. NHY is a leading energy, aluminum and fertilizer firm, with 50,000 employees in 60 countries worldwide, with well known operations, a long history, and a large amount of available information, including experts' analysis. OPC, on the other hand, is a relatively young IT firm with under 100 employees, and very uncertain future income prospects. However, our results could also be explained within a Glosten (1994) model with homogeneous liquidity suppliers. In the Glosten model, different order book shapes arise because liquidity providers condition their limit orders on the probability of informed market orders. Hence, for a given level of liquidity-motivated trading and a given probability of informed trading, the slopes will be more gentle the more volatile assets are.⁴

Several empirical papers investigating order book data are related to our work. Biais et al.

⁴We are grateful to an anonymous referee for pointing this out to us.

(1995) analyze in detail the interaction between the order book and order flow on the Paris Bourse. One relevant finding is that the status of the order book is important for order flows and trading volume. While Biais et al. (1995) focus their analysis on the inner quotes of the order books, we ask whether variations in the liquidity provision from the full order book are relevant for price formation and trading activity. Kalay et al. (2003) estimate the demand and supply elasticities for stocks on the Tel Aviv Stock Exchange. Their main findings are that the order book is more elastic at the beginning of the day, and that the demand side is more elastic than the bid side.⁵ We find similar patterns in our data, however, our focus is on variations in a measure of the average slope of the order book on a daily frequency. There are also two papers that focus on the shape of the order book in the case of financial crises. Goldstein and Kavajecz (2004) provide evidence of a negative relation between the shape of the order book and volatility during a case of an extreme market movement. Kim et al. (2004) study the changes in estimated demand and supply elasticities for different investor types around several Asian financial crises. The demand and supply curves of domestic individuals are found to be more inelastic than the demand and supply curves of domestic institutions and foreign investors. Our focus is on systematic patterns between order book slope, trading volume, and volatility over a longer time period with varying trading conditions.

The paper is organized in the following way. Section 1 describes the data, including the method for estimating the order book slope. Section 2 presents the analysis of the relationship between order book slope, price volatility and trading activity. In section 3, we analyze what factors may explain the order book slope and discuss various interpretations of our findings. Section 4 concludes the paper.

1 The Data

1.1 The Norwegian Stock Market

The data are from the Oslo Stock Exchange (OSE) in Norway. Norway is a member of the European Economic Area, and its equity market is among the 30 largest world equity markets by market capitalization. The OSE is the only regulated marketplace for securities trading in Norway. Since January 1999, it has operated as a fully computerized centralized limit order book system similar to the limit order book systems in e.g. Paris, Toronto, Stockholm and Hong

⁵The first result is interpreted as supportive to sequential trading models with asymmetric information which predict higher adverse selection at the opening (Glosten and Milgrom, 1985). The second result is interpreted as supportive to the empirical finding that buy orders have larger price impacts than sell orders.

Kong. At the end of our sample period (June 2001), 212 securities were listed on the exchange with a total market value of about NOK 677 bill.⁶ The market has grown substantially during the last 10 years. Measured in real terms, the total market capitalization at the end of 2003 was more than double the value at the end of 1994. Another notable characteristic of the market is a very high concentration of values and trading activity in a few large firms.⁷

As is normal in most electronic order-driven markets, the order handling rule at the OSE follows a strict price-time priority. All orders are submitted at prices constrained by the minimum tick size for the respective securities which is determined by the price level of the security. The trading day comprises two sessions: the “pre-trade” session starting at 9.30am and ending with an opening auction at 10am, and the “continuous trading” session from 10am until the trading closes at 4pm. During the continuous trading session, electronic matching of orders with crossing or equal price generates transactions. Orders without a limit price (market orders) have automatic price priority and are immediately executed at the best available prices.

1.2 The data

The data consist of every order and trade that occurred at the OSE in the period from February 1999 through June 2001. Every trade is linked to the underlying orders through an order ID. Thus, if a large order is executed against many smaller orders resulting in several smaller trades, we can trace each executed part back to the initial order.

Many listed securities in the sample are infrequently traded. To avoid noise in intraday measures from illiquid securities, we filter the securities based on two criteria: (i) the security must have been traded in at least 400 out of the total 597 trading days in the sample, and (ii) the security must have an average of 5 trades per day. After applying these filters, we are left with 108 securities. Table 1 provides some descriptive statistics of the trading activity in these securities. The table shows numbers for the whole sample as well as for four sub-samples based on market cap. The firms are assigned to a market capitalization group based on their market capitalization value at the beginning of each year. As found in dealer markets, the average effective spreads are lower than the average quoted spreads. In limit order markets this difference cannot be explained by dealers providing price improvements. Instead, since quoted spreads are measured immediately after each trade, the difference is most likely due to the fact that each trade removes liquidity. The number of trades, the trading volume, the prices and the

⁶During the sample period, USD 1 was worth between NOK 8 and NOK 9.

⁷Notable Norwegian listings include Norsk Hydro, Telenor, and Statoil.

Table 1: Descriptive statistics of the trading activity

The table provides some descriptive statistics of trades for the whole sample, and for four market capitalization groups. Group 1 consists of the 25% smallest firms while group 4 consists of the 25% largest firms. Some firms have experienced large changes in capitalization value during the sample period. To take account of this, we re-sort the market capitalization groups at the beginning of each year. The quoted spread is calculated as a percent of the spread midpoint. The effective spread is calculated as the difference between the execution price and the spread midpoint (as a percentage of the spread midpoint) multiplied by two.

Aggregate statistics:	All firms	MCAP quartiles			
		1 (small)	2	3	4 (large)
Number of firms	108	27	27	27	27
Sum trades (in thousands)	3724	390	522	504	2309
Sum shares traded (mill.)	9585	1707	1922	919	5037
Sum NOK volume (bill.NOK)	648	21	44	68	516
Cross-sectional averages:					
Market cap (mill.NOK)	5259	354	938	2339	13978
Price	88.4	23.34	62.43	105.66	150.73
Daily volatility (%)	2.71 %	3.49 %	2.98 %	2.30 %	2.29 %
Shares traded (V) in thousands	151	116	171	78	288
Trades (N)	58	28	41	41	148
Trade size (AV) in shares	2890	4859	2684	1549	1912
Quoted spread (NOK)	1.65	0.94	1.63	2.11	1.57
Effective spread (NOK)	1.22	0.68	1.20	1.59	1.16
Quoted % spread (% of midpoint)	3.04 %	4.74 %	2.77 %	2.40 %	1.34 %
Effective spread (% of midpoint)	2.22 %	3.38 %	2.03 %	1.85 %	0.99 %

quoted spread increase with market cap, while the average daily volatility, the average trade size, and the quoted percentage spread decrease.

1.3 Rebuilding the order book

The order data include all order submissions, deletions and amendments of existing orders. We also know if an order has a hidden part, and whether it was a buy or a sell order. Thus, for each security in the data, we are able to reconstruct the full order book, with and without the hidden part, at any point in time.

To rebuild the order book, we start at the beginning of the trading day with the orders still remaining after the opening auction has been executed. Then we track all types of orders submitted throughout the day, and update the order book accordingly. Thus, all deletions and amendments of earlier orders, new orders and trade executions are accounted for when we update the book. To remove stale and erroneous orders, we exclude order book volumes above and below 100 ticks from the inner quotes. For a security trading at NOK 100 with a minimum tick size of NOK 0.5 this would mean that orders above NOK 150 and below NOK 50 are excluded from our calculations. The limit on 100 +/- ticks implies that we disregard less

Table 2: The distribution of order book volume

The table provides descriptive statistics on the distribution of the order book volume. The numbers are daily average fractions of accumulated volume, and are reported for all securities, for the bid and ask side separately, for minimum tick sizes, and for the four market capitalization groups.

	ATQ	+/- 1 tick	+/- 5 tick	+/- 10 tick	+/- 20 tick	+/-50 tick	+/-100 tick
All firms	20.9 %	34.7 %	56.8 %	69.4 %	78.4 %	88.6 %	100.0 %
Bid side	23.0 %	40.0 %	62.9 %	73.8 %	81.4 %	89.7 %	100.0 %
Ask side	20.8 %	29.3 %	50.8 %	64.9 %	75.5 %	87.4 %	100.0 %
<i>Minimum tick size</i>							
0.01	20.2 %	30.8 %	37.8 %	49.0 %	60.1 %	82.2 %	100.0 %
0.1	22.2 %	34.2 %	53.2 %	67.4 %	79.4 %	91.7 %	100.0 %
0.5	22.3 %	39.1 %	65.8 %	78.4 %	88.1 %	95.5 %	100.0 %
1	7.0 %	10.7 %	17.6 %	25.1 %	38.8 %	70.0 %	100.0 %
<i>Market capitalization quartiles</i>							
1 (small)	19.1 %	29.7 %	45.2 %	56.6 %	68.2 %	84.0 %	100.0 %
2	21.6 %	34.9 %	56.3 %	69.6 %	79.9 %	91.1 %	100.0 %
3	23.6 %	38.3 %	62.7 %	75.5 %	83.8 %	92.6 %	100.0 %
4 (large)	19.3 %	34.6 %	62.9 %	75.9 %	84.3 %	91.0 %	100.0 %

than 5 percent of our sample.

