ELSEVIER

Journal of Financial Markets 3 (2000) 333–363

Journal of FINANCIAL MARKETS

www.elsevier.nl/locate/econbase

Market structure, informational efficiency and liquidity: An experimental comparison of auction and dealer markets[☆]

Erik Theissen*

Johann Wolfgang Goethe-Universität Frankfurt/Main, Lehrstuhl für Kreditwirtschaft und Finanzierung, Mertonstr. 17-21 (PF 88), 60054 Frankfurt, Germany

Abstract

We report the results of 18 market experiments that were conducted in order to compare the call market, the continuous auction and the dealer market. Transaction prices in the call and continuous auction markets are much more efficient than prices in the dealer markets. The call market shows a tendency towards underreaction to new information. Execution costs are lowest in the call market and highest in the dealer market. The trading volume and Roll's (Journal of Finance (1984) 1127–1139) serial covariance estimator are inappropriate measures of execution costs in the present context. The relation between private signals, trading decisions and trading profits is analyzed. © 2000 Elsevier Science B.V. All rights reserved.

JEL classification: C90; G14

Keywords: Market microstructure; Informational efficiency; Liquidity, Bid–ask spread; Experimental asset markets

☆Earlier versions of this paper were circulated under the title "Market Structure and the Aggregation of Information: An Experimental Investigation". Helpful comments were made by an anonymous referee, by Yakov Amihud, Uri Ben-Zion, Wolfgang Bühler, Lawrence Harris, Christian Rieck, Ailsa Röell, Dirk Schiereck, Reinhard Tietz and seminar participants at the universities of Amsterdam, Berlin (HU), Frankfurt/Main, Giessen and Karlsruhe and at the conference on Organization and Quality of Equity Markets at the Paris Stock Exchange. I am particularly indebted to Jan Krahnen and Martin Weber for valuable advice in all stages of this research. Financial support from Deutsche Forschungsgemeinschaft under grant Kr 1221/3-1 is gratefully acknowledged.

* Tel.: + 49-69-79828429; fax: + 49-69-79828951.

E-mail address: theissen@wiwi.uni-frankfurt.de (E. Theissen).

PII: S1386-4181(00)00010-0

1. Introduction

Growing competition forces stock exchanges to react to customer needs. The design of the trading mechanism is the most important determinant of market performance. The central issues are the computerization of the trading process and the choice between an order-driven and a quote-driven mechanism. This paper focusses on the second issue. Its purpose is to compare the principal alternatives, namely, the call auction, the continuous auction and the competitive dealer market.

Recent changes that occured in major European stock markets indicate that the issue of which trading mechanism is best suited is far from being resolved. For example, the London Stock Exchange replaced its quote-driven trading system with the electronic order-driven system SETS in 1997. On NASDAQ, public limit orders now compete with dealer quotations. In France and Germany, on the other hand, dealers were introduced to provide additional liquidity to the electronic continuous auction markets NSC and XETRA, respectively.

Continuous trading for less liquid stocks in the French CAC system (the predecessor of NSC) was replaced with two daily call auctions in 1992. On the other hand, the stocks listed on the French Nouveau Marché were transferred from a call market to the electronic continuous auction system in 1998. The London Stock Exchange considered to have less liquid stocks traded on SETS while, almost at the same time, Deutsche Börse AG announced that, for a number of less liquid stocks, continuous trading in XETRA would be replaced with call auctions. Call auctions are also frequently used to establish opening prices whereas closing call auctions and intradaily call auctions are less common.

This evidence suggests that more empirical research into the relative advantages of the principal trading mechanisms is needed. The existing empirical literature has mainly focussed on comparing the liquidity of continuous auction and dealer markets. The issue of informational efficiency has rarely been addressed. This is, to a large part, due to data limitations.

These limitations can be overcome in experimental research. In this paper we report the results of a series of 18 market experiments with a total of 216 participants. The experimental method allows to vary the trading mechanism under ceteris-paribus conditions. Different results can therefore be attributed to the design of the trading system. The asset value and the information each market participant holds is known to the experimenter. This allows to directly address the issue of informational efficiency. The experiments described in the present paper are designed in a way that enhances the comparability of the results to those obtained from field studies.

The findings can be summarized as follows. Consistent with the practice of many exchanges to start trading with a call auction, opening prices in the call market are closer to the true value of the asset than opening prices in the continuous markets. The difference to the dealer market is significant whereas the difference to the continuous auction market is not. Generally, transaction prices in the call and continuous auction markets are closer to the asset value than are transaction prices in the dealer market. The results for the latter are largely attributable to high bid–ask spreads. Once the spread is eliminated from the data by averaging prices or by using midquotes instead of prices, the information provided by the time series of prices in the dealer market is very precise. Analysis of the determinants of price changes reveals that prices in the call market underreact to new information. No such pattern is found in the price series from the continuous trading mechanisms.

Transaction costs as measured by the bid–ask spread are lowest in the call market and highest in the dealer market. This result is consistent with evidence obtained from real world stock markets. Both the trading volume and Roll's (1984) serial covariance estimator are shown to be inadequate measures of execution costs in the present context. Higher trading volume is associated with higher rather than lower spreads. The Roll measure systematically overstates the execution costs in the call market.

The quality of the private information traders receive influences the composition of individual portfolios and the trading volume. A relation between signal quality and profits is only found in the dealer market.

The paper is organized as follows. Section 2 provides a brief summary of the literature, Section 3 describes the experimental design and procedures, Section 4 presents the experimental results and Section 5 concludes.

2. Previous research

In a call auction, market and limit orders are batched and are executed at discrete points in time. The auctioneer (or a computer algorithm) sets a price that clears the market, and all orders that are executable are filled at that price. Among the advantages of this trading mechanism are the high liquidity due to the temporal consolidation of the order flow, the non-existence of an explicit bid–ask spread and the lower market impact of large orders. The simultaneous execution of a large number of orders may lead to better price discovery. However, this advantage is partially offset by the fact that no information is conveyed between the calls. The major drawbacks of the call market are the lack of immediacy and the impossibility to simultaneously insure against price and execution risk.

In a continuous auction market participants may submit buy and sell orders or accept orders of other participants at any point in time. Liquidity is provided by the voluntary exposure of limit orders to the market. Liquidity suppliers are compensated by the bid–ask spread which, in turn, is the price of immediacy that liquidity demanders have to pay. As no trader or group of traders is obliged to submit limit orders, immediate execution of market orders is not guaranteed.

Dealer markets rely on designated market makers as suppliers of liquidity. They are obliged to continuously quote bid and ask prices and thus provide a possibility for immediate order execution at predetermined prices. As in a continuous auction the bid–ask spread is the compensation for the cost of this service. Pure dealer markets rely on competition between multiple market makers to keep the price for the immediacy service close to its cost.

Most existing theoretical research concentrates on modelling one trading mechanism (see e.g. Mendelson, 1982; Ho et al., 1985; Satterthwaite and Williams, 1993; Rustichini et al., 1994 for models of call auctions, Friedman, 1984, 1991; Wilson, 1987; Easley and Ledyard, 1993; Glosten, 1994 for continuous auction markets and O'Hara, 1995 for a survey on models of dealer markets).

These models use specific modelling approaches and sets of assumptions. This makes it difficult to compare the results. There have, however, been several attempts at comparing different trading mechanisms in a unified framework. Building on the model developed in Mendelson (1982), Mendelson (1987) finds that gross gains from trade are higher in a centralized call market than in a competitive dealer market but that communication costs are also higher. Kyle (1985) models the interaction of a single insider, random noise traders and competitive market makers. He shows that noise trader losses under a continuous trading regime are twice as large as in a single auction equilibrium. Pagano and Röell (1992) show that trading costs are lowest in a call market and highest in a dealer market. Pagano and Röell (1996) focus on the degree of transparency as the main difference between trading mechanisms. They confirm the finding that uninformed traders' transaction costs are higher in the dealer market. Madhavan (1992) derives noisy rational expectations equilibria for a call market, a continuous auction and a dealer market. He shows that the call market is more robust in the presence of severe information asymmetries.

