ELSEVIER Journal of Banking & Finance 24 (2000) 1323–1357

Journal of BANKING & FINANCE

www.elsevier.com/locate/econbase

An exploratory analysis of the order book, and order flow and execution on the Saudi stock market

Mohammad Al-Suhaibani [a], Lawrence Kryzanowski [b,*]

[a] Department of Economics, Imam University, Riyadh, Saudi Arabia

[b] Department of Finance, Faculty of Commerce, Concordia University, 1455 De Maisonneuve Blvd. West, Montreal, Que., Canada

Received 27 October 1998; accepted 22 June 1999

Abstract

The microstructure of the Saudi Stock Market (SSM) under the new computerized trading system, ESIS, is described, and order and other generated data sets are used to examine the patterns in the order book, the dynamics of order flow, and the probability of executing limit orders. Although the SSM has a distinct structure, its intraday patterns are surprisingly similar to those found in other markets with different structures. We find that liquidity, as commonly measured by width and depth, is relatively low on the SSM. However, liquidity is exceptionally high when measured by immediacy. Limit orders that are priced reasonably, on average, have a short duration before being executed, and have a high probability of subsequent execution. © 2000 Elsevier Science B.V. All rights reserved.

JEL classification: G15

Keywords: Market microstructure; Limit order book; Intraday patterns; Order execution

* Corresponding author. Tel.: +1-514-848-2782; fax: +1-514-848-4500.

E-mail addresses: mohisuh@alumni.concordia.ca (M. Al-Suhaibani), ladfi53@vax2.concordia.ca (L. Kryzanowski).

PII: S0378-4266(99)00075-8

1. Introduction

The recent availability of order, quote, and transaction data from stock markets around the world has stimulated research on intraday stock market phenomena. Intraday patterns identified in the data of US and other developed countries include the persistent U-shaped patterns in returns, number of shares traded, volumes, bid–ask spreads, and volatility. [1,2] Other studies that examine order-driven markets provide new evidence on patterns in the order book, order flow, and the interaction between the order book and order flow. [3]

In this paper, we study the Saudi Stock Market (SSM) which uses a computerized trading mechanism known as Electronic Securities Information System (ESIS). The objective is to examine the behavior of market participants in the SSM to understand better the effect of order placement on market liquidity, and to determine whether certain patterns identified in earlier studies can be generalized to other trading structures. Our paper has several unique aspects. First, the SSM, which is described in detail in the next section, is a pure order-driven market with no physical trading floor, regulated brokers or market makers, and it is closed to foreign portfolio investments. The market also is differentiated by a long mid-day break, partially hidden order book, and a constant tick size. Second, the unique data set provided by the Saudi Arabian Monetary Agency (SAMA) includes all orders for listed stocks submitted during the period from 31 October 1996 to 14 January 1997. This order data set allows for the construction of the complete limit order book for this order-driven market. The data set includes information that allows for the identification of market and limit orders, and what we called order packages. Third, we believe that our study is the first to examine the market microstructure of the SSM. We provide evidence on several issues related to the interaction between the order book and order flow, which adds to the existing empirical literature on order-driven markets. Finally, our paper examines a number of new issues associated with order-driven markets. The literature on market microstructure often discusses liquidity measures such as width, depth, resil-

---

[1] U-shaped patterns refer to the heavy trading activity on financial markets at the beginning and at the end of the trading day, and the relatively light trading activity over the middle of the day (Admati and Pfleiderer (1988)).

[2] For the US markets, these include studies by Wood et al. (1985), Jain and Joh (1988), McInish and Wood (1991, 1992), Brock and Kleidon (1992), Gerety and Mulherin (1992), Foster and Viswanathan (1993) and Chan et al. (1995a,b). McInish and Wood (1990) report similar results for the Toronto Stock Exchange and Lehmann and Modest (1994) find U-shaped patterns in trading for the Tokyo Stock Exchange.

[3] A representative example is the empirical analysis by Biais et al. (1995) of the limit order book and order flow on the Paris Bourse. Niemeyer and Sandås (1995), Hedvall and Niemeyer (1996), Niemeyer and Sandås (1996) and Hedvall et al. (1997) perform similar analyses for stock markets in Stockholm and Helsinki.

iency, and immediacy that may have more relevance for market-order traders. Our unique data set allows us to examine liquidity measures that are relevant for limit order traders, the only suppliers of liquidity on the SSM. Using order duration and logit regressions, we present new evidence on the probability of executing a limit order on the SSM.

The remainder of this paper is structured as follows. Section 2 presents a detailed description of the current trading system. The data sets are described in Section 3. Sections 4 and 5 analyze the limit order book and order flow, respectively. Section 6 presents and analyzes the empirical findings on limit order execution. Section 7 concludes the paper.

2. Market description

The SSM is relatively new in age compared to the stock markets in the developed countries. The first company went public in Saudi Arabia in 1954. By the end of 1982, 48 companies traded in the Saudi market, which was completely unregulated by the government. [4] The collapse of the unregulated stock market in Kuwait motivated the Saudi government to take regulatory action in 1984. [5] The new regulations transferred share trading, which occurred in the over-the-counter market, from the hands of the unofficial brokers to the banks. Because of low volume and lack of coordination between the banks, a delay of several days or weeks often occurred before orders were filled. Several other restrictions resulted in lengthy delays. Banks could neither hold positions in stocks nor break up large blocks of shares to accommodate buyers. [6]

A major development in trading on the SSM post-market-regulation was the establishment in 1990 of an electronic trading system known as ESIS. [7] After the startup of ESIS, the banks established twelve Central Trading Units (CTUs). All the CTUs are connected to the central system at SAMA. The bank CTUs, and designated bank branches throughout the country that are connected to the CTU (ESISNET branches), are the only locations where buy and sell orders can be entered directly into ESIS.

[4] Due to religious considerations, only stocks are traded in the market. From the viewpoint of *sharia* (Islamic law), interest on bonds is regarded as usury.

[5] More information on the Kuwaiti financial crises, which is known as the "Souq al-Manakh" crisis, is found in Darwiche (1986).

[6] In 1992, SAMA allowed the banks to manage open-ended mutual funds for public investors. However, the banks are still not allowed to invest directly or indirectly, through the mutual funds, in Saudi stocks.

[7] More on the history of the SSM up to 1990 is found in Malaikah (1990), Wilson (1991), and Butler and Malaikah (1992).

Trading on the SSM consists of four hours per day, divided into two daily sessions for Saturday through Wednesday. The trading day consists of one two-hour session on Thursday. Table 1 summarizes trading hours and trading days on the SSM. During the morning and evening hours no trading occurs, but *wasata* can add and maintain order packages and orders that were entered through their CTU or ESISNET branches. The *wasata* are neither brokers nor dealers. They are order clerks whose assigned job is merely to receive and verify orders from public traders at the CTU, and then to enter these orders into the system. Conditional on SAMA approval, the banks hire and pay the *wasata*. Sell and buy orders are generated from the incoming sell and buy order packages. If an order package has many firm orders, each is differentiated by parameters such as quantity, price and validity period. [8] Order packages entered into the system may be valid for a period from 1 to 12 days. [9]

At some point of time during the first five-minute opening period, all firm buy and sell orders participate in a *call market*. [10] Orders are executed at an equilibrium price calculated to be the best possible price for executing the maximum number of shares available in the market at the open. This is followed by a *continuous auction market,* where marketable orders by public investors are transacted with the limit orders of other public investors. [11] In the post-trading period, trades are routed to settlement, trading statistics are printed, and no order package or order can be added or maintained.

Only limit orders with a specified price and firm quantity are permitted. Firm orders are eligible for execution during the opening and continuous trading periods according to price-then-time priority rules. An investor can

---

[8] In ESIS terms, order packages are called orders, and orders are called quotes. These definitions differ from those usually used in the literature. Order in the literature usually refers to order with a firm quote that leads instantly to a bid or ask if it is a limit order, or to a trade if it is a market order. The firm quotes (as defined by the ESIS) are more like orders as usually defined in the literature. In the market, generating a firm quote is the same as placing an order. To be consistent with the literature, orders are referred to as order packages, and quotes are referred to as orders.

[9] Before 28 May 1994, the validity period for an order package was either 1, 5 or 10 days. Subsequently, the validity period became 1, 6 or 12 days. From 1 October 1994, the validity period was allowed to be any period from 1 to 12 days.

[10] In a call market, orders for a stock are batched over time and executed at a particular point in time.

[11] A limit order is an order with specific quantity and price and for a given period of time. For a limit buy (sell) order, the price is below (above) the current ask (bid). Marketable limit order is a limit order with a limit price at or better than the prevailing counterparty quote. For a marketable buy (sell) order, the price must equal or better the current ask (bid). Notice that the standard market order (order to buy or sell a given quantity for immediate execution at the current market price, without specifying it) is not accepted by the system. Since marketable and market orders are essentially similar, we use the term market order when referring to marketable orders in the remainder of the paper.

Table 1
Trading hours and trading days on the SSM[a]

Days	From Saturday to Wednesday		Thursday	
Time	From	To	From	To
Morning period[b]	8:15 AM	10:00 AM	8:15 AM	10:00 AM
The first opening period	10:00 AM	10:05 AM	10:00 AM	10:05 AM
The first continuous trading session[c]	10:05 AM	12:00 AM	10:05 AM	12:00 AM
The second opening period	4:25 PM	4:30 PM	None	None
The second continuous trading session[c]	4:30 PM	6:30 PM	None	None
Post-trading period	6:30 PM	7:00 PM	12:00 AM	12:30 PM
Evening period[b]	7:00 PM	8:00 PM	12:30 PM	1:30 PM

[a] *Source:* SAMA, ESIS: Instructions to Central Trading Units.
[b] No trading occurs during these periods. However, *wasta* can add and maintain order packages and orders that were entered through their CTU or ESISNET branches.
[c] The first and second continuous trading periods are 115 and 120 minutes in elapsed time, respectively. Thus, the second continuous trading period is 5 minutes longer than the first continuous trading period.

adjust order prices and their quantities, or change a firm order to on-hold at any time. [12] With each change, the order loses its time priority. When adjusted, the order price must be within its order package quantity and price limit. Aggressive sell (buy) orders can walk down (up) the limit order book. [13] When an order is partially executed, any unexecuted balance is automatically placed in a new order at the same price and with the same execution priority as the original order. The order package can be executed fully or partially through more than one transaction at different times, with different orders, and even with different prices.