The descriptive statistics discussed below are based on 6 hourly spaced snapshots of the order books of each security each day, at 10.30am, 11.30am, 12.30pm, 1.30pm, 2.30pm and 3.30pm. The first snapshot is half an hour after the regular trading session starts. Alternatively, we could have ended the last snapshot at 4pm, but then the order book would have been affected by the large number of order cancellations at the end of the trading day.

In Table 2, we show the distribution of volume in the order book averaged over all securities and dates. At each tick level, the fraction of total shares in the order book is averaged over the 6 hourly order book snapshots. The table shows the order book distribution over minimum tick sizes and market capitalization quartiles.⁸ Around 35 percent of the order book depth is concentrated at the quotes or plus/minus one tick from the quotes. This is quite stable both across tick sizes and across market cap quartiles, although the smallest firms tend to have more dispersed order books. When we separate the bid and ask sides, we find that the volume on the bid side is more concentrated at the inner quotes than the volume on the ask side. This is in line with the findings in other studies, and is consistent with the interpretation that the price impact is larger for buy orders than for sell orders.⁹ There do not seem to be large differences in order book depth across market capitalization quartiles. The largest tick size category is

⁸If a security trades across two minimum tick size regimes on the same day, we remove the security from the sample on this day. Our results do not change if we include these observations.

⁹See Chan and Lakonishok (1993), Chan and Lakonishok (1995), and Kalay et al. (2003)

special in that it only contains one, highly volatile and very actively traded firm, which during the sample period was traded in both the NOK 0.5 and the NOK 1 tick price range. The firm is Opticom, cf. Figure 1.

Hidden orders

When an order is submitted as a hidden order, only a specified fraction of the order is visible to the market. When the visible part of a hidden order is executed, it is replaced by a new part of the hidden order at the same price and of equal size, but with loss of time priority. The statistics presented in Table 2 include the hidden orders. In Table 3, we provide some statistics when orders are split into visible and hidden orders.

Table 3: Descriptive statistics of the hidden orders

The table shows some descriptive statistics for the hidden orders over the whole sample period, half-year periods and market capitalization (MCAP) quartiles. The first section shows the total number of visible orders and hidden orders, and the average order size for the two types of orders. All order sizes are reported in thousands of shares. For the hidden orders, we also calculate the average order size for the hidden (undisclosed) part and visible (disclosed) part of the orders. The second section shows hidden orders as a fraction of the total number of orders and the total share volume in all orders. The last section shows the volume of executed trades where at least one side of the trade comes from an order with hidden volume. It also shows the fraction of total trading volume where a part of a hidden order is involved.

Regular and hidden orders	Full sample	Subperiods (half-years)					MCAP quartiles			
		1999.1	1999.2	2000.1	2000.2	2001.1	1	2	3	4
<i>Number of orders:</i>										
regular (thousands)	5429	536	853	1362	1353	1325	636	827	768	3197
hidden (thousands)	291	26	40	68	78	79	24	44	48	174
<i>Order size:</i>										
regular (1000 shares)	5	7	7	4	4	4	8	7	3	4
hidden (1000 shares):	22	30	31	20	19	18	43	30	16	18
- hidden part	16	22	22	15	14	13	32	22	11	13
- visible part	6	9	9	5	5	5	10	8	4	5
Fraction which is hidden										
Total number of orders	0.05	0.05	0.04	0.05	0.05	0.06	0.04	0.05	0.06	0.05
Total share volume in all orders	0.19	0.17	0.17	0.19	0.19	0.23	0.16	0.20	0.24	0.19
Trades involving hidden orders										
<i>Trade volume:</i>										
Number of trades (thousands)	951	63	107	231	269	282	81	133	136	601
Shares volume (mill. shares)	2436	262	473	528	591	583	367	464	268	1336
NOK volume (bill. NOK)	188	15	28	44	50	51	5	13	22	148
<i>Fraction of trade volume:</i>										
number of trades	0.26	0.19	0.20	0.24	0.28	0.30	0.21	0.25	0.27	0.26
share volume of trades	0.25	0.20	0.21	0.26	0.29	0.32	0.21	0.24	0.29	0.27
NOK volume of trades	0.29	0.22	0.21	0.29	0.33	0.35	0.23	0.30	0.32	0.29

Hidden orders are about 4 times larger than visible orders. The visible part of the hidden orders is also slightly larger than visible orders. Apart from that, hidden orders are quite similar to visible orders: the total number of orders increases with market cap, order size decreases with market cap, and average order size has declined during the sample period. A similar pattern is also evident across market cap quartiles and over time. 5 % of all submitted orders (whole sample) has a hidden part and about 19 % of the share volume of all submitted orders is hidden. Both fractions have increased over time, and seem to be increasing with market cap. With respect to executed trades, 951 thousand trades (of a total of 3724 thousand trades) involved at least one order (one side of the trade) with a hidden volume. This constitutes about 26 % of all trades for the whole sample period. The fraction is highest for the largest firms and has increased by 10 percentage points during the sample period. Overall, hidden orders are shown to be an important component of total market liquidity.

1.4 Measuring the slope of the order book

We estimate the average daily slope of the order book based on the snapshots of the book at the end of the six time intervals described in section 1. Essentially, what we want to measure is the elasticity $\partial q/\partial p$ describing how quantity (q) supplied in the order book changes as a function of the price (p). These elasticities will generally change along the demand and supply curves, and we want a measure that describes the average elasticity across all price levels with positive volumes in the order book. To obtain an average measure, our calculation is carried out in the following steps:

1. Firstly, for each side of the order book, and each snapshot, we accumulate the aggregate number of shares supplied/demanded at each price level. This leaves us with the total volume supplied (demanded) at that price or lower (higher).
2. Secondly, we calculate the local slope at each price level (explained in more detail below).
3. Thirdly, we average the local slopes across all price levels for the bid and ask side separately to obtain an average slope for the bid and ask side for the respective snapshot.
4. Finally, we take the average of the bid and ask slope to get one slope measure for the respective snapshot, and average across the 6 snapshots to obtain one daily average slope for each security.¹⁰

¹⁰Having one daily average slope overcomes the problem that one snapshot (e.g. at noon) is biased, either due

More formally, let N_A be the total number of ask prices (tick levels) containing orders and let τ denote the tick level where $\tau = 0$ represents the bid-ask midpoint and $\tau = 1$ represents the best ask quote (inner tick) with a positive share volume. Furthermore, let p_0^A denote the bid-ask midpoint and let v_τ^A be the natural logarithm of accumulated total share volume at each tick level τ . Since our calculation of the average order book slope is similar for the bid and ask side, we only describe the ask side. The average slope for the ask side, SE, on day t for each snapshot $s \in [1..6]$ for security i is calculated as

$$SE_{i,t}^s = \frac{1}{N_A} \left\{ \frac{v_1^A}{p_1^A/p_0^A - 1} + \sum_{\tau=1}^{N_A} \frac{v_{\tau+1}^A/v_\tau^A - 1}{p_{\tau+1}^A/p_\tau^A - 1} \right\} \quad (1)$$

Finally, the average daily slope of the order book for security i on date t is calculated as

$$SLOPE_{i,t} = \frac{1}{6} \sum_{s=1}^6 \left\{ \frac{SE_{i,t}^s + DE_{i,t}^s}{2} \right\} \quad (2)$$

where $DE_{i,t}^s$ is the average slope for the bid side. The more gentle (steeper) the slope, the more widely distributed (concentrated) the volumes in the order book are. Note also that we use the inverse of the demand and supply schedules, with prices on the x-axis and accumulated volumes on the y-axis, as in Biais et al. (1995).

The first term inside the brackets of equation 1 measures the slope from the bid-ask midpoint to the best ask quote, while the second term is the sum of the local slopes across the remaining levels of the order book. This means that the first and second terms are not measured in the same units. The reason for this is that we are unable to calculate an elasticity for the first level of the book since there is no volume at the bid-ask midpoint. We apply several alternative measures to examine the importance of the choice of slope measurement. One alternative slope measure is the measure applied in Kalay et al. (2003)

$$SE_{i,t}^s = \frac{1}{N_A} \left\{ \sum_{\tau=0}^{N_A} \frac{(V_{\tau+1}^A - V_\tau^A)/NOSH_{i,t}}{p_{\tau+1}^A/p_\tau^A - 1} \right\} \quad (3)$$

where $NOSH_{i,t}$ is the number of outstanding shares for firm i on day t and V_τ^A is the total share volume at tick τ . Another measure is the measure applied in Kim et al. (2004), where the share volume in the order book is normalized with respect to the total share volume of all to large trades having temporarily reduced the liquidity of one side of the book or due to systematic time of day effects.

orders coming into the market on day t . A description of all the alternative slope measures we use is provided in Appendix A.

In the analysis that follows, we use a slope measure calculated from order books which *includes* hidden order volume. An illustration of order books with and without hidden order volume is provided for the two firms in Figure 1. To check whether our results are affected by the hidden order volume, we rebuild the order books without the hidden volume and calculate new slopes from the visible book. For approximately half of the daily slope estimates, the slope based on the visible book differs from the slope based on the total book although the differences are very small. Our results do not change when we re-run the regression models with the visible order book slope.