Most empirical comparisons of different trading mechanisms focus on market liquidity. In continuous auction and dealer markets explicit bid–ask spreads exist and can be used as a measure of liquidity. Spreads have generally been found to be lower in continuous auction markets (e.g. Bessembinder and Kauffman, 1996; Christie and Huang, 1994; de Jong et al., 1995; Huang and Stoll, 1996; Keim and Madhavan, 1996; Lee, 1993; Pagano and Röell, 1990; Stoll, 1993). Barclay et al. (1999) and Bessembinder (1999) document that spreads on NASDAQ fell after competition by limit orders has been introduced in 1997. The trade execution costs remain larger, however, than those on the NYSE (Bessembinder, 1999).

In a call market no explicit bid–ask spread exists. Some authors have therefore used the serial covariance measure developed by Roll (1984) to calculate implied spreads for call markets and to compare transaction costs in call and continuous markets (Haller and Stoll, 1989; Pagano and Röell, 1990; Stoll and Whaley, 1990). However, as will be shown in Section 4.2, the Roll measure is likely to overstate the execution costs in the call market. Amihud et al. (1997) find that the introduction of continuous trading on the Tel Aviv Stock Exchange was associated with a positive abnormal return which they attribute to the higher liquidity of the continuous trading system. They also find evidence of an increase in market efficiency.

Summarizing the empirical results, it can be concluded that there is sound evidence of higher execution costs in dealer markets as compared to continuous auction markets. No clear picture emerges as to how the call market compares to these continuous trading mechanisms. Furthermore, little can be said about the informational efficiency of the different trading mechanisms. This is in large part due to the fact that no direct measure of informational efficiency is available because the true value of an asset is unobservable.

This problem can be overcome in experimental asset markets because the asset value and the information each market participant holds is known to the experimenter. Identical assets can be traded in different environments. There is thus no need to control for different asset characteristics by, e.g., constructing a matched sample. In addition, there is no potential for a selection bias. In the field such a bias may exist because it is possible that certain assets are traded under a particular market regime precisely because this regime is best suited for the asset. Finally, no brokerage fees and no differences in the quality of the service provided (e.g., clearing and settlement) exist. This eliminates another potential source of bias present in field studies. Therefore, experimental asset markets are well suited for comparisons of different trading mechanisms.

Experiments relating to financial markets usually incorporate information asymmetries. A conveniant way to achieve this is to define a small number of states of nature that define the assets' payoffs to traders (Plott and Sunder, 1982). Traders then are given different access to information about the actual state. There are two principal ways to introduce asymmetric information. The first is to provide some traders with superior information as is done in most information-based microstructure models (e.g. Glosten and Milgrom, 1985; Kyle, 1985; Admati and Pfleiderer, 1988). The alternative is to provide all traders with signals of identical ex-ante quality. This is termed heterogeneous information in the sequel. This way of modelling asymmetric information is used in the rational expectations literature (e.g. Diamond and Verrecchia, 1981; Hellwig, 1980) and in some microstructure models (e.g. Kyle, 1989).

Using a design that incorporates asymmetric information, Friedman (1993a) finds that the continuous auction yields higher allocational efficiency but lower market depth than the call market. Informational efficiency is similar for the two trading mechanisms. Friedman (1993b) compares the continuous auction with the dealer market. He concludes that the latter is less efficient and less deep. The market maker position turns out to be profitable.

In the experiments just cited traders hold private valuations for the asset. This is achieved by dividing the traders into two (or more) groups which receive different payoffs for the assets they hold in the *same* state of nature. This allows the analysis of the allocational efficiency of the market because the group with the highest valuation in the given state of nature should eventually hold all assets. Although different valuations may be justified by different tax brackets or different background risk, a common valuation by all traders seems to be more appropriate for market microstructure experiments.[1] The first to use a common value design were Smith et al. (1988). In their continuous auction markets with symmetric information they find evidence of significant and persistent mispricing. Van Boening et al. (1993) get similar results from a series of experimental call markets.

Schnitzlein (1996) provides an experimental implementation of the Kyle (1985) model.[2]

Comparing call and continuous auction trading in the presence of a single informed insider, he finds that the call market provides higher liquidity and lower costs to uninformed noise traders at the expense of slightly less efficient prices.

Krahnen and Weber (1999, 2000) compare the continuous auction with the dealer market. They find that information aggregation in the dealer market is at least as good as in the continuous auction. They also vary the number of market makers and find that a larger number of market makers is associated with lower spreads and higher trading volume.[3]

To summarize, the experimental evidence so far suggests that call markets are not inferior to continuous markets. However, a detailed experimental comparison of call, continuous auction and dealer markets has, with the exception of Friedman (1993a, b), not yet been attempted. In addition, a heterogeneous information design as defined above has not yet been used for comparisons of different trading mechanisms. The present paper fills this gap by reporting results of 18 market experiments, six for each trading mechanism. The design incorporates heterogeneous information and thus focusses on the aggregation of diverse information rather than on insider trading.

[1] The use of a common value design is not without a cost. A private value design induces a rational trading motive because assets held by traders with low valuations should be sold to traders with high valuations. In a common value design (like the one used for the experiments described in this paper) this trading motive is absent. Different degrees of risk aversion, however, still constitute a rational trading motive.

[2] Although capturing basic features of call and dealer markets, the trading mechanisms modeled by Kyle (1985) and implemented in the experiments of Schnitzlein (1996) are not identical to those encountered in real-world markets. Only market orders are allowed and prices in the auction market are set by market makers rather than determined by the interaction of the investors.

[3] Some recent papers (e.g. Bloomfield, 1996; Bloomfield and O'Hara, 1999, 2000; Flood et al., 1997, 1998, 1999; Lamoureux and Schnitzlein, 1996) report results of experiments that were designed to address specific microstructure issues like market transparency, market maker behavior, and off-exchange trading. In these papers, no attempt is made at comparing call and continuous markets.

3. Experimental design and procedures

We conducted a total of 18 experiments. In each experiment 12 subjects traded one risky asset. Each subject participated in only one experiment.[4] An experiment consists of a sequence of 15 trading periods. The first three periods were used to train the subjects and are excluded from the analysis. Each trading period lasts 5 min.[5] The asset value is identical for all subjects (i.e., we use a common value design) and changes from period to period. Prior to the start of a period each subject privately receives a noisy signal about the asset value. This and the distribution the signals are drawn from are common knowledge.

The 18 experiments can be divided into six series of three experiments. In each series the realizations of the asset value process and the private signals are held constant while the trading mechanism is varied. This procedure assures a maximum of comparability between the experiments conducted with the different trading mechanisms.[6]

Each trader received an initial endowment of 100 shares and 50.000 currency units in cash. In addition, there was a credit limit (at a zero interest rate) of 50.000 and short sales up to 100 assets were allowed. The endowments were reinitialized at the beginning of each trading period. This has three advantages. First, traders who sold all their assets or spent all their cash in one period are not restricted to acting only on one side of the market in the subsequent period. Second, the equality of the endowments at the beginning of each period further enhances the comparability between the experiments. Third, reinitialization induces a trading motive because subjects have to balance their portfolio according to their risk preferences in each period.

3.1. Information structure

Denote the asset value in period t by v_t. The value for the first period was set by the experimenter. In subsequent periods the value was determined randomly by a draw from an uniform distribution of the integer values in the intervall $[0.8v_{t-1}; 1.2v_{t-1}]$. The value change between two periods is thus limited to $\pm$ 20%. Each trader receives a private noisy signal about the asset value prior

[4] An alternative approach is to have several groups (cohorts) of subjects trade in all environments and to systematically vary the order in which the cohorts are confronted with the treatments. Given a duration of approximately 3 hours for an experiment and given that we held the asset value constant across treatments (which would not have been feasible had the same subjects participated), we decided against this alternative.

[5] An exception is the first call market experiment where the periods lasted 6 min.