To reduce adverse selection problems, the system has some negotiation capability beside the automatic routing and execution. [14] A transaction only with large value (usually SR 1/2 million [US$133,333] or more) can be executed

[12] All or part of an order package can be canceled by putting it "on-hold" or returning it back to the market at any time. "On-hold" orders are out of the market but still in the system. As a result, they have no price or time priority, and do not become automatically firm after executing all or part of the outstanding firm quantity in the order package.

[13] The limit order book ('the order book') is the collection of all firm limit orders generated from all order packages arrayed in descending prices for bids and in ascending prices for asks.

[14] Adverse selection problems exist if some traders have superior information and cannot be identified. In such situations, the uninformed traders lose on average to informed traders. Without uncertainty, the uninformed traders would trade with each other and not trade with the informed traders.

as a *put-through* transaction outside the system under SAMA supervision. [15] The parties to *put-through* transactions have no obligation to trade at the current quotes or clear the limit orders in between. After execution, the transaction is immediately reported to the market.

The minimum price variation, or tick size, for all stocks in the market is SR 1 ($\approx$27 cents). Transaction fees are charged on each side of the trade, and have a minimum of SR 25 ($\approx$\$6.66). Transaction fees range between 0.5% and 0.1% of the trade value depending on the number of shares executed. The commission is distributed in two parts: 95% to the banks, and 5% to the SSRC for settlement and transfer services. [16] During continuous trading periods, firm orders must be priced within $\pm$10% of the opening price of the given trading period. In turn, the opening price must lie within a price range that is within $\pm$10% of the previous day's closing price. If no opening price exists for that period, the opening price defaults to the previous day's closing price. Occasionally SAMA can allow the price to exceed the present fluctuation limit provided the new price is reasonably justified by the earnings or prospects of the company.

The electronic limit order book is not fully visible to investors since information is displayed publicly in an aggregate format (i.e., only the best quote with all quantities available at that quote). The status of the best quotes and quantities is updated (almost instantaneously) on bank screens each time an order arrives, is canceled, or is executed. Public investors can view the price, quantity, and time of last trade. The terminals and big screens where traders can monitor the market are only available in the CTUs and ESISNET branches of the banks. In the early releases of ESIS, only the *wasata* in the CTUs could view the best five bids and asks, and valued bank customers could easily learn this information by calling their bank's CTU. To prevent this type of unfair access to market information and related front-running problems, SAMA on 1 October 1994 restricted both the *wasata* in the CTUs and the public to viewing only the best two bids and asks. The

[15] *Put-through* transactions (so-called block trades) are not common on the SSM, and usually are handled in an informal manner. In most cases, big traders agree in advance on the transaction and ask SAMA to handle it as a *put-through* transaction. For this reason, the price of the transaction may not reflect current market conditions. If this is the case, SAMA sends a message communicating this information about the trade to the market. Occasionally, an unofficial broker brings in both sides of the *put-through* transaction. In rare cases, an uninformed trader appeals for SAMA supervision to minimize the transaction costs associated with a very large order by handling it as a *put-through*. To facilitate the transaction by this verified uninformed trader, SAMA sends a massage to the CTUs asking for counterparties to complete the transaction.

[16] The SSRC (Saudi Share Registration Company) was formed in 1985 by the Saudi banks to serve as a clearing system for executed trades. Under ESIS, the major role of SSRC is to keep up-to-date records of shareholdings in stock companies.

wasata still have more information about the order book since they know the details of every order placed through their CTU or their ESISNET branches connected to it. This includes the identification of investors, the price and quantities of firm and on-hold orders, and the type of ownership document for sell orders. Details of every order are only observable to surveillance officials. This level of transparency on the SSM hides all firm orders outside the two best quotes. Unlike on-hold orders, hidden orders have price and time priority and can be revealed to the market or executed at any time. For example, a firm order to buy with a price less than the second best bid is hidden but becomes visible when all the quantity at the first best quote is executed. The order also can be executed while it is hidden by an aggressive market sell order. [17]

Only the *wasata* in the CTUs have the right to enter orders directly into the system. Investors in the SSM consist of public investors and bank phone customers. [18] Bank phone customers have an agreement with the banks to change the price and firm quantity of their submitted orders at any time simply by calling their Bank's CTU. As a result, they are less affected than other public traders by the free trading option associated with limit orders since they can change the condition of their orders very quickly before they are "picked off" when new public information arrives. [19] This group of traders includes the institutional investors (e.g. mutual funds) and many technical traders who have trading and no fundamental information.

The date and time of transfer of beneficial ownership for each transaction is the date and time of execution in the system. [20] Transaction confirmation slips are usually printed at CTUs and ESISNET branches and distributed to the clients after each trading session. Following the second trading session, transactions are routed for settlement. The settlement date depends on the type of ownership document. *Ishaar*, which can be retained in the system for

[17] Unlike some trading systems, ESIS does not allow traders to intentionally hide orders that are part of the best two quotes.

[18] SAMA does not allow banks to grant their customers access to the system via any computer network.

[19] As Stoll (1992) explains, a limit order provides the rest of the market with a free option. The trader who places a buy (sell) limit order has written a free put (call) option to the market. For example, suppose the trader submits a buy limit order at $100. If public information causes the share price to fall below $100, this put option will be exercised and the public trader loses because he cannot adjust the limit price quickly. The ability to change limit price more quickly by bank phone customer makes the effective maturity of his limit order very short, and hence the value of the put option associated with this order is almost zero.

[20] The ex-dividend day usually comes before the company closes its record for dividend payments. The company and SAMA agree in advance on this date, and communicate the date to the CTUs.

future sale or printed and given to the investor, are delivered next day morning. [21] In contrast, certificates take from two days to one week or more to be delivered. *Ishaar* takes less time because it can be handled electronically through ESIS Fully Automated Share Transfer (ESISFAST), while the new certificate has to be issued from the company's share registration department. The goal is to abolish all existing share certificates at some future point in time. [22] Because of the difference in settlement dates, and to prevent the creation of two markets for every security, the type of ownership document is not visible to market participants prior to a transaction.

3. The data sets

The data set provided by SAMA consists of intraday data on firm orders for all stocks listed on the market for 65 trading days (31 October 1996 to 14 January 1997). Four of the 71 stocks are excluded due to an absence of orders, three stocks are excluded because they have no transactions, and eight stocks are excluded because they have a small number of transactions. The final data set includes 267,517 orders for the remaining 56 stocks. For each order, the data set reports security code, the date and time of creation, buy–sell indicator, limit price, quantity, and date and time when the order was terminated (canceled, expired, or executed). Because the data uniquely identify the order package that generates the order, the order package data set can be easily constructed from the order data set. Our data set has 86,425 order packages. [23]

Given the information in our order data set, we construct another (a third) data set containing the end-of-minute best five quotes and their associated depths on both sides of the market for all 13,955 minutes of trading. [24] Subsequent references to quotes (bids and asks) are reserved for this data set. We use the date and time of termination, price and quantity of orders along with

[21] On March 19, 1994, SAMA reduced the *ishaar* delivery date to one day instead of two days. Starting from October 1, 1994, *ishaar* was allowed to be issued in the same branch where the order was submitted. Since September 1995, the buyer can know the type of ownership document immediately after executing his buy order. The latest version of ESIS released in June 1997 permits real time settlement for *ishaar* (i.e., execution and settlement times are the same).

[22] During the sample period, around 95% to 97% of trades have *ishaar* documents.

[23] Chan and Lakonishok (1995) use the trading package terminology to describe the trader's successive purchases of a stock. The correspondence between their definition of a trade package and an ex ante order is approximate. In contrast, for our data set, we have more information about orders since we know the set of orders that was generated from an order package. However, we still are unable to confirm that two orders belong to the same ex ante order if the investor broke up a large order into two submitted order packages.

[24] The depth is the number of shares offered or demanded at a given bid or ask.

published daily statistics to identify the order that was part of a transaction (trade data set). The number of transactions in our sample is 84,382. Table 2 presents some summary statistics for each of our four data sets.

Panel A in Table 2 reports summary statistics for the order data set. Limit orders account for 71% of the orders in the sample. The percentage of buy and sell orders is almost equal for most stocks. Most orders (63%) are executed. Based on Panel B, most of the order packages are to sell. Execution rates are similar and evolve around 0.5. Based on Panel C, the public limit order traders supply immediacy to the market nearly all the time with an average inside spread equal to SR 2.24.

Panel D reports the summary statistics for the transaction data set which includes all market orders, the limit orders executed against them, and the orders executed against each other during the call market at the opening. Because two orders constitute each trade, the number of observations in this data set are twice the number of transactions as conventionally reported. Less than 10% of the trades occur during the opening period, and a very small percentage (0.015%) of the trades are executed outside of the system (in the so-called upstairs market). The average returns are positive since the market rose 9.23% over the sample period.

4. Descriptive statistics about the order book

The order book collects all limit orders at any given point of time. Orders come into the book throughout the day at the time they are submitted to the market, and are removed from the book as they are executed, canceled, or expired. Using the quote data set, this section presents and discusses various descriptive statistics concerning the order book. Although our subsequent analyses are based on the five best quotes, it is important to remember that market participants only observe the first two best quotes.