To get a picture of the difference between the slope of the full order book and the slope of the book at the inner quotes, we re-calculate the slope measure based on two different subsets of the book. The distributions of the slope estimates are shown in Figure 2. Figure 2a shows the frequency distribution of slope estimates calculated from an order book which is truncated to 5 ticks from the best quotes.¹¹ Figure 2b shows the distribution of daily slope estimates when we calculate the average slope based on ± 10 ticks of the order book. Finally, Figure 2c shows the frequency distribution when we base our slope estimates on the entire order book (± 100 ticks). The slope decreases the more of the order book we use. This is expected, if the supply and demand curves in the order book are concave.¹² The mean slope when we use the full order book is about 304 (median 231), while it increases to 417 (median 355) and 584 (median 496) when we calculate it from the order book truncated to ± 10 and ± 5 ticks respectively.

Table 4 shows how the order book slopes are related to typical measures of liquidity. The full order book slope has a positive correlation of 0.42 with market cap and a negative correlation of -0.33 with quoted spreads. This pattern indicates that larger firms are generally more liquid, with a smaller spread and a steeper slope.

¹¹That is, we use only the cumulative volume at the five first ticks on each side of the order book when we calculate the average slope.

¹²Concave when we have price on the x-axis and volume on the y-axis.

Figure 2: Frequency distribution of slope estimates

The figures show the frequency distributions for (average) daily equally weighted normalized slope estimates for all securities for the entire sample period. In Figure (a) the slope calculations are calculated using only the first 5 levels of the order book, in Figure (b) we use the first 10 levels of the order book and in Figure (c) we use the entire order book up to 100 tick levels.

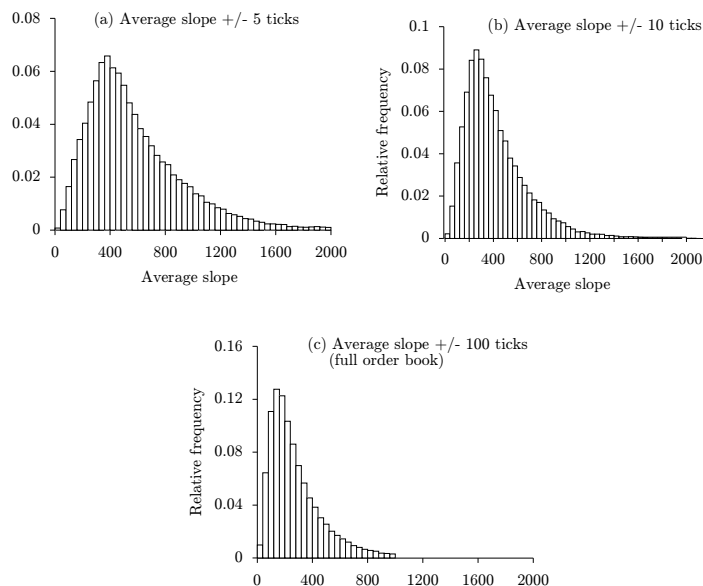


Table 4: Correlation structure for slope measures and liquidity variables

The table presents Pearson's correlation coefficients between slope measures and various measures of liquidity.

	SLOPE (Full order book)	SLOPE10 (+/- 10 ticks)	SLOPE5 (+/- 5 ticks)
Trades (N)	-0.02	-0.02	-0.02
Trade volume shares (V)	0.03	0.11	0.15
Trade size shares (AV)	-0.08	-0.06	-0.07
Market cap (MCAP)	0.42	0.43	0.46
Spread (SPR)	-0.33	-0.36	-0.38
Order volume shares (OV)	0.01	0.05	0.07
Slope (full)	1.00	0.82	0.71
Slope (+/- 10 ticks)	0.82	1.00	0.95
Slope (+/- 5 ticks)	0.71	0.95	1.00

2 Are order book characteristics informative?

In this section, we examine empirically the interaction between the order book slope and the volume-volatility relation. We start out by documenting that a standard volume-volatility relation exists in the limit order market. Then, to study whether there is any relevant information about the relation in the order book, we estimate three groups of models:

1. Price change $_{i,t} = f(\text{SLOPE}_{i,t}, ..)$
2. Number of trades $_{i,t} = f(\text{SLOPE}_{i,t}, ..)$
3. Corr(Price change $_{i,t}$, Number of trades $_{i,t}) = f(\text{SLOPE}_{i,t}, ..)$

In the first group of models, we examine whether the daily average order book slope has any explanatory power in a standard volume-volatility regression model with price volatility as the dependent variable. In the second group of models, we ask whether there is a systematic interaction between trading activity and order book slope. Finally, we investigate whether the order book slope contains information about the volume-volatility correlation, measured by the correlation coefficient between the daily number of trades and the daily price changes. All models are estimated with several control variables related to liquidity.

2.1 The volume-volatility relation in a limit order market

To investigate whether the volume-volatility relation in the Norwegian equity market is similar to the relation found for the US by e.g. Jones et al. (1994) and in the UK by Huang and Masulis (2003), we follow the regression approach in Jones et al. (1994). First, we measure the daily return volatility by running the following regression for each security i

$$R_{i,t} = \sum_{k=1}^5 \alpha_{i,k} D_{k,t} + \sum_{j=1}^{12} \beta_{i,j} R_{i,t-j} + \hat{\epsilon}_{i,t} \quad (4)$$

where $R_{i,t}$ is the return on security i on day t , and $D_{k,t}$ is a day-of-the-week dummy for day k . To avoid measurement errors due to the bid-ask bounce, we calculate returns from the average of bid-ask prices at the close. The 12 lagged return regressors estimate short-term movements in conditional expected returns. The residual, $\hat{\epsilon}_{i,t}$, is our estimate of the unexpected return on security i on date t . The absolute value of this measure constitutes our measure of volatility.

Next, we estimate the following regression equations suggested in Jones et al. (1994) to determine the relative effects on volatility of number of trades (N) and trade size (AV),

$$\text{Model 1: } |\hat{\epsilon}_{i,t}| = \alpha_i + \alpha_{i,m}M_t + \beta_i AV_{i,t} + \sum_{j=1}^{12} \rho_{i,j} |\hat{\epsilon}_{i,t-j}| + \eta_{i,t} \quad (5)$$

$$\text{Model 2: } |\hat{\epsilon}_{i,t}| = \alpha_i + \alpha_{i,m}M_t + \gamma_i N_{i,t} + \sum_{j=1}^{12} \rho_{i,j} |\hat{\epsilon}_{i,t-j}| + \eta_{i,t} \quad (6)$$

$$\text{Model 3: } |\hat{\epsilon}_{i,t}| = \alpha_i + \alpha_{i,m}M_t + \beta_i AV_{i,t} + \gamma_i N_{i,t} + \sum_{j=1}^{12} \rho_{i,j} |\hat{\epsilon}_{i,t-j}| + \eta_{i,t} \quad (7)$$

The $\rho_{i,j}$ terms measure the persistence in volatility across 12 lags. M_t is a dummy variable that is equal to 1 for Mondays and 0 otherwise, $AV_{i,t}$ is the average trade size (total number of shares traded divided by the number of transactions for security i on date t), and $N_{i,t}$ is the number of transactions in security i on date t . The regressions are run for each security and then the parameter estimates are averaged across securities.

The first part of Table 5 provides the results from the estimation of regression equations 5-7 using daily returns for all securities in our filtered sample. Overall, our results are very much in line with the results in Jones et al. (1994). The explanatory power of model 2, where volume is measured by the average number of daily trades, is almost double the explanatory power of model 1, where volume is measured by the average trade size. Moreover, the average trade size has little marginal explanatory power when volatility is conditioned on the number of transactions in model 3. These results are further supported by the characteristics of the distribution of individual security coefficients and t-statistics of the two variables. In model 3, 95.4 percent of the coefficients for the average number of trades are statistically significant, and 99.1 percent of the coefficients for the average number of trades are greater than zero. Similar numbers for the average trade size are respectively 24.1 percent and 57.4 percent. As a robustness check we also estimate the equations for sub-periods of half-years. Although not reported in a table, the results in Table 5 are confirmed in the sub-sample regressions.¹³ Hence the two components of share volume seem to contain quite different information.

Jones et al. (1994) find that trade size has some information content for some of the smaller NASDAQ-NMS firms. This finding is interpreted as supportive of the notion that private information-based trading is important only for the smallest firms on the stock market. To check for similar features in our data, we re-estimate the three regression models for the four size portfolios. The results from these estimations are presented in the second part of Table 5.

¹³Most notably, the $\hat{\gamma}$ estimates of the effect of trades (N), as well as their distributional properties, are very stable across sub-periods. The $\hat{\beta}$ estimates, however, vary considerably across sub-periods and are less significant than $\hat{\gamma}$ for model 1 relative to model 3.

In general, the results from estimating separate regression models for each size portfolio are similar to the results from running one regression for the whole sample. However, we find the opposite result from Jones et al. (1994) that the explanatory power of trade size is the strongest for the largest firms. On the other hand, only about half of the parameter estimates for trade size in the single security regressions are greater than zero, indicating that the effect may not be very systematic across firms.

2.2 Order book slope and price volatility

We now turn to the examination of the informativeness of the order book slope, starting with the first group of models described at the beginning of this section. Specifically, we estimate modified versions of the volume-volatility regression equations in sub-section 2.1, using a panel regression approach with one-way fixed effects

$$|\hat{\epsilon}_{i,t}| = \sum_{k=1}^K \mathbf{X}_{i,t,k} \beta_k + \eta_{i,t} \quad (8)$$

where $|\hat{\epsilon}_{i,t}|$ is the daily volatility estimate from equation 4, $\mathbf{X}_{i,t,k}$ is the matrix of explanatory variables (k) across time (t) for each security (i) and $\eta_{i,t} = \nu_i + \epsilon_{i,t}$ defines the error structure with ν_i as the non-random fixed, security-specific, effect.