[6] One feature of this design is that the amount of information about the asset value is constant across trading institutions. If the trading institution itself affects the incentives to acquire information, this condition may cease to hold once the decision to acquire information is endogenized.

to the start of each period. The signals are independent draws from an uniform distribution of all integer values in the intervall $[0.9v_t; 1.1v_t]$. The signal trader i receives in period t is denoted by $s_{i,t}$. Given the signal a trader can calculate an upper and a lower bound for the asset value:

$$s_{i,t}/1.1 \leqslant v_t \leqslant s_{i,t}/0.9. \tag{1}$$

Due to the stochastic nature of the signals, the average signal $\bar{s}_t = (1/12)\sum_i s_{i,t}$ may deviate from the value v_t. In the analysis of the series 1–3 the value v_t was replaced by $\bar{s}_t$. In series 4–6 the signals were determined in a way that guarantees $\bar{s}_t = v_t$.[7]

The determination of the individual signals assures that all traders receive information of identical ex-ante quality whereas the realizations in a specific trading period are of different quality. The design thus abstracts from the existence of market participants with systematically superior information. This allows to concentrate on the performance of the different trading mechanisms and their ability to aggregate diverse information.

The asset value was not revealed to the participants after completion of a trading period for the following reasons. First, true values are also unobservable in real world markets. Second, revealing the value precludes the analysis of the persistence of pricing errors.

3.2. The trading mechanisms

All experiments were conducted as computerized market experiments. The software used allows to implement the different trading mechanisms that are of interest for the present study. In what follows the details of the trading mechanisms are described.

3.2.1. Call market

Each trading period in the call market experiments was divided into two subperiods of equal length. There was one call in each subperiod. The resulting prices are referred to as the opening and the closing price, respectively.

Traders were allowed to submit buy and/or sell orders. The only restriction was the budget constraint described above. Only limit orders were allowed. A market buy or sell order could, however, be mimicked by setting the price limit sufficiently high or low, respectively. The minimum tick size was set to one,

[7] Replacing v_t with the average signal assumes that the arithmetic mean of the signals is a sufficient summary statistic for the individual signals. This is not true since the cross-sectional distribution of the realizations of the signals provides additional information. This is, however, ignored in the analysis. Repeating the analysis with the true values instead of the average of the signals does not yield different conclusions. The procedure used to generate the signals for experiments 4–6 consisted in reinitialising the random number generator used for the determination of the signals until the average signal, rounded to the nearest integer, was equal to the value. This procedure does not interfere with the random character of the signals from the point of view of the participants.

i.e., the price could only take on integer values. This minimum tick size is comparable to those in real-world markets. Even for the lowest prices observed in the experiments — well above 300 — a minimum tick size of one is equivalent in percentage terms to a tick size of 1/8 for a stock trading at 40.

The computer aggregates the orders and determines a market clearing price such that the share trading volume is maximized. If this results in an interval of possible prices the midpoint of the interval is chosen. If the price does not completely clear the market, rationing takes place. In this case the following priority rules are used: Buy orders with higher price limits and sell orders with lower price limits are filled with priority. If limit prices are equal then the order which specifies the higher quantity gets priority. If quantities are equal, orders are selected randomly for rationing.

The order book is closed and no indicative prices are provided. Therefore, traders' information sets are restricted to their private signals and the history of transaction prices and trading volumes. This is the situation encountered by retail investors without direct access to the trading floor.

3.2.2. Continuous auction

In the continuous auction markets only the best bid and the best ask are displayed. New orders are only accepted when they raise the bid or lower the ask. The standing bid and ask can be accepted at any point in time. Order execution follows price and time priority.

To submit or accept an order, subjects have to click a "buy" or "sell" button, respectively. Price and quantity of the standing ask or bid are displayed automatically. If the subject wants to accept the order she just has to confirm. Otherwise she can accept the price but lower the quantity or change the price. In the latter case she submits an own order that replaces the standing best bid or ask.

The subjects' information sets consist of their private signals and the history of best bids, best asks, and transaction prices and volumes. Trading is anonymous, i.e. the identity of the trader who submitted an order is never revealed.

3.2.3. Dealer market

The dealer market is a variation of the continuous auction market. The number of market makers is set to three in order to assure sufficient competition. Only the market makers are allowed to submit limit orders. The other traders can only accept the prices set by the market makers. The market maker position was assigned to three randomly chosen subjects.

The short sale and credit restrictions for the market makers are relaxed. Each of them had a credit limit of 250.000 and could short sell 500 assets. This assures that the market makers can accomodate the trading demand of the nine other traders.

The three market makers receive information of the same quality as the other traders. This design feature stands in contrast to most information-based models of market maker behavior. However, these models usually assume the existence of informed insiders whereas the present paper focuses on the aggregation of different signals of identical ex-ante quality.

3.3. Experimental procedures

The experiments were conducted at the University of Frankfurt between February and May 1996. Participants were students of economics and business administration. An experiment lasted approximately three hours. Subjects received a set of written instructions.[8] After having read these, each participant had to answer a short test that contained a couple of questions relating to the instructions. The experiment did not start until each subject had correctly answered these questions.

Participants were rewarded for their participation in the experiment. One of the periods 4–15 was chosen randomly. The value of the end-of-period portfolio was transformed into Deutsche Mark and paid out. The average payoff for a 3 hour experiment was DM 48.03, individual payoffs ranged from DM 29.42 to DM 84.58.[9]

After the last period participants had to answer a questionnaire containing questions related to the instructions, to the experimental design and to the strategies the subjects employed and thought to be optimal.

4. Experimental results

This section presents the results of the experiments. Only the periods 4–15 are analyzed because the first three periods were training periods that were not rewarded.[10] Results are presented in three parts. The first section deals with the

[8] The instructions are available upon request.

[9] Situations may exist in which experimental subjects engage in very risky actions because the experimenter cannot credibly collect the losses they might incur. The figures indicate that the probability of incurring losses was low enough to largely eliminate this incentive.

[10] In the continuous auction and dealer markets some transactions occurred at "unrealistic" prices. The most likely reason are typing errors. On the one hand, one would like to eliminate these transactions. On the other hand, however, it is difficult to judge which transactions are due to typing errors and which are not. For the following analysis all transactions at prices that fulfilled one of the following conditions were excluded: $p > Max_i(s_{i,t}/0, 9)$ or $p < Min_i(s_{i,t}/1, 1)$. These prices are above the upper bound or below the lower bound for the asset value for *all* traders. Therefore, all traders would agree that this transaction price cannot be rationally justified. Only if a trader expects to be able to close the position at a more favorable price later in the period, transactions at very low or high prices may be rationalized. However, in this case one would expect the transaction to be part of a sequence of transactions at rising or declining prices. This has not been the case. Furthermore, since endowments were reinitialized after each period, expectations of higher or lower asset values in subsequent periods can also not be used to justify the transactions in question. According to the above criterion, 52 transactions were eliminated. This corresponds to an average of 0.36 transactions per trading period.

informational efficiency of the prices, the second presents results on market liquidity and in the third we analyze how the quality of the private signals influences the traders' portfolio structures and trading profits.

4.1. Informational efficiency

In experimental asset markets measures of informational efficiency can be based on a comparison of asset values and market prices because the true asset value is known. We use three related measures of efficiency. The first is the mean absolute deviation between transaction price and asset value,

$$MAE = \frac{1}{12}\sum_{t=4}^{15} |p_{t,j} - v_t|; \quad j \in o, a, c \tag{2}$$

The index j characterizes the price. All measures are calculated for the opening price (o), the closing price (c) and the volume-weighted average of all prices of the period (a).[11] In the call market the latter is a weighted average of the opening and closing price since there are only two prices in a period. The second measure is the mean relative error,

$$MRE = \frac{1}{12}\sum_{t=4}^{15} \frac{|p_{j,t} - v_t|}{v_t}; \quad j \in o, a, c. \tag{3}$$

Here, the absolute deviation between price and value is divided by the asset value before it is averaged over the periods. This measure makes experiments with different asset value levels comparable. The last measure is the root mean squared error,

$$RMSE = \sqrt{\frac{1}{12}\sum_{t=4}^{15} (p_{j,t} - v_t)^2}. \tag{4}$$

The RMSE measure gives more weight to larger deviations. Table 1 shows the results. The rankings of the trading mechanisms produced by the three measures are very similar. The only deviations occur in sessions two, three and five.