4.1. Relative spreads and depths in the order book

Table 3 reports the time series means and medians of relative spreads between adjacent quotes in the book, and depths at all levels for the 56 stocks in the sample. The spread is usually one, two or three ticks in our sample. Based on Panel A , the mean (median) relative inside spread is 1.79% (1.6%) which is high compared to other markets. [25] Angel (1997) uses data on the bid–ask

[25] The inside spread is the difference between the first best ask ($A1$) and the first best bid ($B1$). The relative inside spread is the inside spread divided by the quote midpoint, or: $2(A1 - B1)/(A1 + B1)$.

Table 2
Summary statistics for each of the four data sets[a]

Order or trade characteristic	All observations	Cross-sectional distribution across the 56 stocks					
		Mean	Min	First quartile	Median	Third quartile	Max
Panel A: Orders data set							
Number of observations	267,517	4,777	411	1,104	3,027	6,946	26,240
Buy (%)	48.88	50.10	44.74	47.96	49.22	52.00	59.72
Limit (%)	71.24	73.84	67.44	71.34	72.63	77.09	83.10
Limit Buy (% of limit orders)	46.24	49.38	41.00	45.91	48.57	52.20	63.39
Market Buy (% of Market orders)	55.41	51.09	32.71	48.53	53.86	56.35	61.36
Executed orders (%)	63.09	58.87	36.31	56.43	60.80	62.43	77.09
Order size	843.40	814.79	113.61	464.99	700.12	1,076.30	2,972.80
Large order (%)	0.62	0.28	0.00	0.00	0.13	0.38	2.01
Panel B: Order packages data set							
Number of observations	86,425	1543	138	396	1109	1900	8180
Buy (%)	38.52	39.93	13.75	33.64	40.82	44.81	63.04
Package size	2,610.64	2,359.90	272.02	1,341.20	2,157.40	3,080.20	8,409.40
Orders per package	3.095	2.969	2.015	2.637	2.909	3.206	4.350
Execution rate	0.5711	0.548	0.343	0.516	0.546	0.590	0.793

Panel C: Quotes data set							
Number of observations	778,593	13,903	11,960	13,955	13,955	13,955	13,955
Availability of immediacy (%)	97.66	97.66	80.75	97.67	100.00	100.00	100.00
Inside spread	2.247	2.274	1.038	1.278	1.533	2.541	10.351
Quote midpoint return (×1000)	–	0.005	−0.015	0.002	0.005	0.008	0.019
Panel D: Transactions data set							
Number of observations	168,764	3,014	154	656	2,045	4,281	17,438
Trades at open (%)	8.81	10.96	3.10	5.28	7.72	14.52	34.48
Trade size	560.88	518.78	52.03	284.76	483.58	721.98	1,372.10
Transaction price	267.3091	196.07	23.62	70.65	111.81	253.94	959.17
Trade-to-trade return (×1000)	–	0.114	−0.615	0.007	0.051	0.147	2.039
Put-through trades (%)	0.15	0.10	0.00	0.00	0.00	0.13	1.04

[a] For the 65 trading days over the period between October 31, 1996 and January 14, 1997, the first column reports various order and trade characteristics after pooling all stocks. The other columns report the cross-sectional distribution of these statistics across the 56 stocks in the sample. All the reported statistics are mean values except for the number of observations and the percentages. The size statistics are computed using the number of shares. The large orders and put-through trades are those with volumes larger than SR 0.5 million. Immediacy is considered available when both bid and ask are established. Inside spread is the difference between the first best ask and the first best bid. The quote midpoint returns are based on the end-of-minute quote midpoints, while trade-to-trade returns are computed using the time series of transaction prices. Execution rate is the number of shares that are filled divided by the total number of shares submitted as a package.

Table 3
The relative spreads and depths in the book[a]

Panel A: The relative spreads between successive levels of the limit order book (x100)

Relative spread	B4–B5	B3–B4	B2–B3	B1–B2	A1–B1	A2–A1	A3–A2	A4–A3	A5–A4
Mean	1.271	1.297	1.240	1.193	1.790	1.281	1.337	1.412	1.348
Median	1.288	1.205	1.115	1.057	1.600	1.246	1.251	1.393	1.436

Panel B: The average volumes at different levels of the limit order book

Depth	B5	B4	B3	B2	B1	A1	A2	A3	A4	A5
Mean	4394	5741	8321	10,319	5616	4072	6926	6374	5672	4410
Median	2081	2711	3370	3448	1910	1514	2764	2949	2665	2443

Panel C: Test of equality of spreads and depths across levels in the order book

Hypothesis	Test statistic	Calculated	F-probability
All relative spreads are equal	$F(8,492) =$	1.9380	0.0526
All relative spreads excluding inside spread are equal	$F(7,492) =$	0.2884	0.9698
All depths are equal	$F(9,550) =$	2.6379	0.0054
All depths are equal (excluding the depths at the second best quotes)	$F(7,550) =$	1.3255	0.2203

[a] Using the best bids and asks and their associated depths, this table reports the means and medians of the relative spreads between adjacent quotes and the quantities offered or demanded at these quotes. The reported depth is the original number of shares divided by 100. A and B denote ask and bid, respectively. B1 is the first best bid, and A1–B1 is the relative inside spread [(first best ask – first best bid)/Quote midpoint] times 100. The quote midpoint is calculated as (first best ask + first best bid)/2.

spread for major market indices for fifteen countries and finds that the median relative spread equals 0.65%. The relative tick size, as is shown in the next section, is the major contributing factor to this high relative spread. The relative inside spread is larger than all other relative spreads on either side of the book. The other relative spreads are moderately constant. In contrast, the average numbers of shares at the first best quote are small (and the smallest on the ask side), are the largest at the second best quote, and decrease beyond the second quotes. [26]

Based on the test results reported in Panel C, the hypotheses that all relative spreads and all depths are equal are rejected, but not rejected when we exclude the inside relative spread, and the depth at the second quotes. [27] The liquidity provision is greater on the bid side. On average, depths are larger and relative spreads are smaller on the bid side.

Our results lie somewhat between those of Biais et al. (1994) and Niemeyer and Sandås (1995). Using data from the Paris Bourse, Biais et al. find that the order book is slightly concave, with an inside spread more than twice as large as the difference between the other levels of the book (which is similar to our results). They also find that the volumes offered or demanded at the first best quotes are smaller than the volumes further away from the best levels. In contrast, Niemeyer and Sandås find that the order book on the Stockholm Stock Exchange is convex. Spreads are wider further away from the inside spread, and volumes are larger close to the inside spread. In fact, they find as we do that the average volumes at the second best quote are the largest.

As Fig. 1 shows, the slope of the order book in our market does not depart strongly from linearity. [28] It is slightly concave near the second quote and convex thereafter. One possible interpretation for this shape is that the adverse selection problem is more pronounced closer to the inside spread. This leads to a higher inside spread, and smaller volumes at the first best quotes. Since all of the five best quotes are available to market participants on the Paris Bourse, and only the best two on the SSM, the contradiction between our results and

[26] The number of orders contributing to each quote (not reported) also has the same pattern as the volumes. Namely, they exhibit an inverted U-shape. They are largest at the second best quotes, and smaller for the other quotes.

[27] The test is conducted using dummy variable regressions of the form $y = b_1 d_1 + \cdots + b_p d_p$, where y is the relative inside spread (or the depth) for all stocks after we stack all observations; d_i, $i = 1, \ldots, p$, is a dummy equal to one if the observation y belongs to the book level i; and p equals 9 for relative spread tests and 10 for the depth tests. We perform the reported equality tests using different sets of linear restrictions.

[28] On the SSM, large trades that execute against several limit orders at different prices will have two prices: marginal and average prices. The plot of price changes for trades of different sizes (as in Fig. 1) is an approximation of the marginal price function of the limit order book or of the slope of the book.

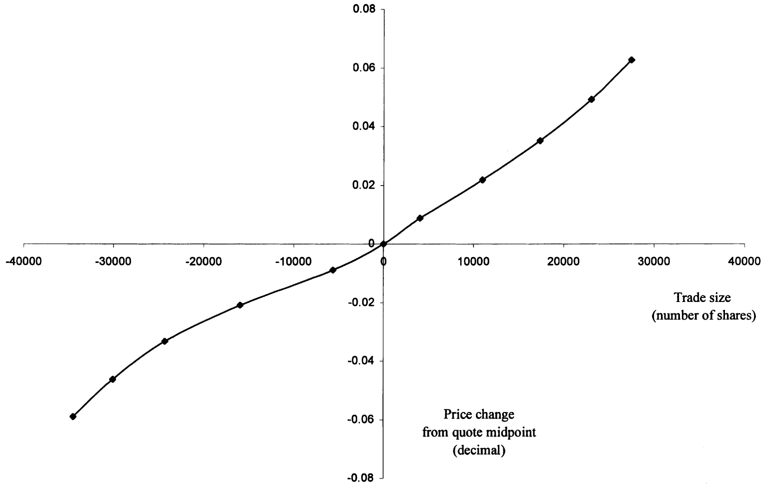

Fig. 1. The average price schedule on the SSM. Using the average relative spreads and depths at various levels of the order book, this figure plots the decimal changes in the transaction price relative to the quote midpoint for trades of different sizes. Negative trade sizes represent sell transactions.

those of Biais et al. may be due to the difference in the information available, which can affect the strategies of market participants. However, our data do not allow us to determine how the volume would be distributed for a different information disclosure structure.

Because the relative inside spread is larger and the depth lower, market liquidity as usually measured by width and depth is relatively low. [29] Market order traders can buy or sell a large number of shares but only at high transaction costs.