In addition to the order book slope, we include several variables to control for liquidity. Less liquid securities generally have a higher volatility since the order book does not contain enough volume to absorb large trades without moving prices too much. In addition, less liquid stocks generally have a higher spread since investors require a discount to buy and a premium to sell the security. Thus, a positive relationship between order book slope and volatility is expected *a priori*.

Results from estimating these models over the full sample period are presented in Table 6. Model 4 is equal to model 3 in equation 7 with the addition of order book slope, market cap, spread, and order book volume as explanatory variables. In model 5, we exclude the two liquidity variables in model 4 with the highest correlation coefficient with the slope, spread (SPR) and market cap (MCAP), and in model 6 we exclude the two trading activity (mixing) variables.

The first thing to note is that the slope variable (SLOPE) is negative and highly significant across all three model specifications. Thus, volatility increases the more gentle the slope is. This may be linked to differences of opinion about public news, “noise trading” from uninformed

Table 5: A volume-volatility regression model

The table reports the results from the estimation of the following three regression models based on Jones et al. (1994):

$$\text{Model 1: } |\hat{\epsilon}_{i,t}| = \alpha_i + \alpha_{i,m} M_t + \beta_i AV_{i,t} + \sum_{j=1}^{12} \rho_{i,j} |\hat{\epsilon}_{i,t-j}| + \eta_{i,t}$$

$$\text{Model 2: } |\hat{\epsilon}_{i,t}| = \alpha_i + \alpha_{i,m} M_t + \gamma_i N_{i,t} + \sum_{j=1}^{12} \rho_{i,j} |\hat{\epsilon}_{i,t-j}| + \eta_{i,t}$$

$$\text{Model 3: } |\hat{\epsilon}_{i,t}| = \alpha_i + \alpha_{i,m} M_t + \beta_i AV_{i,t} + \gamma_i N_{i,t} + \sum_{j=1}^{12} \rho_{i,j} |\hat{\epsilon}_{i,t-j}| + \eta_{i,t}$$

Using the Jones et al. (1994) notation we have that $|\hat{\epsilon}_{i,t}|$ is the absolute value of the return on security i in period t , conditional on its own 12 lags and day-of-week dummies, M_t is a dummy variable that is equal to 1 for Mondays and 0 otherwise, $AV_{i,t}$ is the average trade size, $N_{i,t}$ is the number of transactions for security i on day t , and the coefficients $\rho_{i,t}$ measure the persistence in volatility. Columns 3-5 show parameter estimates averaged across all individual security regression equations, while columns 6-9 show the parameter distribution across securities. $\hat{\beta}$ is the average parameter estimate for the average trade size variable (AV), $\hat{\gamma}$ is the average parameter estimate for the number of trades variable (N). In the distribution of estimates columns we report, respectively, the percentage of $\hat{\beta}$ and $\hat{\gamma}$ estimates over all single security regression equations that are significant. In the last two columns we report the percentage of parameter estimates that are greater than zero. The first part of the table shows the results from running the regression equations over the whole sample. The second part of the table shows the similar results when we split the sample into four size portfolios.

Model	Firms	Parameter estimates			Distribution of estimates			
		$\hat{\beta}$ (AV)	$\hat{\gamma}$ (N)	adj. R ²	% t($\hat{\beta}$)>2	% t($\hat{\gamma}$)>2	% $\hat{\beta}$ >0	% $\hat{\gamma}$ >0
Model 1 (AV)	108	0.145	-	0.057	26.9 %	-	81.5 %	-
Model 2 (N)	108	-	0.031	0.145	-	95.4 %	-	100.0 %
Model 3 (AV,N)	108	0.053	0.031	0.149	22.2 %	94.4 %	58.3 %	100.0 %
<hr/>								
Model 1 (AV)								
1 (small)	27	0.145	-	0.080	16.2 %	-	78.4 %	-
2	27	0.219	-	0.055	18.2 %	-	77.3 %	-
3	27	0.274	-	0.048	19.0 %	-	64.3 %	-
4 (large)	27	1.021	-	0.038	30.8 %	-	79.5 %	-
Model 2 (N)								
1 (small)	27	-	0.052	0.174	-	89.2 %	-	97.3 %
2	27	-	0.028	0.147	-	75.0 %	-	95.5 %
3	27	-	0.036	0.136	-	81.0 %	-	95.2 %
4 (large)	27	-	0.014	0.174	-	79.5 %	-	92.3 %
Model 3 (AV,N)								
1 (small)	27	0.079	0.053	0.175	10.8 %	86.5 %	64.9 %	97.3 %
2	27	0.076	0.030	0.148	4.5 %	75.0 %	54.5 %	95.5 %
3	27	0.075	0.036	0.140	16.7 %	78.6 %	45.2 %	95.2 %
4 (large)	27	0.237	0.014	0.179	30.8 %	82.1 %	35.9 %	94.9 %

Table 6: Volatility and the slope of the total order book

The table shows the results from estimating different versions of the model (Model 4),

$$|\hat{\epsilon}_{i,t}| = \beta_0 M_{i,t} + \beta_1 N_{i,t} + \beta_2 AV_{i,t} + \beta_3 MCAP_{i,t} + \beta_4 SPR_{i,t} + \beta_5 OV_{i,t} + \beta_6 SLOPE_{i,t} + \sum_{j=1}^{12} \rho_{i,j} |\hat{\epsilon}_{i,t-j}| + \eta_{i,t}.$$

using a panel regression approach with fixed effects. $\eta_{i,t} = \nu_i + \epsilon_{i,t}$ defines the error structure with ν_i as the non-random fixed, security specific, effects. $|\hat{\epsilon}_{i,t}|$ is the absolute return adjusted for day-of-week effects and autocorrelation in returns. M is a dummy variable for Monday, N is the number of transactions, AV is the average trade size in shares, $MCAP$ is the natural log of the market capitalization, SPR is the relative spread (quoted spread in % of the midpoint price), OV is the total number of shares in the order book (sum of all orders on bid and ask side of the order book) and $SLOPE$ is the average slope of the bid and offer side of the order book. Panel A, shows parameter estimates, t-values and standard errors, for 3 variations of the model estimated over the total sample period. In model 5, we exclude two highly correlated liquidity variables ($MCAP$ and SPR), and in model 6 we exclude trading activity (N) and trade size (AV). The autoregressive estimates have been excluded from the table. For the F-tests, ** denotes significance at the 1 percent level.

Variables	MODEL 4			MODEL 5			MODEL 6		
	Est.	t-value	std.err	Est.	t-value	std.err	Est.	t-value	std.err
M (Monday)	0.035	0.30	0.034	0.041	1.18	0.035	-0.001	-0.03	0.035
N (trades)	0.005	43.71	0.000	0.004	39.05	0.000	-	-	-
AV (trade size)	0.018	4.62	0.004	0.024	6.03	0.000	-	-	-
MCAP (firm size)	-0.046	-1.37	0.034	-	-	-	0.211	6.33	0.033
SPR (% spread)	0.203	38.88	0.005	-	-	-	0.184	34.61	0.005
SLOPE	-0.085	-8.92	0.010	-0.152	-15.92	0.010	-0.111	-11.49	0.010
OV (order volume)	0.026	20.92	0.004	0.024	6.27	0.004	0.053	14.24	0.004
Adj. R ²	0.230			0.205			0.199		
N (firms)	108			108			108		
T (time series)	572			572			572		
F-test (no fixed effects)	12.65**			13.06**			10.61**		

investors¹⁴, or pick-off risk.¹⁵ We will discuss several interpretations of our findings at the end of section 3.2.

The total volume in the order book (OV) is shown to have a significant positive effect on volatility. This result is consistent with the result in Biais et al. (1995) that more trades are executed when the order book is thick. The correlation coefficients between order book volume and trade volume (0.44) and between order book volume and trades (0.19) also suggest that the volume-volatility relation depends on the incoming order flow and the state of the order book. The estimation results also show that larger firms have less volatile prices than smaller firms, and that higher spreads coincide with higher volatility.

The estimation results for models 5 and 6 in Table 6 are essentially the same as for model 4.

¹⁴A problem could be that a steeper slope implies a less pronounced bid-ask bounce, and thus a lower volatility. However, as outlined in sub-section 2.1, we try to avoid measurement errors due to the bid-ask bounce by calculating returns using the average of bid-ask prices.

¹⁵If some liquidity suppliers are informed about the volatility, as in the Foucault et al. (2003) model, they may find it optimal to bid less aggressively when they know that the volatility is high.

The important thing to note is that the parameter estimate for SLOPE is significantly negative and relatively stable across all three model specifications. The slope parameter is most negative and most significant in model 2, when we remove the spread (SPR) and market capitalization (MCAP) variables. This suggests that the slope captures liquidity effects on volatility.

Both the F-test of no security-specific effects (security-specific constants) and the Hausman specification test of whether a random-effects model would be more appropriate relative to the fixed effects specification are rejected at the 1 percent level for all three models.¹⁶ Since not all securities are traded every day, our sample is unbalanced. However, results from estimating the same models on a balanced sample, including 27 securities which were traded every day during the sample period, are quantitatively similar. See appendix B for estimation results for the balanced sample.