4.1.1. Descriptive Statistics

The figures in Table 1 were obtained by averaging over the 12 periods of each experiment. Before proceeding it is therefore worthwhile to ask whether this aggregation masks systematic patterns which may be caused by learning effects.

[11] Note that calculating the deviation between the asset value and the weighted average of the prices of a period is not the same as the average of the deviation between the prices and the asset value. Most importantly, the latter measure incorporates the bid–ask spread whereas the former largely eliminates the spread by averaging over all prices.

We performed two tests to answer this question. First, we compared the MRE measure of informational efficiency (separately for the opening, the closing, and the weighted average prices) for the first half and the second half of the experiments using non-parametric Wilcoxon tests. We found no evidence of significantly higher efficiency in the second half of the experiment. Second, we used a Page test to test the null hypothesis that the pricing errors for different periods are drawn from the same distribution against the alternative of monotonically decreasing pricing errors (with at least one strict inequality). The null hypothesis was rejected in only one of the nine cases (closing prices in the call market). From these results we conclude that there are no significant learning effects beyond the three initial training periods and that it is therefore safe to pool the data from different periods in the following analysis.

The first question to ask is whether information aggregation takes place at all. With no aggregation, prices can be expected to have the same precision as the private signals. Table 1 shows the average precision of the signals for each of the measures discussed above. It is apparent that aggregation takes place in the call markets and (with the exception of opening and closing prices in series 4) in the continuous auction markets. In the dealer markets the volume-weighted average prices clearly aggregate information whereas the opening and closing prices do not. These results are confirmed by a non-parametric Wilcoxon test. All price series with the exception of opening and closing prices in the dealer market have a median MRE which is smaller than the median MRE of the signals on at least the 10% level of significance. The results for the dealer market will be discussed in more detail in the sequel.

The figures in Table 1 document that opening prices in the call market are more precise than those in the dealer market in all six sessions and more precise than those in the continuous auction in four (MAE and MRE measure) or five (RMSE) sessions. The greater accuracy of the call market prices is likely to be a consequence of the temporal consolidation of the order flow.

Comparing the measures for the opening and the closing prices reveals that there is no evidence of increasing price efficiency in the course of the trading period in both the call market and the continuous auction. In the dealer market, on the other hand, efficiency increases. Despite this increase, however, closing prices in the dealer markets are the least accurate whereas closing prices are most accurate in the continuous auction.

One striking observation concerning the dealer market is the large discrepancy between the pricing errors for the opening and closing prices on the one hand and the volume-weighted average price on the other hand. Irrespective of the measure used, the average price is more efficient than either the opening or the closing price. However, this finding can easily be explained. As will be shown later, bid–ask spreads in the dealer markets are large. If a single transaction is used for the calculation of the pricing error then the measured error on average

Table 1
Informational efficiency: descriptive analysis

Table 1 contains measures of price quality in the experimental markets. MAE is the mean absolute pricing error. MRE is the mean of the absolute pricing errors, divided by the asset value. RMSE is the root mean squared error. The rows denoted "signals" give the average quality of the private signals. CM is the call market, DA the continuous (double) auction and MM the dealer (market maker) market. Open and close refers to the first, and last, transaction of each period, respectively. Av. is the average of all transaction prices in a period, weighted by the trading volume. A shaded cell marks the lowest deviation between price and value within each series and for each of the price series.

Experiment		MAE			MRE in %			RMSE		
Series	Institution	Open	Av.	Close	Open	Av.	Close	Open	Av.	Close
1	Signals		21.73			4.78			25.44	
1	CM	12.84	13.13	15.80	2.95	2.92	3.43	16.33	16.53	19.98
1	DA	15.33	14.29	17.04	3.27	2.94	3.67	17.40	18.43	22.09
1	MM	29.01	8.80	22.45	6.20	1.79	4.83	34.50	11.67	26.88
2	Signals		23.05			4.93			26.60	
2	CM	10.58	10.17	13.76	2.24	2.11	2.74	13.28	12.44	16.83
2	DA	6.99	12.30	10.13	1.51	2.63	2.08	10.09	14.82	14.19
2	MM	29.09	11.05	15.61	6.18	2.33	3.26	34.32	11.93	20.82
3	Signals		21.16			4.89			23.99	
3	CM	16.67	15.25	13.5	3.88	3.56	3.14	18.04	16.57	15.00
3	DA	15.57	9.39	11.31	3.58	2.18	2.62	18.77	12.27	13.94
3	MM	17.72	9.51	16.31	4.07	2.17	3.66	22.35	11.29	20.14
4	Signals		22.86			5.01			26.15	
4	CM	17.08	16.92	19.67	3.63	3.58	4.14	21.35	21.87	25.54
4	DA	23.42	15.33	24.42	5.07	3.33	5.31	30.03	19.38	30.58
4	MM	21.83	9.25	16.08	4.87	2.00	3.50	25.60	11.74	21.40
5	Signals		18.90			4.67			21.96	
5	CM	9.08	9.08	9.5	2.28	2.29	2.41	11.66	11.43	11.64
5	DA	12.83	6.17	5.92	3.44	1.48	1.41	18.61	8.54	8.43
5	MM	25.92	8.25	13.17	6.27	2.06	3.24	30.87	11.70	16.17
6	Signals		25.14			5.03			28.68	
6	CM	12.50	12.42	12.67	2.57	2.55	2.59	15.59	15.51	15.65
6	DA	13.58	7.67	10.92	2.66	1.55	2.29	17.12	9.57	13.18
6	MM	45.5	11.25	27.75	9.25	2.29	5.75	52.39	14.84	32.14
av.	Signals					4.89				
av.	CM				2.93	2.84	3.08			
av.	DA				3.26	2.35	2.90			
av.	MM				6.14	2.11	4.04			

contains the half spread. Calculation of the average price largely eliminates the spread and thus leads to lower pricing errors.

In both theoretical and empirical microstructure research it is often assumed that the midpoint of the bid and the ask quote is the best available estimate of the true asset value. We therefore calculated the measures of price efficiency on the basis of midquotes for both the continuous auction and the dealer market. Results are shown in Table 2. With one exception (DA3) the deviations between the midquotes and the asset value are larger than those (shown in Table 1) between average prices and value. This suggests that the spread is not set symmetrical around the fundamental value and that transactions are more likely to occur on the side that is closer to the asset value. This is an important observation because it implies that the effective bid–ask spread will overstate the transaction costs actually incurred by the traders.

4.1.2. Tests of efficiency

The analysis so far has been largely descriptive. Formal tests can be performed by comparing the pricing errors on a period-by-period basis. For this purpose the data from the six series were pooled, resulting in 72 observations for each of the three trading mechanisms and each of the three price series.[12] We first used an ANOVA to test whether there are significant differences in the mean MRE measures. The MRE measure was used because it expresses the pricing error as a percentage of the asset value and is therefore best suited for the pooled data set. The null hypothesis of equal means is rejected at the 5% level for the opening and the closing prices (p-values < 0.001 and 0.027, respectively) and at the 10% level for the average prices (p-value 0.070).

Next, a pairwise comparison of the efficiency measures was conducted. Both a t-test for paired samples and a non-parametric Wilcoxon matched-pairs signed-rank test were used. Since the results of the latter are similar, only the results of the t-test will be discussed. Panel A of Table 3 shows the results.