4.2. Tick size and price discreteness

The SSM has one tick size of SR 1, which imposes price discreteness and forms a lower bound on the spread. The prices of the stocks in our sample range from 24 to 956 SR implying a minimum relative spread (or relative tick

[29] Four dimensions are often associated with liquidity in the market microstructure literature: width, depth, immediacy and resiliency. According to Harris (1990), width refers to the spread for a given number of shares, depth refers to the number of shares that can be traded at given quotes, immediacy refers to how quickly trades of a given size can be done at a given cost, and resiliency refers to how quickly prices revert to former levels after they change in response to large order flow imbalance initiated by uninformed traders. Overall, a market is liquid if traders can quickly buy or sell large numbers of shares when they want at a low transaction cost.

size = 1/price) between 4.21% and 0.1%. The median relative tick size is 0.9% which is relatively large compared to the median relative tick size for major stock markets. Using data for 2517 stocks that constitute the majority of the capitalization in the world equity market, Angel (1997) finds that the median relative tick size is equal to 0.259%.

Theoretically, a large tick size encourages limit order traders to provide liquidity to the market, and imposes higher transaction costs on market order traders. Given the price and time priority rules, the limit order trader has a first mover advantage only if the tick size is large enough to prevent quote matching. [30] If the tick size is small, then the quote matcher obtains time precedence by submitting an order at a price slightly better than the standing quote.

Based on the summary statistics on tick size reported in Table 4, 53.77% of the inside spreads are binding (the inside spread equals one tick), 22.48% equal two ticks, and 23.75% equal three or more ticks. Tick size is more important for lower priced stocks. The tick size is binding for 76.7% of the observations for stocks in the lowest price category, and for only 25.86% of the stocks in the highest price category. In unreported results, we find that the majority of the other spreads are binding even for highly priced stocks. The last row of Table 4 supports the assertion that large tick size encourages limit order traders to provide liquidity to the market. The percentage of limit orders submitted to the market increases as the relative tick size increases. This might suggest that a larger tick induces liquidity. A larger tick however increases transaction costs for market order traders, which may reduce overall liquidity for stocks. The optimal tick, as Angel (1997) concludes, is not zero. Its optimal size represents a trade-off between the benefits of a nonzero tick for limit order traders and the cost that a tick imposes on market order traders.

4.3. Availability of immediacy

Immediacy is available in the market when a market order can be instantaneously executed. In an order-driven market as the SSM, the availability of immediacy depends upon the limit order traders. Immediacy will be unavailable if no public limit orders are present. Table 5 summarizes the percentages of time when immediacy is unavailable at all levels of the book. Despite the absence of market makers, market liquidity measured by immediacy is notably high. On average, the immediacy at the first best bid and ask is unavailable for only 1.51% and 1.19% of the total trading time,

[30] Quote-matchers are traders whose willingness to supply liquidity depends on the limit orders of other liquidity suppliers. Harris (1990) discusses the quote-matcher problem in detail.

Table 4
Tick size statistics for the SSM[a]

Variable	All stocks	Price level sub-samples				
		1 (Lowest)	2	3	4	5 (Highest)
Number of quotes at all levels (in millions)	5.688	0.913	1.111	1.120	1.255	1.164
Quote midpoint range	23.62 to 956.15	23.62 to 64.71	64.71 to 93.48	93.48 to 167.71	167.71 to 329.57	329.57 to 956.15
Average quote midpoint	195.27	46.37	77.94	118.72	226.32	469.73
Inside spreads that equal one tick (%)	53.77	76.70	62.10	52.77	52.09	25.86
Inside spreads that equal 2 ticks (%)	22.48	16.89	21.98	25.34	25.97	21.97
Inside spreads that equal 3 or more ticks (%)	23.75	6.41	15.91	21.89	21.93	52.17
Spread (in ticks)	2.278	1.336	1.825	1.965	2.193	4.196
Relative inside spread	1.79%	3.12%	2.27%	1.70%	1.02%	0.91%
Relative tick size	1.04%	2.38%	1.30%	0.87%	0.46%	0.22%
Limit order (%)	59.4	64.2	61.4	60.9	58.1	56.8

[a] This table presents statistics on tick sizes on the SSM. The statistics are computed for all 56 stocks in the sample and for five sub-samples classified by the mean of stock price during the sample period. We classify the sample using price because the tick is constant and equal to SR 1 for all stocks, which implies that the relative tick size can be measured by the inverse of price. Since the tick size is one, the spread (in ticks) is the same as the observed spread in the market. The relative inside spread is (first best ask – first best bid)/quote midpoint. Quote midpoint = (first best ask + first best bid)/2. The relative tick size is 1/quote midpoint. Limit order is the percentage of limit orders to the total number of orders.

respectively. [31] As expected, most active stocks have even lower percentages. The difference between the five categories becomes more evident as we move away from the first best quotes.

4.4. Intraday pattern in the order book

In this section we examine the intraday patterns in the relative inside spread, depth and the squared quote midpoint return. [32] As shown in Fig. 2, the relative inside spread decreases over the first trading session, and is fairly constant over the second. The test results reported in Panel A of Table 6 support this result. In the first session, the last trading interval has the lowest relative spread (1.74%). The regression is constructed so that the slopes represent the difference between the mean relative spread in this interval and the other intervals in the

[31] We should keep in mind that these statistics are for the more active stocks in the market since we eliminated the most thinly traded stocks from our sample.

[32] The quote midpoint is the average of the best bid and ask quotes.

Table 5
The availability of immediacy at all levels of the book on the SSM[a]

Variable	All stocks	Order frequency sub-samples				
		1 (Lowest)	2	3	4	5 (Highest)
Mean number of orders	4777	564	1536	3157	5897	11,544
Immediacy is unavailable (%)						
B5	43.66	73.98	70.47	45.04	19.92	8.75
B4	30.36	58.83	55.50	23.25	12.02	2.88
B3	16.25	39.85	29.13	8.93	3.55	0.46
B2	5.30	15.78	5.86	3.60	1.39	0.04
B1	1.51	4.35	0.64	1.50	1.07	0.00
A1	1.19	4.59	0.16	1.21	0.01	0.00
A2	4.68	16.68	4.24	1.71	1.03	0.00
A3	13.42	40.97	23.14	2.73	1.22	0.02
A4	23.30	64.18	44.05	8.10	1.40	0.12
A5	32.73	80.04	60.22	21.84	2.04	0.48

[a] This table summarizes the relative durations of times when immediacy is unavailable at the best five quotes on both sides of the market. Immediacy will be unavailable whenever there is no limit order to buy or sell. Relative duration is the total time that immediacy was impaired as a percentage of the time that the SSM was open over the sample period. B and A denote bid and ask, respectively. B1 and A1 are the first best bid and the first best ask, respectively.

session. As constructed, the *t*-statistics are direct tests of whether any differences exist in mean relative spreads. Moving from the first to the seventh coefficient estimate one finds that both the difference and significance decrease. We also reject the hypothesis that all differences are zero. In contrast, no significant patterns are identified in the second trading session.

While many studies document a U-shaped intraday pattern for the spread, [33] other studies report patterns similar to that found in our market. Chan et al. (1995a) find that NASDAQ spreads are at their highest at the open and narrow over the trading day. Similar results are reported by Chan et al. (1995b) for the CBOE options, and by Niemeyer and Sandås (1995) and Hedvall (1995) for two order-driven markets, the Stockholm Stock Exchange and the Helsinki Stock Exchange, respectively.

If the spread is a good proxy for transaction costs, the relative inside spread pattern together with patterns found in trading activities (see Section 5.3) is not supportive of most of the models for explaining trade concentration. Admati and Pfleiderer (1988) present a model where concentration of trading may be generated at an arbitrary time of the day. Liquidity traders, particularly traders who have to trade within a given time period, pool their trades in an effort to

[33] Studies which find a U-shaped pattern in the spread include Brock and Kleidon (1992), McInish and Wood (1992), Foster and Viswanathan (1993) and Lehmann and Modest (1994).

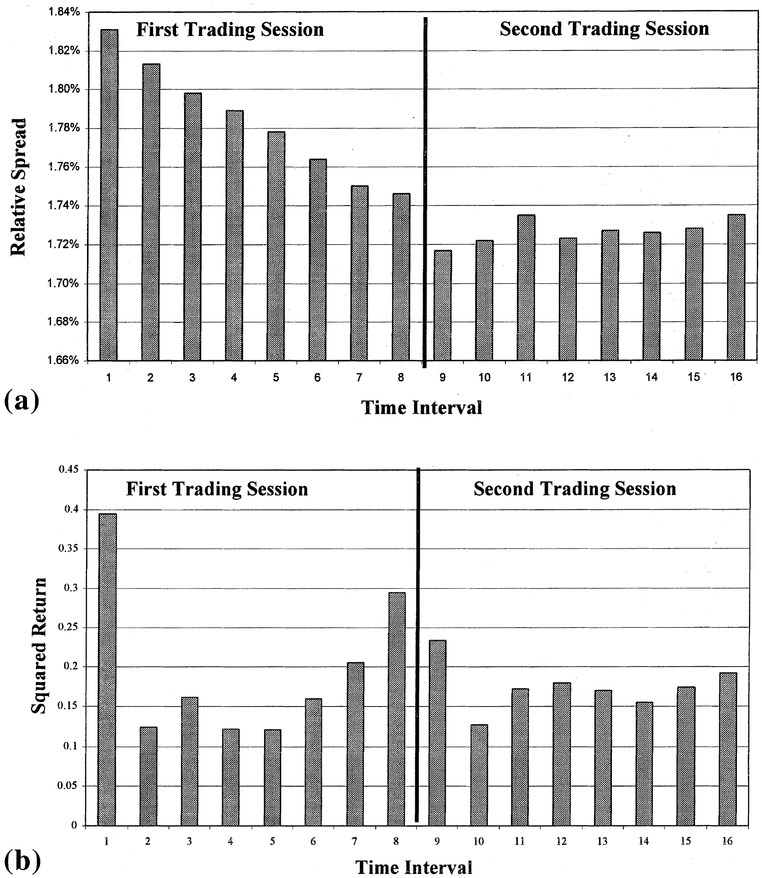

Fig. 2. Intraday patterns in the order book. This figure reports the intraday relative inside spreads and squared quote midpoint returns. Each trading session is divided into eight intervals, and the daily relative spread and squared midpoint return are computed for each interval for all stocks in the sample. The bars are the averages over the 65 days in the sample. The relative inside spread = (best ask − best bid)/QMP, where QMP denotes quote midpoint = (best ask + best bid)/2. The quote midpoint return is calculated as $\log(\mathrm{QMP}_t) - \log(\mathrm{QMP}_{t-1})$. (a) Intraday relative spread. (b) Intraday squared return (×100,000).

reduce their transaction costs. Informed traders, in an attempt to hide their trading intentions, also trade at the same time. The model predicts that traded volume should be highest when transaction costs are lowest. Similarly, Brock and Kleidon (1992) conjecture that periodic market closure results in greater liquidity demand at the open and close. In response, liquidity suppliers may practice price discrimination by changing their quotes during these periods of high demand. This model implies high transaction volumes and concurrent wide spreads at both the open and close.