To examine the stability of the results, we estimate model 4 for non-overlapping sub-periods of three-month intervals through the entire sample period (not reported). The SLOPE parameter is remarkably stable across the sub-samples, and is significant at the 1 percent level within all sub-samples, except for the first. The number of trades parameter is also highly significant across all sub-periods while the average trade size parameter is significant at the 1 percent level only in half of the sub-sample regressions. These results suggest that the number of trades is the important component of volume in the volume-volatility relation, as was also found in section 2.1.

An interesting question is whether the slope calculated from different sets of the order book contain different information about volatility. To examine this, we re-estimate the regression models in Table 6 with slope measures calculated from the two sub-sets of the order book described in sub-section 1.4. Table 7 reports the estimation results. The results when we use the slope calculated from the +/- 10 ticks from the best quotes (SLOPE10) are reported in model 7a, while the results when we use the +/- 5 ticks from the best quotes (SLOPE5) are reported in model 7b. The main result is that the slope parameter remains negative and significant. The parameter estimates become smaller compared to the case where we used the full order book, but this is mainly due to the fact that the mean of the slope estimates increases (as shown in Figure 2) while the dependent variable remains unchanged.

We also estimate the models using the truncated slope measures across sub-periods (not

¹⁶To check for autocorrelation in the residuals, we run simple OLS regressions for each security. The hypothesis of no autocorrelation cannot be rejected for any of the single security regression models. When we test for normality, we find that the assumption is violated for less than half of the securities. The violations are due to outliers in the sample.

Table 7: Volatility and the slope of the inner book

The table presents the results from estimating two versions of the following model using a panel regression approach with one-way fixed effects,

$$|\hat{\epsilon}_{i,t}| = \beta_0 M_{i,t} + \beta_1 N_{i,t} + \beta_2 AV_{i,t} + \beta_3 MCAP_{i,t} + \beta_4 SPR_{i,t} + \beta_5 OV_{i,t} + \beta_6 SLOPE_{i,t} + \sum_{j=1}^{12} \rho_{i,j} |\hat{\epsilon}_{i,t-j}| + \eta_{i,t}$$

where $\eta_{i,t} = \nu_i + \epsilon_{i,t}$ defines the error structure with ν_i as the non-random fixed, security-specific, effects. $|\hat{\epsilon}_{i,t}|$ is the absolute return adjusted for day of week effects and autocorrelation in returns. M is a dummy variable for Monday, N is the number of transactions, AV is the average trade size in shares, $MCAP$ is the natural logarithm of market capitalization values, SPR is the relative spread (quoted spread as % of the midpoint price), OV is the total number of shares in the order book (sum of all orders on bid and ask side of the order book) and $SLOPE$ is the average slope of the bid and offer side of the truncated order book. In Model 1a (Model 1b), the average slope of the book is calculated using only the 10 (5) first tick levels of each side of the order book. The autoregressive estimates have been excluded from the table.

Variables	Model 7a (+/- 10 ticks)			Model 7b (+/- 5 ticks)		
	Estimate	t-value	std.err	Estimate	t-value	std.err
M (Monday)	0.035	0.31	0.034	0.035	0.30	0.034
N (trades)	0.005	44.48	0.000	0.005	44.60	0.000
AV (trade size)	0.017	4.34	0.004	0.017	4.29	0.004
MCAP (firm size)	-0.026	-0.78	0.034	-0.020	-0.59	0.034
SPR (% spread)	0.206	39.16	0.005	0.207	39.28	0.005
SLOPE10 (+/- 10 ticks)	-0.041	-5.93	0.007	-	-	-
SLOPE5 (+/-5 ticks)	-	-	-	-0.024	-4.82	0.005
OV (order volume)	0.027	21.30	0.004	0.027	21.40	0.004
Adj. R ²	0.23			0.23		
N (firms)	108			108		
T (time series)	572			572		
F-test fixed effects	12.55**			12.61**		

reported). The significance of the slope variable is greatly reduced within sub-periods when we only use the inner part of the book. This suggests that the inner part of the book and the other liquidity variables capture mainly the same effects, while the full order book captures additional characteristics. This is also evident from the correlation coefficients between the three slope measures and different liquidity variables presented in Table 4. Overall, our results suggest that, in addition to trading volume, both the order flow and the status of the order book are significantly related to contemporaneous volatility. In appendix A, we show that the main results presented here are invariant to several alternative ways of measuring the slope.

A caveat on causality

A potential problem with the above analysis is the indeterminacy with respect to the causality between volatility and several of the explanatory variables such as the average order book volume, number of trades, the spread and the slope measure. Although this probably is most

important at the transaction level, several of our measures are averages across hourly snapshots. Thus, dynamic interactions between order submissions and the status of the order book, as examined in detail by Biais et al. (1995), are left out of our regression model. Biais et al. (1995) find that a thin book attracts new orders while a thick book increases trading activity. On the other hand, if some liquidity suppliers are informed about the volatility in the next period, as in the Foucault et al. (2003) model, they may find it optimal to bid less aggressively when they know that the volatility will be high. To examine these issues, we run simple Granger causality tests between our slope measure and various order types and trading activity variables, both on an hourly and a daily frequency. Overall, we are unable to determine a clear one-way Granger causality relation between the variables; instead we find a two-way Granger causality for most variable combinations.

2.3 Order book slope and trading activity

In this section, we examine the relationship between order book slope and contemporaneous trading volume. Table 8 presents the results from the estimation of the second group of models described at the beginning of the section. More specifically, we estimate a cross-sectional time series regression model with the number of trades as the dependent variable, and the order book slope and several control variables related to liquidity as explanatory variables. The main result is that order book slope calculated from the total order book (model 8) is significantly negative related to the number of trades. Thus, the more gentle the order book slope, the higher the trading volume represented by the number of trades. Models 9 and 10 in Table 8 are estimated with slope measures calculated from respectively the ± 10 and ± 5 best quotes. Interestingly, we find that the parameter estimate switches sign and becomes more positive the closer we get to the inner quotes. Thus, the slope at the inner quotes is positively related to the number of trades, while the average slope for the full book is negatively related to the number of trades. In other words, the order book slope seems to contain different information about trading activity depending on what subset of the order book we use to calculate the slope. This result is in contrast to the finding in the previous sub-section that the relationship between price volatility and the full order book slope is well proxied by a slope measured over the inner part of the book.

We also find that the number of trades is lower on Mondays, that the average trade size is unrelated to the number of trades, and that larger firms are more frequently traded. In addition, we find that there is less trading when the quoted percentage spread is large, and that there

are more trades when the volume of shares in the order book is high. Again, one caveat with respect to the analysis is that we do not take into account the dynamic interactions between the order flow and status of the order book. For example, as found by Biais et al. (1995), a thinner book may attract new orders which in the next step increases the number of transactions.

Table 8: Trading activity and the order book slope

The table presents the results from estimating the following model which relates trading activity to the slope of the order book,

$$N_{i,t} = \beta_0 M_{i,t} + \beta_1 AV_{i,t} + \beta_2 MCAP_{i,t} + \beta_3 SPR_{i,t} + \beta_4 SLOPE_{i,t} + \beta_5 OV_{i,t} + \eta_{i,t}.$$

The model is estimated as a panel regression model with one-way fixed effects. The dependent variable, N is the number of transactions, M is a dummy variable for Monday, AV is the average trade size in shares, $MCAP$ is the natural logarithm of market capitalization values, SPR is the relative spread (quoted spread as % of the midpoint price), OV is the total number of shares in the order book (sum of all orders on bid and ask side of the order book) and $SLOPE$ is the average slope of the bid and offer side from the full order book, $SLOPE10$ is the slope calculated from the order book truncated to +/- 10 ticks, $SLOPE5$ is the slope calculated from the order book truncated to +/- 5 ticks. $\eta_{i,t} = \nu_i + \epsilon_{i,t}$ defines the error structure with ν_i as the non-random fixed, security-specific, effects. ** indicate that the F-test for fixed effects is significant at the 1 percent level.

Variables	MODEL 8			MODEL 9			MODEL 10		
	Est.	t-value	std.err	Est.	t-value	std.err	Est.	t-value	std.err
M (Monday)	-6.62	-4.72	1.40	-6.45	-4.58	1.41	-6.23	-4.42	1.41
AV (trade size)	1.10	6.81	0.16	0.92	5.66	0.16	0.87	5.38	0.16
MCAP (firm size)	59.57	43.68	1.36	61.51	44.89	1.37	60.22	43.80	1.37
SPR (% spread)	-2.47	-11.59	0.21	-1.55	-7.21	0.21	-1.08	-5.01	0.22
SLOPE (full book)	-8.97	-23.12	0.39	-	-	-	-	-	-
SLOPE10 (+/- 10 ticks)	-	-	-	-0.41	-1.45	0.28	-	-	-
SLOPE5 (+/- 5 ticks)	-	-	-	-	-	-	1.81	8.64	0.21
OV (order book volume)	4.98	32.83	0.15	5.09	33.44	0.15	5.11	33.55	0.15
Adj. R ²	0.423			0.417			0.418		
N (firms)	108			108			108		
T (time series)	572			572			572		
F-test (no fixed effects)	180.8**			209.6**			203.5**		

2.4 Order book slope and the volume-volatility relation

As a final analysis, we examine whether the order book slope is also interacting with the correlation between trading volume and volatility, cf. the third group of models described at the beginning of the section. The results from an estimation of a model relating the order book slope to the correlation coefficient between trading volume and volatility is presented in Table 9.