[12] It is common practice to analyze experimental data on a period-by-period basis. This assumes that the different trading periods of an experiment are independent observations. Since this is not necessarily true, the test statistics should be interpreted with care (see Friedman, 1993a, p. 428, footnote 17). In the light of this problem we also performed non-parametric tests where we treated each experiment (rather than each period) as one observation. This leaves us with six observations for each trading mechanism. A Kruskal-Wallis ANOVA rejects the null hypothesis of equal pricing errors for the opening prices, but not for the average and closing prices (p-values are 0.13 in both cases). In pairwise comparisons we found, using Wilcoxon tests, that the pricing errors calculated from opening prices are higher in the dealer market than in both the continuous auction and the call market. No significant difference was found between the two auction mechanisms. Pairwise comparisons for average and closing prices are not feasible because the Kruskal-Wallis ANOVA did not reject the null of equal medians. Finally, comparing the pricing errors in opening and closing prices within trading institutions, we found that the quality of prices improves only in the dealer market. These test results are consistent with those reported in Table 3 and are clearly based on independent observations.

Table 2
Informational efficiency measured on the basis of midquotes

Table 2 contains the MRE measures of price quality calculated from the midquotes in the continuous auction and dealer markets. The calculation is based on all midquotes in effect immediately prior to a transaction.

Series	Continuous auction (%)	Dealer market (%)
1	4.13	3.56
2	3.23	3.92
3	1.51	2.51
4	4.27	2.97
5	2.80	3.15
6	3.21	4.70
Average	3.19	3.47

There is no significant difference between the precision of the prices in the call and continuous auction markets.[13] The opening and the closing prices in the dealer markets are significantly less precise than those in the call markets and the continuous auction markets. If, as was discussed above, average prices are considered, the dealer market turns out to be significantly more efficient than the call market whereas the difference to the continuous auction is insignificant.

We also tested whether the quality of prices improves in the course of the trading period. Both a t-test and a Wilcoxon matched-pairs signed-rank test were used to compare the precision of opening and closing prices within trading mechanisms. The results of the t-test are presented in Panel B of Table 3. Only the dealer market shows a significant reduction of the pricing errors in the course of the trading periods. In the call market and the continuous auction, on the other hand, the precision of opening and closing prices is not significantly different from each other. This result can be interpreted in either of two ways. First, it is consistent with efficient price discovery very early in the trading period. It is, however, also consistent with persistent mispricing.

There is no clear answer to the question of which interpretation is correct. The answer depends on what one considers to be "efficient price discovery". In an attempt to discriminate between the two explanations we formed two groups of observations for each experiment. The first contains the six trading periods with

[13] The fact that the difference between the opening pricing errors in the call and continuous auction is insignificant is surprising. Given the temporal consolidation of the order flow in the call auction, one should expect significantly lower pricing errors in the call markets. It has, however, been documented that efficient price discovery in continuous markets may occur in the absence of trading (e.g. Cao et al., 2000; Greene and Watts, 1996).

Table 3
Informational efficiency: test results

Table 3 contains the results of t-tests for paired samples. The measure for the pricing error used for the comparisons is the MRE measure. The entry in each cell is the difference between the MRE of the row and the column price series where each price series consists of 72 observations (12 periods *6 experiments). A negative entry means that the price series defined in the row has a smaller pricing error than the one defined in the column. The differences are measured as a percentage of the asset value. t-values are given in parentheses. An * indicates significance at the 5%-level (2-tailed test).

Panel A: Comparison of the opening, average, and closing prices across trading mechanisms

		Continuous auction			Dealer market		
		Open	Av.	Close	Open	Av.	Close
Call market	open	− 0.330 (− 0.77)			− 3.219* (− 5.54)		
Call market	av.		0.482 (1.38)			0.727* (2.22)	
Call market	close			0.179 (0.44)			− 0.962* (− 2.13)
Continuous auction	open				− 2.889* (− 4.81)		
Continuous auction	av.					0.245 (0.91)	
Continuous auction	close						− 1.141 (− 2.22)

Panel B: Comparison of the opening and closing prices within the trading mechanisms

	Call market	Continuous auction	Dealer market
Open–Close	− 0.152 (− 0.78)	0.3565 (0.90)	2.105* (3.65)

the smallest pricing errors at the opening, the second contains the remaining six observations with the largest pricing errors at the opening. We then looked at the difference between the opening and closing pricing errors separately for both groups. We found that, when the pricing error at the opening was large, there was clear evidence of improvement for the continuous auction and the dealer market. The improvement was somewhat less pronounced for the call market; here, efficiency improved in four cases but worsened in two cases.

In the light of these findings, we interpret our results as follows. Price discovery is not perfectly accurate. Rather, there appears to be a certain level of mispricing which is not further reduced when trading continues. In the call market – where the pricing error at the opening is smallest - this level is often reached in the first transaction. There is thus not much room for improvement. In the continuous auction, further improvement is achieved when the pricing error at the opening is large. On the other hand, when pricing errors at the opening are small, no further improvement is achieved.

4.1.3. The adjustment of prices to new information

An important issue relating to informational efficiency is the question of how prices adjust to new information. Particularly, the issue of over- or underreaction of prices to new information has received much attention in the recent literature (e.g. Barberis et al., 1998; Daniel et al., 1998; Odean, 1998). Experimental research is especially suited to address this issue because the flow of information is controlled by the experimenter. It is therefore possible to directly analyze the pattern of successive price changes and to relate them to value changes. It can thus be tested whether there is evidence of over- or underreaction to new information. In the present case it is of particular interest whether any such pattern is caused by the trading mechanism. To analyze this we regressed the price changes on the value changes and the previous pricing error:

$$(p_t - p_{t-1}) = \alpha + \beta(v_t - v_{t-1}) + \gamma(p_{t-1} - v_{t-1}) + \varepsilon_t. \tag{5}$$

To assure the comparability of the experiments only opening and closing prices were used; other transactions in the continuous auction and dealer markets are thus excluded from the analysis. The regression was estimated for each experiment separately.[14]

In an efficient market, β should not be different from one. A value below one is evidence of an underadjustment to new information whereas a β larger than one can be interpreted as evidence of overreaction. The value of γ is expected to be negative because past pricing errors should be corrected. The regression results are presented in Table 4. The adjusted R^2 shows that the independent variables explain a large part of the variation in price changes. The call market produces the best fit (average R^2 0.87), followed by the continuous auction (0.82) and the dealer market (0.75). This comes as no surprise because prices (and, therefore,

[14] We also estimated the six equations for each of the trading mechanisms jointly and tested the null hypothesis of equal slope coefficients across equations. This hypothesis was rejected for the call market and the dealer market. We therefore estimated the equations for each experiment separately rather than pooling the data for each trading mechanism.

Note also that our specification assumes that the correction of the lagged pricing error as measured by the coefficient γ is the same at the opening (where the lagged pricing error is taken from the previous trading period) and at the close (where the lagged pricing error is from the opening transaction of the same period). We therefore also estimated the model allowing for different coefficients but could not reject the null hypothesis of equal coefficients.

Table 4
Determinants of price changes: regression results

Table 4 presents results of the regression $(p_t - p_{t-1}) = \alpha + \beta(v_t - v_{t-1}) + \gamma(p_{t-1} - v_{t-1}) + \varepsilon_t$. An * indicates that the coefficient is significantly different from zero at the 5% level. A shaded cell in the β-column indicates that the coefficient is significantly different from one at the 5% level. CM is the call market, DA the continuous (double) auction and MM the dealer (market maker) market.

Experiment	α	β	γ	R^2
CM 1	2.273	0.714*	− 0.389*	0.759
CM 2	− 2.808	0.843*	− 0.815*	0.846
CM 3	1.568	0.573*	− 0.363*	0.917
CM 4	− 4.065	0.686*	− 0.184	0.817
CM 5	4.590*	0.848*	− 0.438*	0.950
CM 6	− 0.456	0.761*	− 0.350*	0.955
Average	0.184	0.738	− 0.423	0.874
DA 1	− 8.398	0.996*	− 0.696*	0.840
DA 2	− 3.788	0.944*	− 0.863*	0.920
DA 3	− 7.232*	0.872*	− 0.751*	0.849
DA 4	− 10.122	0.805*	− 0.789*	0.646
DA 5	− 0.958	1.029*	− 0.773*	0.785
DA 6	− 2.396	0.947*	− 0.818*	0.876
Average	− 5.482	0.932	− 0.782	0.819
MM 1	− 8.002	1.272*	− 0.855*	0.643
MM 2	3.517	1.108*	− 0.639*	0.814
MM 3	− 2.591	1.143*	− 1.150*	0.798
MM 4	− 7.571*	0.551*	− 0.813*	0.763
MM 5	− 9.065	0.696*	− 1.012*	0.700
MM 6	6.471	0.457*	− 1.455*	0.803
Average	− 2.874	0.871	− 0.987	0.754

price changes) are affected by the spread in the continuous markets, and spreads are larger in the dealer market.