Table 6
Tests for intraday patterns in the order book for the SSM[a]

		Interval	Coefficient	t-Statistic	P-value
Panel A: Relative inside spread (×100)					
First session					
No. of observations	520				
Omitted interval	8	B_8	1.7456	86.6036	0
$F(7,512)$	2.302	$B_1–B_8$	0.0857	2.9906	0.0029
P-value	0.0256	$B_2–B_8$	0.0673	2.4501	0.0146
		$B_3–B_8$	0.0519	1.8634	0.063
		$B_4–B_8$	0.0438	1.5452	0.1229
		$B_5–B_8$	0.0321	1.1316	0.2583
		$B_6–B_8$	0.0187	0.654	0.5134
		$B_7–B_8$	0.0049	0.1732	0.8626
Second session					
No. of observations	432				
Omitted interval	1	B_1	1.717	73.8843	0
$F(7,512)$	0.0933	$B_2–B_1$	0.0048	0.1493	0.8814
P-value	0.9986	$B_3–B_1$	0.018	0.581	0.5616
		$B_4–B_1$	0.0058	0.1908	0.8487
		$B_5–B_1$	0.0101	0.3287	0.7426
		$B_6–B_1$	0.0089	0.2926	0.7699
		$B_7–B_1$	0.0114	0.3795	0.7045
		$B_8–B_1$	0.0177	0.6111	0.541
Panel B: Squared quote midpoint return (×100,000)					
First session					
No. of observations	520				
Omitted interval	5	B_5	0.1211	5.2417	0
$F(7,512)$	4.2354	$B_1–B_5$	0.2738	6.0002	0
P-value	0.0001	$B_2–B_5$	0.0031	0.1153	0.9083
		$B_3–B_5$	0.0408	1.5114	0.1313
		$B_4–B_5$	0.0012	0.0462	0.9632

Table 6 (Continued)

		Interval	Coefficient	t-Statistic	P-value
		B_6–B_5	0.0384	1.3627	0.1736
		B_7–B_5	0.0836	2.4448	0.0148
		B_8–B_5	0.1734	1.4003	0.162
Second session					
No. of observations	432				
Omitted interval	2	B_2	0.1279	7.8211	0
F (7,512)	1.0904	B_1–B_2	0.1051	3.4395	0.0006
P-value	0.3683	B_3–B_2	0.0444	1.3719	0.1708
		B_4–B_2	0.0516	1.0435	0.2973
		B_5–B_2	0.0416	1.0471	0.2957
		B_6–B_2	0.0274	0.9058	0.3655
		B_7–B_2	0.0458	1.6122	0.1077
		B_8–B_2	0.0642	2.8834	0.0041

[a] This table reports the results from regressing relative spreads and squared midpoint returns on a set of dummy variables. Each trading session is divided into eight intervals. The daily relative spreads and squared midpoint returns are computed for each interval for all stocks in the sample. The regression equation takes the form $Y = C + B_1 D_1 + \cdots + B_8 D_8$, where Y denotes the relative inside spread (or squared quote midpoint return) during all intervals and days after all the observations are stacked; and D_i, $i = 1, \ldots, 8$, is a dummy variable that equals one if the observation y belongs to interval i. To avoid linear dependency among the explanatory variables, only seven of the eight possible dummy variables are used in each regression. The dummy variable belonging to the interval with the lowest mean is deleted for this purpose. In this formulation, the constant term represents the coefficient of the deleted dummy variable, while the other coefficients represent the difference between each of the other intervals and the omitted interval. t-Statistics based on White covariance matrix estimation provide a direct test of whether any intraday differences exist between the omitted interval and the other intervals. F-statistics show the overall significance (all differences are zero). Separate regressions are performed for each trading session.

However, the observed spread pattern for the SSM can be explained using the model of Madhavan (1992). The high spread in the morning is due to greater uncertainty. As information asymmetries are partially resolved, traders become informed by observing the market. This leads to a decline in the spread during the day. The explanation offered by Chan et al. (1995b) attributes such a spread pattern to the absence of specialist market power.

We use the squared midpoint quote returns as a measure of stock return volatility. As shown in Fig. 2 and the regression results reported in Panel B of Table 6, volatility is at its highest during the first trading interval, followed by the last trading interval before the close. [34] Considered in isolation, this finding is consistent with the information-based model of Admati and Pfleiderer (1988), which predicts that high volume periods have more informative and hence more volatile prices. No significant patterns are identified for the number of shares and volume for the first best quotes.

5. Order flow dynamics on the SSM

In this section, we investigate the dynamics of order flow on the SSM. We condition our analysis on order direction (buy or sell), price position, state of the book, and time of the day.

5.1. Order flow and the limit price position

We divide the orders into 14 categories (or events) based on limit price position. On the buy side, the price position of a buy order may be above the prevailing ask (aggressive market buy), at the prevailing ask (market buy), within the existing spread (limit buy within), at the prevailing bid (limit buy at), below the prevailing bid but above or at the second bid (limit buy below), and below the second bid (hidden limit buy). The last event is the cancelation of a previously posted limit buy. Orders on the sell side are categorized similarly. The frequency of each occurrence is documented in the last row of Table 7. With regard to market orders, the most frequent events are market sell and buy orders (11.48% and 13.41%, respectively). The frequency of aggressive orders is very small. On the limit order side, the most frequent events are limit orders at prevailing quotes.

In Table 7, the columns correspond to an event at time *t,* and the rows to events at time $t-1$. Each row reports the percent frequency of each of the

[34] The U-shaped pattern in volatility is documented for other markets by Wood et al. (1985), Harris (1986), McInish and Wood (1992), Foster and Viswanathan (1993), and Lehmann and Modest (1994).

Table 7
Order flow conditional on the position of the last limit price[a]

t	Number of observations	Aggressive market sell	Aggressive market buy	Market sell	Market buy	Limit sell within	Limit buy within	Limit sell at
$t-1$								
Panel A: All observations								
Aggressive market sell	2297	50.63	0.04	3.27	1.31	4.01	0.87	0.44
Aggressive market buy	5307	0.00	50.29	0.96	2.37	0.49	4.22	1.51
Market sell	31580	1.72	0.15	35.46	3.97	8.04	1.56	2.59
Market buy	36901	0.06	1.86	3.25	44.81	2.56	8.92	24.26
Limit sell within	10825	0.06	0.78	6.15	32.47	13.67	2.82	7.17
Limit buy within	10250	0.36	0.39	25.23	6.25	4.01	13.32	5.86
Limit sell at	31911	0.05	0.21	3.07	22.51	1.91	2.38	29.96
Limit buy at	33911	0.04	0.16	21.72	3.77	3.14	1.96	6.75
Limit sell above	12035	0.07	6.12	4.37	7.18	2.68	2.46	9.40
Limit buy below	12413	1.84	0.37	7.79	5.04	3.44	2.77	7.21
Hidden limit sell	27759	0.12	2.03	5.01	5.64	3.14	2.83	8.03
Hidden limit buy	18288	0.71	0.51	5.93	4.59	2.98	2.53	5.17
Cancel limit sell	25382	0.19	0.60	8.56	6.02	3.58	2.84	9.09
Cancel limit buy	16572	0.29	0.42	8.00	5.82	3.48	3.23	8.00
Unconditional	275246	0.83	1.93	11.48	13.41	3.94	3.73	11.60

	Limit buy at	Limit sell above	Limit buy below	Hidden limit sell	Hidden limit buy	Cancel limit sell	Cancel limit buy
Aggressive market sell	1.18	0.30	23.29	0.96	7.10	0.96	5.66
Aggressive market buy	0.89	19.79	0.55	10.97	1.41	5.88	0.70
Market sell	38.22	0.94	1.03	1.94	1.36	1.53	1.62
Market buy	3.35	1.10	0.99	2.55	1.53	3.96	0.87
Limit sell within	6.55	4.72	3.12	9.40	4.91	5.89	2.30

Limit buy within	8.78	3.47	7.06	8.13	8.57	5.68	3.08
Limit sell at	5.61	4.68	3.35	8.82	4.08	11.56	1.83
Limit buy at	22.06	3.36	6.77	8.67	5.98	7.29	8.41
Limit sell above	7.47	17.57	4.15	17.03	5.57	15.0976	0.8475
Limit buy below	12.28	3.89	13.43	10.06	13.73	1.45	16.70
Hidden limit sell	7.38	5.86	4.13	26.46	6.57	21.86	0.95
Hidden limit buy	6.56	3.06	6.04	9.65	21.77	0.88	29.60
Cancel limit sell	8.45	5.31	4.24	15.30	5.45	27.34	3.18
Cancel limit buy	11.24	3.96	7.52	10.42	16.70	3.50	17.62
Unconditional	12.32	4.37	4.51	10.09	6.65	9.22	6.01

Panel B: Diagonal percent frequencies in the sub-samples

	Number of observations	Aggressive market sell	Aggressive market buy	Market sell	Market buy	Limit sell within	Limit buy winthin	Limit sell at
The same trader	79627	86.92	88.94	73.52	77.19	43.92	51.46	22.82
Different traders	195563	2.73	10.41	2.83	8.81	4.79	4.20	31.13

	Limit buy at	Limit sell above	Limit buy below	Hidden limit sell	Hidden limit buy	Cancel limit sell	Cancel limit buy
The same trader	30.21	12.24	8.5	30.45	27.1	18.75	17.16
Different traders	20.79	19.18	15.21	24.57	17.72	28.33	17.71

[a] For all trading days and stocks, this table reports the empirical percent frequencies for 14 events related to limit price position, conditional on the previous event. The events are as they are defined in Section 5.1. Rows correspond to events at time $t-1$, and columns correspond to events at time t. Each row adds up to 100%.

twelve events conditional on the event in that row. The table supports the "*diagonal effect*" found in Biais et al. (1995) that the probability that a given event will occur is larger after this event has just occurred than it would be unconditionally. For example, market sell (buy) orders are most frequent after market sell (buy) orders. [35] Biais et al. put forward three explanations for this correlation. First, the succession of identical types of orders could reflect strategic order splitting, either to reduce the market impact of a non-informational trade, or to get the most from private information about the value of the stock. Second, if different traders are imitating each other, the cause of the correlation is the order flow itself. Finally, traders could react similarly to the same events related to a particular stock or the economy as a whole.