We estimate three different specifications of the model: one where we do not include explanatory variables other than the order book slope, one where we also include average trade size and total order volume as explanatory variables, and one where we use market cap and order book slope as explanatory variables. The slope is significantly negatively related to the

Table 9: The volume-volatility relation and the order book slope

The table presents the results from estimating different versions of the following model,

$$\text{Corr}(N_{i,t}, | \hat{\epsilon}_{i,t} |) = \beta_0 + \text{SLOPE}_{i,t} + \beta_1 \text{SPR}_{i,t} + \beta_2 \text{AV}_{i,t} + \beta_3 \text{OV}_{i,t} + \beta_4 \text{MCAP}_{i,t} + \eta_{i,t}.$$

The model is estimated as a panel regression model with one-way fixed effects. The dependent variable, $\text{Corr}(N_{i,t}, | \hat{\epsilon}_{i,t} |)$ is the average daily correlation coefficient measured over a quarter between the number of trades and the absolute return adjusted for day-of-week effects and autocorrelations in returns, SLOPE is the average slope of the bid and ask side calculated from the total order book, SPR is the quoted spread as % of the midpoint price, AV is the average trade size in shares, OV is the total number of shares in the order book, and MCAP is the natural logarithm of market capitalization values. $\eta_{i,t} = \nu_i + \epsilon_{i,t}$ defines the error structure with ν_i as the non-random fixed, security-specific, effects. ** indicate that the F-test for fixed effects is significant at the 1 percent level.

Variables	Corr(N, $\hat{\epsilon}_{i,t}$)			Corr(N, $\hat{\epsilon}_{i,t}$)			Corr(N, $\hat{\epsilon}_{i,t}$)		
	Est.	t-val.	std.err	Est.	t-val.	std.err	Est.	t-val.	std.err
Intercept	0.404	6.52	0.062	-0.650	-4.92	0.132	-0.215	-0.81	0.265
SLOPE (full book)	-0.021	-4.90	0.004	-0.012	-2.67	0.004	-0.014	-3.33	0.004
SPREAD (%)	-0.030	-11.31	0.003	-0.014	-4.60	0.003	-	-	-
AV	-	-	-	-0.034	-1.93	0.017	-	-	-
OV	-	-	-	0.109	11.62	0.009	-	-	-
MCAP	-	-	-	-	-	-	0.054	1.94	0.028
Adj. R ²	0.15			0.19			0.11		
N (firms)	108			108			108		
T (time series)	29			29			29		
F-test (no fixed effects)	2.05**			1.93**			2.64**		

volume-volatility relation under all three model specifications. Thus, we find evidence that the positive correlation between trading volume and price volatility relation is stronger the more gentle the order book.

3 What explains the order book slope?

In this section, we first provide some further analysis of the information content of the order book by examining variables suspected to interact with the order book slope. We then suggest some possible interpretations of our findings about the relationship between the order book slope and the volume-volatility relation.

3.1 Additional analysis of the order book slope

As there are no theoretical models available that address the relationship between the order book slope and the volume-volatility relation, we approach the task of interpreting our results indirectly by investigating empirically what factors can explain variations in the order book slope.

Firstly, we look at the hypothesized explanations for the volume-volatility relation. There

are two complementary theoretical explanations for this phenomenon. The main explanation is that new information about asset values acts as the driving force (or mixing variable) for both market prices and trading volume, see Karpoff (1987) for a review of this literature.¹⁷ The other explanation is that prices do not change merely because of new information about asset values but also because investors have dispersed beliefs about asset values. This dispersion may be due to asymmetric information or to differences of opinion about symmetric information. In any case, theoretical models by Shalen (1993) (asymmetric information) and Harris and Raviv (1993) (symmetric information) show that dispersion of beliefs will intensify the volume-volatility relation, by increasing both trading volume and volatility.¹⁸ To investigate this issue, we include as an explanatory variable the variation in analysts' earnings forecasts. This variable is used as a proxy for dispersion in beliefs in Ghysels and Juergens (2001). The analysts' earnings forecast data are obtained from Datastream/IBES, and the variation in analysts' earnings forecasts is calculated as the standard deviation of the estimates for the month as a percent of the mean estimate for that month. A small dispersion implies a strong consensus among the analysts, while a large dispersion indicates that analysts disagree on the security's future earnings.

A natural conjecture is that the order book slope is partly determined by competition among the traders, i.e. the order book will be steeper when there are many traders submitting orders than when there are few traders. To investigate this, we include as an explanatory variable the number of unique brokers who have placed orders in the book. Finally, the order book slope is also closely related to liquidity. To control for liquidity, we include market cap as an explanatory variable.

Panel A of Table 10 shows the results from the estimation of two versions of a panel regression model where we relate the average monthly order book slope for each security to the variables described above. As shown in panel B of the table, the market cap is highly correlated with the number of unique brokers who submit orders, i.e. there are more unique brokers trading in large firms than in small firms. Thus, as a robustness check, we run the model without the liquidity control. Removing market cap does not alter our results.

The order book slope is found to be significantly related to all the explanatory variables in the model. The coefficient for the number of unique brokers submitting orders is negative,

¹⁷Several empirical studies find support for this explanation under the assumption that the arrival rate of information can be proxied by the daily number of transactions (Harris, 1987; Jones et al., 1994).

¹⁸Daigler and Wiley (1999) perform an indirect test of Shalen (1993)'s model and find evidence that uninformed traders contribute to price volatility. Using the spread in analysts' earnings forecasts as a proxy for dispersion of beliefs, Ghysels and Juergens (2001) also find that dispersion is positively related to volatility.

i.e. the book is steeper when the number of unique brokers is lower. Thus, it does not seem to be the case that the order book is determined by competition among traders. The analysis also reveals that there is a negative relation between the average order book slope and the variation in analysts' earnings forecasts. In other words, the greater the disagreement among analysts, the more gentle the average slope of the order book is. This is interesting because it suggests that traders' order submission strategies are related to dispersion of beliefs about the asset values. To further investigate this, we provide some additional analysis of the spread in analysts' earnings forecasts.

Table 10: What factors explain the order book slope?

Panel A of the table shows the results from a monthly cross-sectional time series regression where we relate our slope measure, SLOPE, to the spread in analysts' earnings forecasts (coeff of variation), the natural logarithm of the number of brokers trading in the firm ($\ln(\text{brokers})$), and the natural logarithm of market capitalization values ($\log(\text{mcap})$). The spread in analysts' earnings forecasts is calculated as the standard deviation of the estimates for the month as a percent of the mean estimate for that month. This is a measure of the relative dispersion of estimates around the mean estimate. A small dispersion implies a strong consensus among the analysts. A large dispersion indicates that analysts disagree on the security's future earnings. Panel B shows Pearson's correlation coefficients between different explanatory variables.

PANEL A: Analyst forecast regression						
Variables	SLOPE			SLOPE		
	Estimate	Std.err	p-value	Estimate	Std.err	p-value
Intercept	4.351	0.481	< .01	4.474	0.467	< .01
$\log(\text{brokers})$	-0.780	0.133	< .01	-0.813	0.129	< .01
$\log(\text{mcap})$	-0.116	0.107	0.28	-	-	-
coeff of variation	-0.007	0.003	< .01	-0.007	0.003	< .01
R-sq.	0.56			0.56		
F-test	22.0			24.4		
Time series obs.	65			29		
Cross-sections	29			65		

PANEL B: Variable correlations						
	$\log(\text{mcap})$	$\log(\text{analysts})$	$\log(\text{brokers})$	slope	volatility	hidden fraction
$\log(\text{analysts})$	0.64	-	-	-	-	-
$\log(\text{brokers})$	0.51	0.45	-	-	-	-
slope	0.21	0.21	-0.11	-	-	-
volatility	-0.09	-0.18	-0.18	-0.50	-	-
hidden fraction	-0.24	-0.21	-0.71	0.03	-0.16	-
coeff of variation	-0.11	-0.06	0.06	-0.26	0.18	-0.04

In section 1, we documented that an important component of the total market liquidity consists of hidden orders. Thus, an interesting question is whether the use of hidden orders interacts with the dispersion of beliefs, i.e. do traders hide more of their order volume in periods when it seems that beliefs are more dispersed? To investigate this, we include as an explanatory variable the non-visible fraction of the daily order volume. We also include our

volatility variable, the order book slope, and the number of analysts covering the security.

Table 11: What factors explain the spread in analysts' earnings forecasts?

The regression examines the relationship between the spread in analysts' earnings forecasts and possible explanatory variables. The independent variables are the number of unique brokers submitting orders, the number of analysts covering the security, the price volatility, the fraction of share volume that is hidden, and the order book slope.

Variables	Coeff. of variation			Coeff. of variation		
	Estimate	Std.err	p-value	Estimate	Std.err	p-value
Intercept	40.276	5.751	< .01	45.095	6.131	< .01
log(brokers)	1.346	1.362	0.32	0.586	1.405	0.68
log(analysts)	-5.374	1.731	< .01	-5.945	1.756	< .01
volatility	0.639	0.456	0.16	0.434	0.476	0.36
hidden fraction	0.122	4.566	0.98	0.447	4.655	0.92
SLOPE	-	-	-	-0.623	0.256	< .01
R-sq.	0.48			0.49		
F-test	20.2			18.75		
Time series obs.	65			65		
Cross-sections	29			29		

The results from the estimations are provided in Table 11. As expected, the number of analysts covering the security and the order book slope have significant effects on the dispersion of beliefs among analysts. However, neither volatility nor the use of hidden orders seem to be related to dispersion of beliefs about asset values. A potential problem with the model is that the number of analysts is highly positively correlated with the number of unique brokers who submit orders, cf. Panel B of Table 10. We therefore re-estimate the model without the number of analysts. Removing this variable does not alter the results.