The most striking results are those for the call market. In all six series this trading mechanism exhibits significant underadjustment to new information. This result is in accordance with the model of Kyle (1985) which predicts an estimated coefficient of 0.5 for the single auction equilibrium.[15] It is also

[15] Underreaction may also be caused by behavioral biases such as overconfidence (see, e.g., Odean, 1998). In this case one should expect overconfident traders with inaccurate signals to incur losses. As will be shown later, however, we do not find a relation between signal accuracy and trading profits in the call markets. In the present context we therefore favor a microstructure explanation of the underreaction.

consistent with the lagged adjustment of call market prices to market information documented by Amihud et al. (1997). The absolute values of γ are smallest for the call market. This indicates that pricing errors are not corrected immediately but show a tendency to persist. Results for the two continuous trading mechanisms give more support to the hypothesis of efficient information processing. The continuous auction appears to incorporate new information very precisely. β estimates for the dealer markets show more variability (and have higher standard errors) although the hypothesis of a β equal to one can only be rejected in two cases.

The regression results show that the call market, although depicting pricing errors that are not larger than those of the other trading institutions, leads to inefficiencies because of the underreaction of prices to new information.[16]

4.2. Liquidity

A market is considered to be liquid when assets can be bought or sold rapidly and at a price close to the equilibrium value. This definition incorporates two dimensions, time and cost. The trading mechanisms differ with respect to the time dimension because in the call market access to the market is limited to discrete points in time. This makes it difficult to incorporate both the time and the cost dimension of liquidity into a single numeric measure. The following analysis is, therefore, restricted to the cost dimension.

The empirical evidence from field markets suggests that execution costs are higher in the dealer market than in the continuous auction market. The temporal consolidation of the order flow in the call market may lead to a reduction of the execution costs below the level of the continuous auction. Several liquidity measures have been proposed in the literature. In the sequel the trading volume, the bid–ask spread and Roll's (1984) serial covariance estimator will be analyzed.

4.2.1. Trading volume

Trading volume is often considered to be a good proxy for liquidity. This can be justified with the high correlation usually observed between the trading volume for different stocks and their bid–ask spreads or other measures of liquidity. Therefore, Table 5 reports the trading volume in the 18 experimental markets. It is apparent from the figures that trading volume is lowest in the call

[16] We re-estimated the model including the lagged value change as an additional explanatory variable. The estimated coefficient was positive for all six call market experiments, two of the estimated coefficients were significant at the 5% level. For the continuous auction and dealer markets, we obtained coefficient estimates of varying signs. None of the coefficients was significantly different from zero. These additional results strengthen the conclusion of underreaction in the call market.

Table 5
Trading volume and net change in asset holdings

Table 5 shows the gross trading volume and the net change in asset holdings, both averaged over the periods 4–15. In the call market volume is the total volume from both calls within the period and the net change in asset holdings is calculated from the portfolios after the second call.

Series	(Gross) trading volume			Net change in asset holdings		
	Call market	Continuous auction	Dealer market	Call market	Continuous auction	Dealer market
1	913.3	1339.7	1135.9	595.7	591.1	685.8
2	587.3	637.5	1365.8	511.0	432.3	612.3
3	884	833.9	1835.2	555.0	440.0	747.4
4	304.3	1142.3	2211.3	174.3	626.4	1202.8
5	693.4	1215.8	1490.7	595.3	626.4	809.6
6	613.3	909.2	1367.9	483.9	444.8	433.2
Average	665.9	1013.1	1567.8	485.9	526.8	748.5

market and highest in the dealer market. A t-test and a non-parametric Mann-Whitney u-test (results are not shown), both for the pooled sample, reveal that these differences are highly significant.

This result is, however, not surprising. The call market with two calls per period offers very limited possibilities to buy and sell assets within the same period. On the other hand, volume in the dealer market is likely to be systematically overstated due to the double-counting of transactions. A way to correct both effects is to replace the (gross) trading volume with the net change in asset holdings also shown in Table 5. This measure is calculated by summing the absolute deviation between end-of-period and initial asset holdings over the participants and then dividing by 2. The net change in asset holdings does not contain that part of the trading volume that is caused by buying and reselling shares (or vice versa) within a single period. As a consequence, double-counting in the dealer market is also eliminated.

The results show that there is no systematic difference between the call market and the continuous auction market whereas the net change in asset holdings tends to be higher in the dealer market. This does, however, not in itself constitute an advantage because higher trading volume is not necessarily associated with lower execution costs.

4.2.2. Bid–ask spread

The most widespread measure of execution costs is the bid–ask spread. Measurement of the spread is straightforward in the continuous auction and the dealer market. We calculate the effective spread which measures the difference between the transaction price and the midquote. The latter is usually considered to be the best estimate of the true asset value available to investors without private information.[17]

No explicit spread exists in the call market. Given the availability of order flow data, it is, however, nevertheless possible to construct a spread measure for the call market. The cost of transacting can be assessed after each transaction by considering the sell order with the lowest ask price and the buy order with the highest bid price that were not executed. The bid–ask spread in the call market can be defined as the difference between these two price limits (Mendelson, 1982; Friedman, 1993a).[18] The rationale for this measure is the fact that an additional order would have been filled at a price no worse than the prices constituting the spread. Therefore, the spread as defined above can be interpreted as a measure of the price impact of an additional order.

Average bid–ask spreads for the 18 experiments, expressed as a percentage of the asset value, are shown in Table 6.[19] The spreads are lowest in the call market and highest in the dealer market.[20] Both a one-way ANOVA and a non-parametric Kruskal-Wallis one-way ANOVA (results are not shown) reveal that the differences in the means are highly significant. This conclusion is reinforced when the data is disaggregated and each period is taken as one observation.

Consistent with the finding that spreads in the dealer market are high, the market maker position is profitable. On average, profits of subjects acting as market makers are 5–10% higher than average profits. Since the market maker positions were assigned randomly, different trading skills can not explain this result.

---

[17] Transactions at prices inside the quotes were not possible. Therefore, the effective spread is identical to the quoted spread immediately prior to the transaction (The *current spread* in the terminology of Neal, 1992). In spite of this, the average quoted spread in the experimental markets is higher than the average effective spread. The reason is that transactions tend to occur when the spread is low. This phenomenon is also observed in real world stock markets.

[18] Note that in the case of rationing one of the prices used to calculate the spread is the market price because in this case not all orders that were executable at the market clearing price could actually be executed.

[19] The result for the dealer market in session 6 must be regarded as an outlier. Here, two subjects consistently managed to buy high and sell low. Market makers thus had no incentive to lower their spreads.

[20] The effective spread measures execution costs relative to the quote midpoint. As outlined earlier, this may overstate the true cost if the spread is not set symmetrically around the fundamental value and transactions occur with greater probability on the side of the market where the spread is closer to the fundamental value. To correct for this potential bias we also estimated the execution costs relative to the fundamental value. Although the differences between the trading institutions are smaller, the qualitative results remain unchanged.

Table 6
Market liquidity: bid–ask spreads

Table 6 shows the average bid–ask spreads in the 18 experiments. For the continuous auction and dealer markets average effective spreads are shown. The spread in the call market is the difference between the limit prices of the lowest ask and the highest bid that were not executed. The spread for each period was calculated as a percentage of the asset value; then, the average spread for each experiment was calculated.