Since our data sets do not identify traders, we cannot explicitly investigate the three hypotheses concerning individual order submission behavior. However, we know that orders originating from the same order package certainly belong to one trader, and this allows us to infer a subset of orders belonging to the same trader. The fraction of observations where the same trader acted in two subsequent events is 28.94% of all of the order flow events. [36] If the order-splitting hypothesis is the dominant factor in explaining order flow correlation, then we should observe higher percentages of subsequent events that are initiated by the same trader. This is indeed the case as shown in Panel B in Table 7. The percentages of the same trader subsequent events are larger for most events, which indicates that the "diagonal effect" is more common in the same trader subset. Hedvall and Niemeyer (1996) use a data set from the Helsinki Stock Exchange that includes dealer identities and find, as in our market, that strategic order splitting is more common than imitation. Further, the imitation hypothesis cannot explain the diagonal effect in hidden orders. Since the traders have no incentive to split hidden orders, the only possible explanation is traders reacting to similar information events.

The diagonal effect in the case of limit orders within the best quotes, not conditional on trader identity, has been explained by the undercutting and overbidding behavior of traders competing to supply liquidity to the market (Biais et al., 1995). The results in Panel B of Table 7 do not support this explanation. The gradual narrowing of the spread, as a result of placing quotes within the spread, comes mainly from the same trader and not from compe-

[35] The diagonal effect is present beyond one lag. When we account for additional lags, we find similar effects.

[36] Given the limited information concerning trader identification for our data set, the frequencies of subsequent order events on different sides of the market from one trader are always zero. In reality, these frequencies may not be zero. However, the fact that market regulation does not match and execute two orders if they are generated from the same trader makes this possibility less likely. One trader can make a market in one or more stocks by posting limit orders on both sides of the market, but he can not make a false market by executing his market orders against his limit orders.

tition between different traders. However, the succession of cancelation is consistent with the explanation that traders imitate each other or react similarly to the same events.

Based on Panel A of Table 7, we find that market buys (sells) are exceptionally frequent after asks (bids) at and within the best quotes. Traders prefer to wait for additional liquidity to be provided, and preferably at a better price, before deciding to trade. In contrast, limit orders to buy (sell) at the quotes are particularly frequent after market sell (buy) orders. Since a market sell (buy) order consumes the existing liquidity and may lead to a downward (upward) shift in the book, the observed behavior may reflect competition between limit order traders to restore liquidity. Market liquidity in terms of resiliency is considerable.

Several other observations are consistent with information effects in the order process. After aggressive and market sell (buy) orders, there are often new limit sell (buy) orders placed within the quotes. Furthermore, limit buy (sell) orders placed away from the quote and cancelations on the buy (sell) side of the book are more frequent after aggressive market sell (buy) orders. The order book tends to shift downward (upward) after aggressive market sell (buy) orders. This behavior could reflect the adjustment in market expectations to the information content of these trades. Biais et al. observe a similar effect after large trades, and attribute their observation to the information effect.

Using χ^2 tests for the significance of the equality between the conditional and unconditional probability for all stocks, we reject the hypothesis at the 1% level.

5.2. Order flow and the state of the order book

Table 8 reports the probabilities of different types of orders and trades occurring given the previous state of the book. The state of the book is summarized by the size of the inside bid–ask spread and the depth at the first best quotes. Both the spread and depth for a given stock are defined to be large (small) when they are larger (smaller) than their respective time series medians over the sample period. Consistent with earlier theoretical and empirical findings for order flow, market orders occur more frequently when the spread is tight. Limit orders occur within the spread more frequently when the spread is large. Limit orders "offer liquidity when it is scarce" and market orders "consume it when it is plentiful" (Biais et al., 1995).

Limit orders within the spread occur more frequently when the depth at the quote is large, and limit orders at the quotes are relatively more frequent when the depth is small. Given the price and time priority rules, the only way to increase the probability of execution when the depth is large (and especially when the spread is large) is to undercut or overbid the best quote. Based on χ^2

Table 8
Order flow conditional on the state of the order book[a]

Spread	Depth	Number of observations	Aggressive market sell	Aggressive market buy	Market sell	Market buy	Limit sell within	Limit buy within	Limit sell at	Limit buy at	Limit sell above	Limit buy below	Cancel limit sell	Cancel limit buy
Large	Large	128,486	0.73	1.71	11.02	13.46	5.31	5.11	11.31	11.11	12.64	9.81	10.81	6.98
	Small	104,145	0.65	1.54	9.75	10.87	3.04	2.63	13.26	14.98	17.44	13.40	7.50	4.93
Small	Large	17,304	2.06	4.63	13.75	16.68	2.70	3.63	8.62	8.76	11.06	9.73	10.60	7.78
	Small	25,311	1.30	2.80	19.29	21.35	1.46	1.28	8.25	9.98	12.83	9.73	7.31	4.41
Unconditional		275,246	0.83	1.93	11.47	13.41	3.94	3.73	11.60	12.32	14.37	11.16	9.22	6.02

[a] For all trading days and stocks, this table reports the empirical percent frequencies for twelve events related to limit price position conditional upon the state of the book (summarized by spread and depth) one second before submitting the order. Aggressive (market) order is an order with price better than (equal to) the opposite prevailing quote. Limit order within (at, above or below) is the order priced within (at, away from) the inside spread. For each stock, the spread and depth are defined to be large, if they are larger than their time-series medians during the sample period.

tests, we reject the null hypothesis at the 1% level of the independence between the order and trade events and the state of the book.

5.3. Order flow and the time of the day

In this section, we examine the pattern of number and volume of all, limit and sell orders, and all, small and large transactions. As Fig. 3 shows, the number and volume of all new orders and transactions exhibit a U-shaped pattern during each within-day session, and a W-shaped pattern over the trading day. The proportions of orders and trades submitted are largest in the morning. The proportions in the first trading interval in the second session are usually larger than the proportions at the end of the day. The concentration around the open and close are like those observed in many stock markets with different microstructures. [37]

The call market may be a contributing factor to the concentration at the opening. Since all qualifying orders are executed at a single price at the open, traders benefit from their orders being executed at a price better than their quote. Limit order traders are less affected by free option problems at the open, and lose less to informed traders if they trade during the call market. [38] The large proportion of limit orders during the first interval of each session supports this explanation. The high level of limit orders at the end of every trading session could result from limit price adjusting. [39] Less patient traders start to adjust their prices as the end of the session approaches in order to induce other traders to execute against them (Niemeyer and Sandås, 1995).

A larger proportion of small orders is executed at the opening, whereas larger proportions of large orders are executed during and at the end of the session. One possible explanation for this behavior put forward by Biais et al. (1995) is that small traders at the opening contribute to price discovery, while large trades tend to occur after price discovery has already occurred.

We test for the significance of the patterns in number and volume for new orders and transactions using a similar regression to the regression used in Section 4.4. The unreported results indicate significant U-shapes. No significant intraday pattern was identified for transaction price.

[37] See, for example, Jain and Joh (1988), McInish and Wood (1990,1991), Gerety and Mulherin (1992), Foster and Viswanathan (1993), Biais et al. (1995) and Niemeyer and Sandås (1995).

[38] If a trader trades only at the call, then the option value of his order is smaller because it is good at the time of the call.

[39] Adjusting the limit order price or quantity results in the order receiving new date and time stamps. Accordingly, an order adjustment leads to two events: canceling an existing order, and submitting a new one.