3.2 Interpretation of results

If we are to interpret our results based on the existing theoretical literature on limit order markets, such as Glosten (1994), the systematic patterns found should be related to a situation where homogeneous liquidity suppliers condition their limit orders on the probability of informed trading. Assuming a given level of liquidity-motivated trading and a given probability of informed trading, the result that the slopes are more gentle the more volatile assets are, is consistent with this model. We also find that smaller firms have order books with a gentler slope than larger firms. Assuming that there is more private information in smaller than in larger firms, the gentler slopes in small firms may indicate that there is a greater pick-off risk in these firms.

On the other hand, Sandås (2001) tests the predictions in Glosten (1994) in the Swedish market, which is a pure limit order market similar to the Norwegian market examined in this

study. He finds strong evidence that there is insufficient depth in the observed order book relative to the theoretical prediction. In other words, the slopes of the empirical demand and supply schedules at the inner quotes seem much too gentle to be explained by the theory.

Another interpretation, which is supported by the results in sub-section 3.1, is that different shapes of the limit order book across securities reflect dispersion of beliefs about security values. If uncertainty about the value of a security is high and traders have different private valuations, they should submit their orders over a wider range of prices relative to a situation where there is greater agreement among the traders about the security's value. This interpretation provides an interesting explanation for why the order volumes observed in the limit order book are more dispersed than predicted by Glosten (1994). Although no models exist that offer any predictions of how the full limit order book would look like in a market with heterogenous traders, two different types of models provide an interesting framework which could motivate such an interpretation.

Shalen (1993) shows that strategic behavior of uninformed investors may be an important contributor to both volume and volatility in addition to information arrivals. In her model, uninformed investors are faced with a signal extraction problem, and react to all types of trades in the order flow which may or may not be related to informed trading. Due to this, they increase both trading volume and price volatility above what would be expected in equilibrium. Thus, the relationship between volume and volatility is not merely due to the information arrival process (as in the mixture of distributions framework), but also due to strategic trading by uninformed traders. The higher the fraction of uninformed traders in the population, the greater the dispersion of beliefs, and the greater the excess volume and excess volatility.

Harris and Raviv (1993) explain the volume-volatility relation by a "differences of opinion" model. In their model, investors are assumed to act differently on the same news, i.e. trading is induced by differences of opinion about publicly available information. Beliefs are updated using Bayes rule. All traders are rational, but they view others as having irrational models. Two groups of risk-neutral speculators receive the same information but disagree on the extent to which it is important (but agree to disagree). As long as one of the groups remains more optimistic than the other, there is no trading. Trading occurs only and whenever the cumulative information for one of the trader groups switches from favorable to unfavorable, or vice versa.¹⁹

¹⁹Prices change every period whether or not trading occurs. The volume-volatility relation arises because the price changes are larger on average when trading occurs.

4 Conclusions

This paper has shed light on the informativeness of the limit order book. We show strong evidence that the order book slope is systematically related to price volatility, trading activity, and the dispersion of beliefs about asset values among financial analysts. Our findings are consistent with models where the volume-volatility relation is partly explained by dispersion of beliefs. Firstly, a more gentle order book slope (i.e. greater dispersion of beliefs) increases volatility and trading activity. Secondly, a more gentle slope also coincides with a stronger correlation between volume and volatility.

Our findings provide a motivation for future research on this topic. One direction for further work is theoretical and empirical research on why traders place orders with limit prices far from the inner quotes. In this paper, we have discussed asymmetric information and disagreement as two possible reasons for dispersed order books. Another reason, suggested by Sandås (2001), is that placements of orders deep in the book are based on strategic choices where the gains from obtaining price priority of orders far from the best quotes are traded off the costs of monitoring them.

Another important direction for future work is to look more closely at the causality issues briefly discussed in section 2. Another interesting exercise related to the interpretation of the order book slope could be to examine how the characteristics of the order book change around public information disclosures. This could enhance our understanding of the interaction between price formation and traders' order submission strategies. There is also a new strand of theoretical literature which focuses on liquidity suppliers who act strategically based on some form of information or market power. These models may provide explicit testable restrictions which could help to shed some light on the issues discussed in this paper.

A Alternative slope measures

To examine whether the choice of calculation method for the order book slope affects our results, we calculate four alternative measures in addition to the measure SLOPE used in the main analysis and described in sub-section 1.4. The results from running different regression models with alternative slope measures are shown in Table 12. The “RAW SLOPE” is calculated by using the volume at each tick level instead of its natural logarithm. This measure is calculated as in equations 1-2. The “NOSH SLOPE” is calculated based on an order book which is normalized with respect to the total number of outstanding shares on date t for firm i . This is similar to the slope measure used in Kalay et al. (2003), and assumes that all shares in the firms are for sale and that the prices are too low for the shareholders who do not trade on a particular day. The average slope is calculated as shown in equation 3. “OV SLOPE” is calculated based on an order book which is normalized relative to the total share volume in all orders coming into the market on that day. This measure is similar to the normalization used in Kim et al. (2004). As a final measure we apply the “NORM SLOPE”, where we normalize the order book at each snapshot relative to the total number of shares supplied in the order book at that snapshot. For one snapshot, this measure changes the volume at each price level in the order book to the percentage volume at each tick. Across securities, it implicitly standardizes the order book relative to market cap and firm liquidity. More specifically, for the ask side, let RV_{τ}^A be the fraction of total share volume on, respectively, the ask side of the book at snapshot s at price level τ . This fraction is calculated for each level of the ask side as,

$$RV_{\tau}^A = V_{\tau}^A / \sum_{\tau} V_{\tau}^A \quad (9)$$

where V_{τ}^A is the number of shares supplied at price level τ . The average slope is calculated as in equation 1, but by replacing the volumes with RV_{τ}^A .

B Balanced sample estimation

To examine the robustness of our results, we restrict our sample to securities that were traded every day through the sample period of 572 trading days. This leaves us with a balanced sample of 27 securities with 572 time series observations each. In addition, the filtering leaves us with a sample of the largest, most liquid and actively traded firms on the exchange. If the

Table 12: Estimation with different slope measures

The table shows the estimated relationship between the slope of the order book and volatility for various slope measures calculated from the full order book (+/- 100 ticks) and order books truncated at +/- 10 ticks and +/- 5 ticks. SLOPE is the slope measure we applied in the panel regression in table 6. RAW SLOPE is the slope calculated with the raw share volume at each tick. NOSH SLOPE is the slope when we normalize the share volume in the order book relative to the number of outstanding shares of the firms. OV SLOPE is the slope when we normalize the share volume in the order book relative to the total market-wide share volume in all orders coming into the market on that day. Finally, NORM SLOPE is the slope when the volume at each tick level is the fraction of total shares in the order book at that snapshot for that security.

Full sample						
	Full order book		+/- 10 ticks		+/- 5 ticks	
	Est.	t-value	Est.	t-value	Est.	t-value
SLOPE (%)	-0.085	-8.9	-0.041	-5.9	-0.024	-4.8
RAW SLOPE (mill.)	-0.203	-9.9	-0.114	-8.0	-0.083	-8.1
NOSH SLOPE	-3.886	-8.4	-3.475	-9.1	-3.757	-11.1
OV SLOPE	-5.852	-12.6	-1.630	-7.2	-1.824	-8.5
NORM SLOPE	-0.010	-14.3	-0.003	-6.2	-0.002	-4.6
February 1999 - March 2000 (13 months)						
	Full order book		+/- 10 ticks		+/- 5 ticks	
	Est.	t-value	Est.	t-value	Est.	t-value
SLOPE (%)	-0.084	-6.0	-0.049	-4.6	-0.031	-3.8
RAW SLOPE (mill.)	-0.216	-6.9	-0.099	-5.2	-0.068	-5.1
NOSH SLOPE	-5.091	-7.1	-5.418	-8.5	-5.422	-9.3
OV SLOPE	-5.636	-8.1	-4.103	-8.1	-3.442	-8.1
NORM SLOPE	-0.010	-9.0	-0.004	-5.1	-0.002	-4.0
April 2000 - June 2001 (14 months)						
	Full order book		+/- 10 ticks		+/- 5 ticks	
	Est.	t-value	Est.	t-value	Est.	t-value
SLOPE (%)	-0.114	-7.6	-0.045	-4.9	-0.029	-4.4
RAW SLOPE (mill.)	-0.201	-7.1	-0.159	-7.2	-0.133	-7.9
NOSH SLOPE	-4.520	-6.9	-3.478	-6.9	-3.886	-8.9
OV SLOPE	-5.072	-7.3	-0.770	-2.9	-1.109	-4.3
NORM SLOPE	-0.010	-11.0	-0.003	-4.6	-0.001	-3.2

previous results are mainly due to noise or outliers introduced by small illiquid securities or the unbalanced data, the balancing of the sample should reveal this. In Table 13 we re-estimate model 1 in panel A of Table 6 and model 1a and 1b in panel A of Table 7 for the balanced sample. The estimation results are quantitatively similar to the results when we use the full sample. Most interestingly, the parameter estimate for SLOPE is negative and of similar size as before. In addition, the SLOPE estimate becomes smaller (less negative) the more we truncate the order book. As before, this is mainly due to the increase in the size of the slope estimates the more the order book is truncated. Also the significance of the SLOPE is also reduced the less of the order book we use. The largest difference between the models estimated for the balanced and unbalanced sample is that the R-squared of the models is much higher for the balanced sample, suggesting that there is more noise in the unbalanced sample.