Series	Call market	Continuous auction	Dealer market
1	1.728	4.838	11.597
2	2.482	4.069	11.117
3	1.187	3.225	8.053
4	1.855	5.397	9.437
5	1.799	4.093	9.395
6	1.506	6.481	22.286
Average	1.760	4.684	11.981

4.2.3. *The serial covariance estimator*

The serial covariance estimator developed by Roll (1984) allows to calculate an estimate of the spread when quote data is not available. It has also been used in order to assess the transaction costs in call markets (e.g. Haller and Stoll, 1989; Pagano and Röell, 1990; Stoll and Whaley, 1990). The serial covariance estimator is defined as

$$s_r = 2\sqrt{-\,Cov(\Delta p_t;\Delta p_{t-1})}, \tag{6}$$

where $Cov(\Delta p_t;\Delta p_{t-1})$ is the serial covariance of successive price changes.

As quote data for the experimental markets is available, there is little need to calculate the Roll measure in order to estimate the transaction costs in these markets. We can, however, use the experimental data to assess the accuracy of the serial covariance estimator. This is important because the argument can be made that the serial covariance estimator is not an appropriate measure of the transaction cost in a call market. It will tend to overstate the transaction costs.

To show this, it is assumed that the price process is described by $p_t = v_t + \varepsilon_t$ where v_t is the value of the asset at time t and ε_t is an i.i.d. random disturbance with $E(\varepsilon_t) = 0$ and $Var(\varepsilon_t) = \sigma_\varepsilon^2$. The Roll measure s_r can then be shown to be $s_r = 2\sigma_\varepsilon$ irrespective of the trading mechanism. This is equivalent to the usual formulation since there $\varepsilon_t = s/2$ and $\varepsilon_t = -s/2$, both with probability 0.5, and therefore $\sigma_\varepsilon = s/2$ where s denotes the (explicit) spread. Thus the Roll measure can, in principle, be applied to call markets. However, the interpretation of the resulting estimate is different. In both a continuous auction and a dealer market, $\varepsilon_t = s/2$ for a market participant who buys assets using a market order and

Table 7
Market liquidity: Roll measure versus effective spreads

Table 7 shows the average effective bid–ask spread (in absolute rather than in percentage terms) and the Roll (1984) serial covariance estimator $s_r = 2\sqrt{-Cov(\Delta p_t, \Delta p_{t-1})}$. For the continuous trading mechanisms it has been computed from only opening and closing prices (upper entry) and for the time series of all transaction prices (lower entry). The spread has to be compared with the Roll measure calculated from all prices. n/a indicates that the Roll measure could not be computed because the serial covariance of successive price changes was positive.

Series	Call market		Continuous auction		Dealer market	
	Roll measure	Average spread	Roll measure	Average spread	Roll measure	Average spread
1	22.236	8.292	31.880		67.068	
			21.140	21.672	29.722	53.180
2	35.888	11.75	27.154		51.590	
			20.567	18.994	31.776	52.415
3	n/a	5.125	27.171		72.198	
			11.163	14.180	22.544	35.057
4	11.493	8.375	47.047		10.192	
			16.693	24.234	20.288	42.965
5	8.641	7.208	35.863		32.026	
			12.782	15.964	18.708	36.977
6	17.246	7.708	n/a		94.112	
			29.708	32.426	85.187	108.670

$\varepsilon_t = -s/2$ for a sale. Therefore, the spread estimate can readily be interpreted as a measure of execution costs. In the call market, however, the realization of ε does not depend on whether an individual market participant buys or sells but depends on aggregate demand and supply conditions. Although a single order may have a price impact (as is evidenced by the spread measure discussed above), ε is likely to largely overstate this impact and can thus not be interpreted as a measure of execution cost.

Table 7 contains the results for the Roll measure and compares them with the bid–ask spread. The spread is calculated as defined earlier, but here the absolute level of the spread rather than a percentage is given. The Roll measure has been calculated twice for the continuous trading mechanisms; once on the basis of opening and closing prices (upper entry in the cells) and once on the basis of all transaction prices (lower entry). The latter leads to lower spread estimates in most cases. The main reason is that it sometimes occurs that traders or market makers enter quotes for large quantities that are quickly exhausted by other traders. This leads to a series of transactions at the same price which reduces the negative serial correlation in the return series.

The serial covariance estimator is lower for the continuous auction market than for the dealer market. The values for the call market tend to be even lower. Comparing the bid–ask spread with the Roll measure yields interesting results. For the call market the Roll measure is higher than the spread. This confirms the above analysis which predicts that the Roll measure overstates the execution costs in the call market.

In the continuous auction the Roll measure provides a precise estimate of the spread. Its average value is 18.68 as compared to an average spread of 21.24. The correlation between the two measures is 0.895. In the dealer market, on the other hand, the average spread is much higher than the Roll measure (average over the six experiments 54.88 as compared to 34.70). The correlation is, however, high as evidenced by a rank correlation of 0.77.[21] The large discrepancy between the explicit spread and the serial covariance estimator corroborates the conjecture that the effective spread may be a biased measure of the realized execution cost. Despite the numerical differences, however, all measures clearly indicate that execution costs are highest in the dealer market. This result is consistent with empirical evidence obtained from field data.

4.3. Signal precision, portfolio structure and trading profits

During the experiments subjects receive information of different quality. Because subjects do not know the precision of their signals at the beginning of a period they are unable to adjust their behavior accordingly. It can therefore be expected that the quality of the information influences their trading decisions and the total wealth at the end of the period.[22]

If subjects act on their information and can, at least at the beginning of the period, not judge the quality of their signals, then one should expect to find that subjects observing higher signals will, on average, buy assets while subjects with low signals will sell assets. We computed the correlation and the rank correlation between the signal quality and the number of assets in the end-of-period portfolio for each period of each experiment. The signal quality is defined as the signed difference between signal and asset value. The results are presented in Table 8. As has been expected, there is a clear relationship between the two

[21] The rank correlation coefficient is used because the correlation is heavily influenced by the extreme results from session 6.

[22] In private value experiments as defined in Section 2 (e.g. Plott and Sunder, 1982) different traders value the asset differently. The traders with the highest valuation should eventually hold all assets. By comparing the actual allocation with this rational expectations equilibrium allocation, the allocational efficiency of the markets can be analyzed. This is not possible in a common value design because the equilibrium allocation is not known.

Table 8
Information and asset allocation

Table 8 shows the correlation and rank correlation between the signed deviation between private signal and value on the one hand and the number of assets in the end-of-period portfolio on the other hand. The coefficients have been computed separately for each period and each experiment. This results in a total number of 72 coefficients for each market mechanism. CM is the call market, DA the continuous (double) auction and MM the dealer (market maker) market.

	CM	DA	MM
Average correlation coefficient	0.437	0.423	0.491
Number of positive coefficients	69	62	68
Significant (10%, one-tailed test)	48	47	56
Average rank correlation coefficient	0.398	0.414	0.520
Number of positive coefficients	67	61	68
Significant (10%, one-tailed test)	42	45	55

variables, indicating that subjects with higher [lower] signals buy [sell] assets on average. This relationship seems to be somewhat more pronounced in the dealer market.

The results presented in Table 8 document a relation between the *signed* signal quality and the *signed* asset holdings. We further expect a relation between the *absolute* signal quality and the trading volume. Traders with more extreme signals observe a higher discrepancy between their private signal and the market price. They are thus more likely to trade than traders with more precise signals. This is especially true for the call market in which the first transaction occurs before any information besides the private signals becomes available. We regressed the net trading volume for each subject and each period on the absolute signal quality. The results, shown in Panel A of Table 9, confirm the expectation of a positive relation between absolute signal precision and net trading volume. The coefficient for the call market is larger than the coefficients for the continuous auction and the dealer market. The differences between the coefficients are, however, not significant.