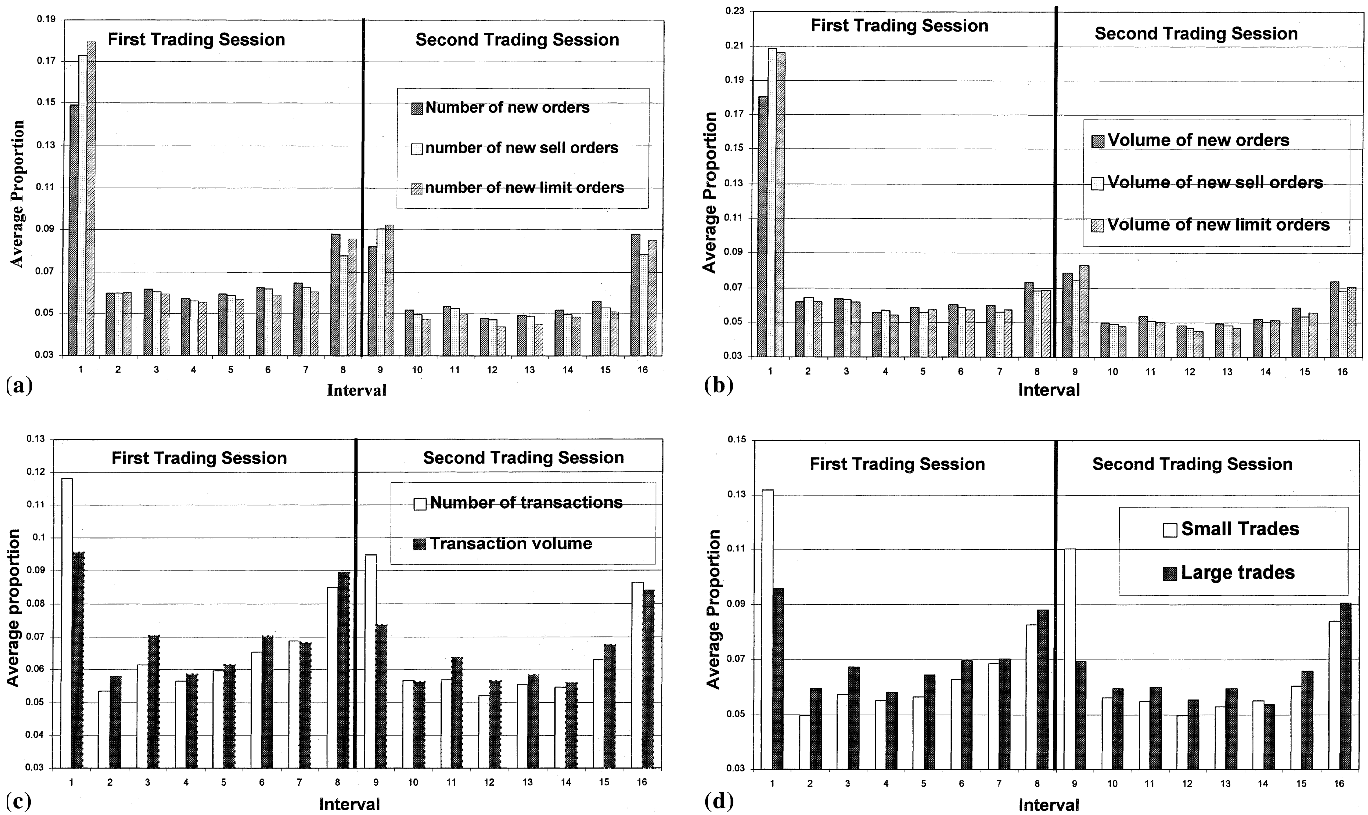

Fig. 3. Intraday patterns in the order flow on the SSM. Various plots of the number and volume of new orders and transactions are provided below. Each trading session is divided into eight trading intervals, and the number and volume of orders (transactions) in each interval are computed as proportions of the total daily number and volume of orders (transactions). Each bar is the average proportion across the 65 trading days in the sample. Transactions are defined to be large if they exceed their time series medians over the sample period. (a) Number of orders. (b) Order volume. (c) Transactions. (d) Trade size.

6. The analysis of order execution

Market liquidity can be measured by the cost of effecting a transaction at a given point of time, or by the time it takes to transact (Lippman and McCall, 1986; Amihud and Mendelson, 1989). In our examination of the order book and order flow, various aspects of the cost measure of market liquidity (such as width, depth, resiliency and the availability of immediacy) were addressed. The cost measure of liquidity is more relevant to market order traders whose objective is to obtain immediacy at a low cost. The time measure of liquidity is more relevant to limit order traders, who supply liquidity on the SSM. In a setting where limit orders provide immediacy and wait for order execution, liquidity is measured by the expected time to execute a limit order at a given price, and more generally, by the probability of limit order execution. In this section, we examine these issues.

6.1. Order duration given limit order characteristics

The duration of an order, T_i, is the length of time until the order is executed, canceled or expired. In this subsection, we analyze order duration data using survival analysis. This statistical technique is very suitable for modeling order duration, since order durations are non-negative and random. This statistical technique allows us to estimate the conditional distribution of limit order execution times, T_i, as a function of explanatory variables, x_i, such as order characteristics and state of the book, $F(T_i < t \mid x_i)$. $F(\cdot)$ is the CDF of the Weibull distribution, $1 - \exp(-\lambda_i t)^p$ and $\lambda_i = \exp(-x_i' b)$. [40] The parameters p and b can be estimated by maximum likelihood. Following Lo et al. (1997), we treat limit orders that are canceled or expired as censored observations. Ignoring the information in non-executed orders can bias the estimator of the conditional distribution of execution times.

We estimate the survival model for buy and sell limit orders. The set of regressors in x includes a constant, an aggressiveness indicator, order size, number of orders per package, the inside spread, order imbalance, shares in the book, prior market order, and a volatility measure. Following Harris (1996), we measure order aggressiveness by $1 - 2(A - P)/(A - B)$ for buy orders and the negative of this quantity for sell orders, where $A(B)$ denotes the first best ask (bid), and P is the limit order price. This measure assigns a value of one to market orders and less than one to limit orders. Limit orders placed at the

[40] We used the Weibull distribution because it allows for duration dependence. The hazard function for the Weilbull distribution can be monotonically increasing or decreasing in t depending on p.

quote have a value of –1, and the difference between the order price and the best quote on the same side increases as this value gets smaller.

The order imbalance variable is defined as $k = Q_b/(Q_b + Q_s)$, for buy orders and $(1 - k)$ for sell orders, where $Q_s(Q_b)$ is the number of shares offered at A (demanded at B) at the time of order entry. Shares in the book is the number of shares in the book ahead of the order which have a higher priority for execution. Prior market order is the ratio of the market orders that are initiated by the same side of the market to the total market orders submitted during the last half-hour. Following Lo et al. (1997), we use the ratio of trades (market orders) in the last half hour to trades in the last one hour as a proxy for high frequency changes in volatility. The sign of the coefficient on each explanatory variable indicates the direction of the effect of that variable on the conditional probability of executing limit orders and on the expected time to execution. We estimate the model using the pooled data for all 56 stocks in the sample.

Table 9 reports the estimates of the coefficients and the median of the Weibull distribution. The negative sign on the aggressiveness indicator implies

Table 9
Survival analysis results[a,b]

	Buy orders	Sell orders
Number of observations		
Number of observations	80257	89987
Censored observations	39.21%	44.35%
Parameter estimate for independent variable		
Constant	6.7729*	1.0914*
Price aggressiveness	−2.0145*	−0.4465*
Number of orders per package	−0.0024*	0.0317*
Number of shares	0.0958*	−0.0061*
Inside spread	0.4332*	0.1766*
Order imbalance	1.2208*	0.6988*
Shares in the book	0.0001*	0.0002*
Prior market order	1.2423*	3.3825*
Volatility	0.274*	0.9144*
P	0.3238*	0.2827*
Median duration	72.5764	36.4784

[a] This table reports the parameter estimates of the survival analysis of order durations using the Weibull model, $F(t|x) = 1 - \exp(-\lambda t)p$ and $\lambda = \exp(-x'b)$. The duration of an order is the length of time in minutes that the order stays active (firm) in the market. Censored observations are limit orders that are canceled or expired. The set of regressors in x includes a constant, an aggressiveness indicator, order size, number of orders per package, the inside spread, order imbalance, shares in the book, prior market order, and a volatility measure as they are defined in Section 6.1 of the body of the paper. The estimate of median duration is calculated using, $\log(2)^{1/p}\lambda$.

[b] Also the hypothesis that the difference between the coefficient estimates for buy and sell orders for each independent variable is equal to zero is rejected using the Wald test for all of the independent variables.

* Statistically significant at the 1% level.

that the more aggressive the limit order price the shorter is the expected time to execution. The result is a natural outcome of the price priority rule. The number of orders per package is used to measure the degree of trader activity in the market, where large numbers indicate more active traders. The effects of this variable and order size on the expected time to execution are small and change in sign. The inside spread has a positive effect on the expected time to execution. These results imply that orders placed when the spread is wide are more difficult to execute. A wider spread implies a higher transaction cost, which provides little incentive for market order traders to execute against the existing limit orders. Consistent with the prediction of the theoretical model of Handa et al. (1996), as the order imbalance increases in favor of the other side of the market, the expected time to execution increases. The sign of the estimated coefficient for the shares in the book variable is as expected. As the number of shares that have higher priority of execution increases, we anticipate an increase in the expected time to execution. However, the magnitude of the effect is nearly zero. The estimated coefficent for the variable of prior market order is positive. This indicates that if the same side of the market initiates most trades in the last half-hour, a longer time to execution is expected. The result is intuitive because rising (falling) markets reduce the probability of executing buy (sell) orders. Finally, the positive sign of the estimated coefficient for the volatility variable implies that a longer time to execution is expected when the market is more volatile.

The estimate of p is 0.32 and 0.28 for buy and sell orders, respectively. This means that the hazard function has negative duration dependence. That is, the likelihood of executing a limit order at time t, conditional upon duration up to time t, is decreasing in t. The longer a limit order waits for execution, the less likely it is that it will be executed within, say, the next trading period.