In Table 14 we re-estimate the model for trading activity (number of trades, N) in panel A of Table 8 for the balanced sample. Although the parameter estimates change more in size than was the case for the volatility models, the parameters are qualitatively similar. Most importantly, the SLOPE parameter estimate is negative when it is calculated using the full order book, and becomes less negative and less significant the more the order book is truncated. Thus, also for the balanced sample, the results suggests that the more dispersed prices are across the order book, the more trades are executed.

Table 13: Volatility/slope regression with balanced data

The table shows the estimation results of a cross-sectional time series model for the relationship between different slope measures and the daily volatility using a balanced sample. The models are similar to those estimated (for the unbalanced sample) in Table 6 (Model 1) and Table 7. The estimated model is,

$$|\hat{\varepsilon}_{i,t}| = \beta_0 M_{i,t} + \beta_1 N_{i,t} + \beta_2 AV_{i,t} + \beta_3 MCAP_{i,t} + \beta_4 SPR_{i,t} + \beta_5 OV_{i,t} + \beta_6 SLOPE_{i,t} + \sum_{j=1}^{12} \rho_{i,j} |\hat{\varepsilon}_{i,t-j}| + \eta_{i,t}.$$

where $\eta_{i,t} = \nu_i + \varepsilon_{i,t}$ defines the error structure with ν_i as the non-random fixed, security-specific, effects. $|\hat{\varepsilon}_{i,t}|$ is the absolute daily return, N is the number of transactions, M is a dummy variable for Monday, AV is the average trade size in shares, $MCAP$ is the natural logarithm of market capitalization values, SPR is the relative spread (quoted spread as % of the midpoint price), OV is the total number of shares in the order book (sum of all orders on bid and ask side of the order book) and $SLOPE$ is the average slope of the bid and offer side from the full order book, $SLOPE10$ is the slope calculated from the order book truncated to +/- 10 ticks, $SLOPE5$ is the slope calculated from the order book truncated to +/- 5 ticks. $|\hat{\varepsilon}_{i,t-j}|$ are lagged absolute returns to take into account autocorrelations.

Variables	Volatility/slope (full book)			Volatility/slope (+/- 10 ticks)			Volatility/slope (+/- 5 ticks)		
	Est.	t-value	std.err	Est.	t-value	std.err	Est.	t-value	std.err
M (Monday dummy)	-0.083	-1.5	0.057	-0.081	-1.4	0.057	-0.078	-1.4	0.057
N (trades)	0.004	34.6	0.000	0.004	35.6	0.000	0.004	35.7	0.000
AV (avg. trade size)	0.017	2.5	0.007	0.014	2.0	0.007	0.013	1.9	0.007
MCAP (market cap.)	-0.015	-2.7	0.006	-0.009	-1.6	0.006	-0.008	-1.4	0.006
SPR (% quoted spread)	0.345	16.3	0.021	0.358	16.9	0.021	0.351	16.6	0.021
SLOPE (full book)	-0.008	-10.0	0.001	-	-	-	-	-	-
SLOPE10 (+/- 10 ticks)	-	-	-	-0.007	-10.6	0.001	-	-	-
SLOPE5 (+/-5 ticks)	-	-	-	-	-	-	-0.005	-9.6	0.001
OV (order book volume)	0.011	3.1	0.004	0.010	3.0	0.004	0.011	3.0	0.004
Adj. R ²	0.38			0.38			0.37		
N (firms)	25			25			25		
T (time series)	572			572			572		
F-test (no fixed effects)	41.6**			41.2**			40.9**		

Table 14: Volume/slope regression for balanced sample

The table shows the results from estimating the relationship between the slope measures based on different truncations of the order book and the number of trades using a balanced sample. The models are similar to those estimated for the full (unbalanced) sample in Table 8. The estimated model in panel B is,

$$N_{i,t} = \beta_0 M_{i,t} + \beta_1 AV_{i,t} + \beta_2 MCAP_{i,t} + \beta_3 SPR_{i,t} + \beta_4 SLOPE_{i,t} + \beta_5 OV_{i,t} + \eta_{i,t}.$$

where $\eta_{i,t} = \nu_i + \varepsilon_{i,t}$ defines the error structure with ν_i as the non-random fixed, security-specific, effects. N is the number of transactions, M is a dummy variable for Monday, AV is the average trade size in shares, $MCAP$ is the natural logarithm of market capitalization values, SPR is the relative spread (quoted spread as % of the midpoint price), OV is the total number of shares in the order book (sum of all orders on bid and ask side of the order book) and $SLOPE$ is the average slope of the bid and offer side from the full order book, $SLOPE10$ is the slope calculated from the order book truncated to +/- 10 ticks, $SLOPE5$ is the slope calculated from the order book truncated to +/- 5 ticks.

	Trades/slope (full book)			Trades/slope (+/- 10 ticks)			Trades/slope (+/- 5 ticks)		
	Est.	t-value	std.err	Est.	t-value	std.err	Est.	t-value	std.err
M (Monday dummy)	-15.05	-3.6	4.224	-14.01	-3.3	4.227	-13.98	-3.3	4.216
AV (avg. trade size)	-0.50	-0.9	0.514	-0.49	-0.9	0.514	-0.34	-0.7	0.513
MCAP (market cap.)	8.10	19.7	0.411	7.41	17.8	0.417	6.89	16.4	0.420
SPR (% quoted spread)	-32.39	-21.0	1.543	-32.47	-21.0	1.545	-32.07	-20.8	1.540
SLOPE (full book)	-0.50	-8.1	0.062	-	-	-	-	-	-
SLOPE10 (+/- 10 ticks)	-	-	-	0.33	6.7	0.049	-	-	-
SLOPE5 (+/- 5 ticks)	-	-	-	-	-	-	0.44	11.3	0.039
OV (order book volume)	3.33	12.8	0.262	3.43	13.1	0.262	3.45	13.2	0.261
Adj. R ²	0.33			0.33			0.34		
N (firms)	25			25			25		
T (time series)	572			572			572		
F-test (no fixed effects)	233.2**			224.0**			211.9**		

References

- 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.
- Bloomfield, R., M. O’Hara, and G. Saar (2003). “The “make or take” decision in an electronic market: Evidence on the evolution of liquidity.” Working Paper, Cornell University. Forthcoming in the *Journal of Financial Economics*.
- Chan, L. and J. Lakonishok (1993). “Institutional trades and intra-day stock price behavior.” *Journal of Financial Economics*, 33, 173–199.
- Chan, L. and J. Lakonishok (1995). “The behavior of stock prices around institutional trades.” *Journal of Finance*, 50, 1147–1174.
- Daigler, R. and M. Wiley (1999). “The impact of trader type on the futures volatility-volume relation.” *Journal of Finance*, 54, 2297–2316.
- Foucault, T., S. Moinas, and E. Theissen (2003). “Does anonymity matter in electronic limit order markets.” Unpublished paper.
- Ghysels, E. and J. L. Juergens (2001). “Stock market fundamentals and heterogeneity of beliefs: tests based on a decomposing of returns and volatility.” Unpublished paper, University of North Carolina.
- Glosten, L. (1994). “Is the electronic open limit order book inevitable?” *Journal of Finance*, 49, 1127–1161.
- Glosten, L. and P. Milgrom (1985). “Bid, ask and transaction prices in a specialist market with heterogeneously informed traders.” *Journal of Financial Economics*, 14, 71–100.
- Goldstein, M. A. and K. A. Kavajecz (2004). “Trading strategies during circuit breakers and extreme market movements.” *Journal of Financial Markets*, 7, 301–333.
- Harris, L. (1987). “Transaction data tests of the mixture of distributions hypothesis.” *Journal of Financial and Quantitative Analysis*, 22(2), 127–141.
- Harris, M. and A. Raviv (1993). “Differences of opinion make a horse race.” *Review of Financial Studies*, 6, 473–506.
- Huang, R. D. and R. W. Masulis (2003). “Trading activity and stock price volatility: evidence from the London Stock Exchange.” *Journal of Empirical Finance*, 10, 249–269.
- Jones, C., G. Kaul, and M. Lipson (1994). “Transactions, volume, and volatility.” *Review of Financial Studies*, 7, 631–651.
- Kalay, A., O. Sade, and A. Wohl (2003). “Measuring stock illiquidity. An investigation of the demand and supply schedules at the TASE.” *Journal of Financial Economics*, 74(3), 461–486.
- Karpoff, J. (1987). “The relation between price changes and trading volume: A survey.” *Journal of Financial and Quantitative Analysis*, 22(1), 109–126.
- Kim, J.-W., J. Lee, and R. Morck (2004). “Heterogeneous investors and their changing demand and supply schedules for individual common stocks.” NBER Working Paper Series, no 10410.

- Sandås, P. (2001). “Adverse selection and competitive market making: empirical evidence from a limit order market.” *Review of Financial Studies*, 14(3), 705–734.
- Shalen, C. T. (1993). “Volume, volatility, and the dispersion of beliefs.” *Review of Financial Studies*, 6(2), 405–434.