6.2. The probability of executing limit orders

When immediacy is available during a continuous trading session, a trader can trade with certainty using a market order and not a limit order. The probability of executing a limit order is always less than one. In this section, we analyze the probability of order execution using a logistic probability model. The dependent variable, y, is the execution indicator, which equals one if the order is executed and zero otherwise. The probability of execution is conditioned on a set of regressors, x, $\text{Prob}[y = 1|x] = \Lambda(x'b)$, where $\Lambda(\cdot)$ is the logistic cumulative distribution function. The marginal effect of x on the probability is $\Lambda(x'b)[1 - \Lambda(x'b)]b$. The set of regressors in x includes the same set used in survival analysis. After pooling the data for all stocks, we estimate the coefficient, b, and the marginal effect (the slope). The results are reported in Table 10. The marginal effect, $\Lambda(\cdot)$, is evaluated at the mean of the variable. Similar to the findings in the previous subsection, price aggressiveness has a

Table 10
Logit regression results[a,b]

Number of observations and R-squared values	Buy orders		Sell orders	
Number of observations	80257		89987	
Dependent variable equals one	47.28%		48.71%	
McFadden pseudo R-squared	0.4149		0.4149	
Parameter estimates for following independent variables	Coefficient	Slope	Coefficient	Slope
Constant	3.3114*		2.7978*	
Price aggressiveness	0.9419*	0.1448*	0.6327*	0.0946*
Number of orders per package	0.0063*	0.001*	0.0028*	0.0004*
Number of shares	−0.0637*	−0.0098*	−0.0323*	−0.0048*
Inside spread	−0.2052*	−0.0315*	−0.1947*	−0.0291*
Order imbalance	−0.7468*	−0.1148*	−0.157*	−0.0235*
Shares in the book	−0.0000*	−0.0000*	−0.0000*	−0.0000*
Prior market order	−0.5502*	−0.0846*	−1.2109*	−0.1811*
Volatility	−0.2434*	−0.0374*	0.2926*	0.0438*

[a] This table reports the results for the logit regressions, $[y|x] = \Lambda(x'b)$, where y is a dummy variable that is equal to one if the order is executed, and zero otherwise. The set of regressors in x includes a constant, an aggressiveness indicator, order size, number of orders per package, the inside spread, order imbalance, shares in the book, prior market order, and a volatility measure as they are defined in Section 6.1 of the text of this paper. $\Lambda(\cdot)$ is the logistic cumulative distribution function. The coefficient is the b estimate. The slope is the marginal effect of x on the probability of execution, as given by $\Lambda(x'b)[1-\Lambda(x'b)]b$, when $\Lambda(x'b)$ is evaluated at the mean of the regressors. McFadden pseudo R-squared is $1-(\ln L/\ln L_0)$, where $\ln L$ and $\ln L_0$ are the log-likelihood functions evaluated at the unrestricted and restricted estimates (all coefficients, except the constant, are zero), respectively.
[b] The hypothesis, that all the coefficients (except the constant term) are equal to zero, is rejected using both the Likelihood Ratio and Wald tests. The hypothesis, that the difference between the estimated coefficients for buy and sell orders for each independent variable is equal to zero, also is rejected using the Wald test for all but one variable (namely, the Number of shares).
* Statistically significant at the 1% level.

positive effect on the probability of execution. Overall, limit orders with "reasonable" prices are highly liquid in terms of executability. The results also indicate that more active traders have higher probabilities of execution. Active traders frequently have standing firm orders at and away from the quote either to make a market, or to seize the free option quickly. Since they also monitor the market more closely, we expect them to adjust their exposed orders more frequently than others. The negative signs on the estimated coefficients for the order size variable suggest that larger orders are more difficult to execute. The signs of the estimated coefficients of the volatility measure variable imply that sell (buy) orders have higher (lower) probabilities of execution when market conditions are more active. This could be attributed to the 9.23% rise in the market index over the sample period. This rising market period also may explain the significant differences between buy and sell orders identified in Tables 9 and 10. Other regression results in Table 10 are generally consistent with the findings that we obtained from the survival analysis.

7. Concluding remarks

In this paper, we describe and analyze the microstructure of the SSM under the computerized trading system, ESIS. We analyze the order book, order flow and order execution using four rich data sets on orders, order packages, quotes and transactions. Although the SSM has a distinct structure, its intraday patterns are surprisingly similar to those found in other markets with different structures. These include U-shaped patterns in traded volume, number of transactions and volatility. Like other order-driven markets, the SSM exhibits a U-shaped pattern in the placement of new orders.

We find that the relative inside spread is higher only at the open and declines gradually afterwards on the SSM. This pattern is similar to the one observed for a number of markets without designated market makers. We find that the average relative inside spread is large compared to other markets, mainly due to a relatively high tick size. Tick size is an important determinant of the inside spread for low priced stocks, and for all other relative spreads. As in other studies, we detect a "diagonal effect" in order flow. Strategic order splitting rather than imitation or competition hypothesis appears to be the dominant factor causing this effect.

We find that liquidity, as commonly measured by width and depth, is relatively low on the SSM. However, it is exceptionally high when measured by immediacy. We also present new evidence on other measures of market liquidity that are more relevant to order-driven markets. For example, we find that limit orders that are priced reasonably, on average, have a shorter expected time to execution, and have a high probability of subsequent execution.

Acknowledgements

Financial support from FCAR (Fonds pour la formation de chercheurs et l'Aide à la Recherche), SSHRC (Social Sciences and Humanities Research Council of Canada) and the Imam University are gratefully acknowledged. We appreciate the comments of Anastasios Anastasopoulos, Abraham Brodt, Gordon Fisher, Michael Sampson and two referees of this journal on earlier versions of this paper. We would like to thank the Saudi Arabian Monetary Agency for assistance in obtaining the data used in this paper. All remaining errors are the authors' responsibility. Comments are welcomed.

References

Admati, A., Pfleiderer, P., 1988. A theory of intraday patterns: Volume and price volatility. The Review of Financial Studies 1 (1), 3–40.

Amihud, Y., Mendelson, H., 1989. The effect of computer based trading on volatility and liquidity. In: Lucas, H., Schwartz, R.A. (Eds.), The Challenge of Information Technology for the Securities Markets: Liquidity, Volatility, and Global Trading. Dow Jones-Irwin, Homewood, Illinois.

Angel, J., 1997. Tick size, share prices, and stock splits. Journal of Finance 52, 655–681.

Biais, B., Hillion, P., Spatt, C., 1995. An empirical analysis of the limit order book and the order flow in the Paris Bourse. Journal of Finance 50, 1655–1690.

Brock, W., Kleidon, A., 1992. Periodic market closure and trading volume: A model of intraday bids and asks. Journal of Economic Dynamics and Control 16, 451–489.

Butler, K., Malaikah, S., 1992. Efficiency and inefficiency in thinly traded stock markets: Kuwait and Saudi Arabia. Journal of Banking and Finance 16 (1), 197–210.

Chan, L., Lakonishok, J., 1995. The behavior of stock prices around institutional trades. Journal of Finance 50, 1147–1174.

Chan, K., Christie, W., Schultz, P., 1995a. Market structure and the intraday pattern of the bid–ask spreads for NASDAQ securities. Journal of Business 68 (1), 35–60.

Chan, K., Chung, P., Johnson, H., 1995b. The intraday behavior of bid–ask spreads for NYSE stocks and CBOE options. Journal of Financial and Quantitative Analysis 30 (3), 329–346.

Darwiche, F., 1986. The Gulf Stock Exchange: The rise and fall of the Souq Al-Manakh. Croom Helm, London.

Foster, F., Viswanathan, S., 1993. Variations in trading volume, return volatility and trading costs: Evidence on recent price formation models. Journal of Finance 48, 187–211.

Gerety, M., Mulherin, H., 1992. Trading halts and market activity: An analysis of volume at the open and the close. Journal of Finance 47, 1765–1784.

Handa, P., Schwartz, R., Tiwari, A., 1996. The determinants of bid–ask spread in an order-driven market, Working paper.

Harris, L., 1986. A transaction data study of weekly and intradaily patterns in stock returns. Journal of Financial Economics 16, 99–117.

Harris, L., 1990. Liquidity, trading rules, and electronic trading systems, New York University Salomon Center, Leonard N. Stern School of Business, New York.

Harris, L., 1996. Does a large minimum price variation encourage order exposure, Working paper, University of South California.

Hedvall, K., 1995. Trade concentration hypotheses: An empirical test of information vs. demand models on the Helsinki Stock Exchange. Journal of International Financial Markets, Institutions and Money 3, 135–163.

Hedvall, K., Niemeyer, J., 1996. Order flow dynamics: Evidence from the Helsinki Stock Exchange, Working paper, Swedish School of Economics and Business Administration.

Hedvall, K., Niemeyer, J., Rosenqvist, G., 1997. Do buyers and sellers behave similarly in a limit order book? A high-frequency data examination of Finnish Stock Market. Journal of Empirical Finance 4, 279–293.

Jain, P., Joh, G., 1988. The dependence between hourly prices and trading volume. Journal of Financial and Quantitative Analysis 23 (3), 269–282.

Lehmann, N., Modest, D., 1994. Trading and liquidity on the Tokyo Stock Exchange: A bird's eye view. Journal of Finance 49, 951–984.

Lippman, S., McCall, J., 1986. An operational measure of liquidity. American Economic Review 76 (1), 43–55.

Lo, A.W., MacKinlay, A.C., Zhang, J., 1997. Econometric models of limit-order executions. Working paper no. W6257, Cambridge, MA.

Madhavan, A., 1992. Trading mechanism in securities markets. Journal of Finance 47, 607–641.
Malaikah, S., 1990. Saudi Arabia and Kuwait: A study of stock market behavior and its policy implications, Ph.D. dissertation, Michigan State University.
McInish, T., Wood, R., 1990. An analysis of transactions data for the Toronto Stock Exchange. Journal of Banking and Finance 14, 441–458.
McInish, T., Wood, R., 1991. Hourly returns, volume, trade size, and number of trades. Journal of Financial Research 14 (4), 303–315.
McInish, T., Wood, R., 1992. An analysis of intraday patterns in bid–ask spreads for NYSE stocks. Journal of Finance 47, 753–764.
Niemeyer, J., Sandås, P., 1995. An empirical analysis of the trading structure at the Stockholm Stock Exchange, Working paper, Stockholm School of Economics.
Stoll, H., 1992. Principles of trading market structure. Journal of Financial Research 6, 75–107.
Wilson, Peter W, 1991. A question of interest: The paralysis of Saudi banking. Westview, Boulder, CO.
Wood, R., McInish, T., Ord, J., 1985. An investigation of transactions data from NYSE stocks. Journal of Finance 40, 723–741.