The relation between signal precision and trading volume should also hold on the aggregate level. In periods with more dispersed signals, differences in opinion between traders are more pronounced and trading volume is likely to be higher. We used the coefficient of variation as our measure of the dispersion of the signals and relate this measure to the net trading volume. The coefficients of correlation between the two variables are presented in Panel B of Table 9. The relation between signal dispersion and net trading volume is strongest in the call markets. The average correlation is 0.24 and all six individual coefficients are

Table 9
Information quality and trading volume

Panel A: Results of a regression of the net change in asset holdings of trader i in period t (defined as the absolute difference between the initial endowment and the number of assets at the end of the period) on the quality of the signal the trader received in that period (defined as the percentage deviation between the signal and the fundamental value, $|s_{i,t} - v_t|/v_t$).

	Intercept (p-value)	Slope (p-value)
Call market	50.69	619.91
	(0.00)	(0.00)
Continuous auction	62.92	509.49
	(0.00)	(0.00)
Dealer market	95.88	591.04
	(0.00)	(0.00)

Panel B: Correlation between the signal dispersion (measured by the coefficient of variation, the ratio of the standard deviation of the signals of the 12 traders in a period to the mean of the signals) and the net change in asset holdings (as defined earler).

	Call market	Continuous auction	Dealer market
Average	0.243	0.136	0.158
Number of positive coefficents	6	4	4

positive. The average coefficients for the continuous auction and the dealer markets are 0.14 and 0.16, respectively, and four of the six coefficients are positive.

Finally, it is interesting to analyze whether subjects with less precise information have lower end-of-period wealth. Due to its stochastic character the quality of the information a subject receives changes from period to period. Averaged over the periods, the signal quality is therefore approximately equal. Tests must thus be performed on a disaggregated basis. The following procedure was chosen. For each period subjects were divided into two groups according to the quality of their signals. The absolute deviation between the signal and the asset value was used as measure of information quality. Then the average end-of-period wealth of both groups was calculated. This procedure was repeated for each period. Due to the random character of the signal quality the composition of the groups changes from period to period. Finally, the results were aggregated over the 12 periods of the experiment. Panel A of Table 10 shows the difference between the average wealth of the groups with high and low signal precision. For the dealer market the calculation has been repeated after exclusion of the market makers. These results are also shown in the table.

Table 10
Information quality and end-of-period wealth

Panel A: Differences in the average end-of-period wealth between the groups of subjects with high and low signal precision. In each period subjects were grouped by the absolute deviation between their private signal and the asset value. Then the average end-of-period wealth was calculated and averaged over the 12 periods. The composition of the groups changes from period to period because the signal quality is a random variable. Results in parentheses in the last column are calculated after elimination of the market makers from the sample.

Average advantage due to high information quality	CM	DA	MM
1	− 225.5	1183.0	1385.9 (618.9)
2	493.0	− 776.6	2254.8 (1133.5)
3	597.1	− 1022.6	2524.6 (432.9)
4	164.6	− 907.1	557.3 (− 522.5)
5	− 108.3	386.0	888.2 (819.5)
6	− 830.1	1836.9	− 241.9 (−1154.0)

Panel B: Results of a regression of the normalized profit of trader i in period t (defined as the relation between the end-of-period wealth of that particular trader and the average end-of-period wealth of all traders in that period) on the quality of the signal the trader received in that period (defined as the percentage deviation between the signal and the fundamental value, $|s_{i,t} - v_t|/v_t$).

	α (p-value)	β (p-value)
Call market	1.000 (0.00)	− 0.0097 (0.65)
Continuous auction	0.997 (0.00)	0.0591 (0.26)
Dealer market	1.014 (0.00)	− 0.2849 (0.02)

In both the call market and the continuous auction there is apparently no relationship between information quality and end-of-period wealth. For the dealer market there is weak evidence of a relationship. These findings are confirmed by a regression analysis for the pooled data set. For this analysis, the end-of-period wealth $w_{i,t}$ of each subject was expressed as a percentage of the average wealth of all subjects at the end of the period. Then the data was pooled over all six experiments conducted with each trading mechanism. Finally, the normalized wealth variable was regressed on the percentage deviation between the private signals and the asset value:

$$w_{i,t} = \alpha + \beta \frac{|s_{i,t} - v_t|}{v_t} + \varepsilon_{i,t} \tag{7}$$

The results, shown in Panel B of Table 10, confirm the earlier finding that a relation between signal quality and end-of-period wealth exists in the dealer market only.

It can thus be concluded that traders with imprecise signals do, on average, not loose in the call and continuous auction markets. The pattern of results depicted in Table 10 is consistent with efficient price discovery in the experimental auction markets. Subjects with high [low] signals buy [sell] assets on average but they transact at fair prices and therefore do not loose.

The results for the dealer markets may be explained by the fact that traders in these markets do not have the choice to supply liquidity. Traders with imprecise signals trade more (as documented in Table 9) and thus incur higher execution costs. This, in turn, reduces their trading profits.

5. Summary

This paper reports the results of a series of market experiments that were conducted in order to compare the call market, the continuous auction market and the dealer market. The design incorporates asymmetric information but guarantees that the ex-ante quality of the private signals of all traders is identical. Therefore, the aggregation of diverse information can be analyzed in the absence of insider trading.

It is found that opening prices in the call market are closer to the true value of the asset than opening prices in the continuous auction and the dealer markets, although the difference to the continuous auction is not significant. The practice of many exchanges to start trading with a call auction can thus be justified. Single transaction prices in the call and continuous auction markets are generally much more efficient than prices in the dealer market. The latter is, however, very efficient when average prices are analyzed. Averaging the prices of a trading period largely eliminates the bid–ask spread. The conclusion, therefore, is that prices in a dealer market convey information of high quality, but at the expense of high transaction costs.

The call market, although depicting small pricing errors, shows a systematic tendency towards underreaction to new information. This result is consistent with the predictions of theoretical models such as Kyle (1985) and the empirical results reported in Amihud et al. (1997).

An analysis of market liquidity leads to the conclusion that execution costs are lowest in the call market and highest in the dealer market. This result does not depend on how execution costs are measured. It is consistent with empirical evidence obtained from field data. The trading volume turns out to be a poor proxy for the execution costs because higher trading volume is associated with higher rather than lower spreads. The experimental results confirm the prediction that Roll's (1984) serial covariance measure overstates the execution costs in call markets.

The accuracy of the signals traders receive influences the composition of individual portfolios and the trading volume. A relation between signal quality and profits is only found in the dealer market. This result is consistent with efficient price discovery in the call and continuous auction markets.

There are several promising ways to extend the research presented in this paper. First, similar comparisons of trading mechanisms can be performed using a design that incorporates insiders. Second, the design of the trading mechanisms can be varied. Possible variations are the introduction of indicative prices or an open order book in the call market or of varying numbers of market makers in the dealer market.

References

Admati, A., Pfleiderer, P., 1988. A theory of intraday patterns. Review of Financial Studies 1, 3–40.

Amihud, Y., Mendelson, H., Lauterbach, B., 1997. Market microstructure and securities values: evidence from the tel aviv stock exchange. Journal of Financial Economics 45, 365–390.

Barberis, N., Shleifer, A., Vishny, R., 1998. A model of investor sentiment. Journal of Financial Economics 49, 307–343.

Bessembinder, H., Kauffman, H., 1996. A comparison of quotations, trading costs, and volatility for NYSE and NASDAQ-listed stocks. NYSE working paper 96–03, May.

Bloomfield, R., 1996. Quotes, prices and estimates in a laboratory market. Journal of Finance 51, 1791–1808.

Bloomfield, R., O'Hara, M., 1999. Market transparency: who wins and who loses? Review of Financial Studies 12, 5–35.

Bloomfield, R., O'Hara, M., 2000. Can transparent markets survive? Journal of Financial Economics 55, 425–459.

Cao, C., Ghysels, E., Hatheway, F., 2000. Price discovery without trading: evidence from the NASDAQ pre-opening. Journal of Finance 55, 1339–1365.

Christie, W., Huang, R., 1994. Market structure and liquidity: a transactions data study of exchange listings. Journal of Financial Intermediation 3, 300–326.

Daniel, K., Hirshleifer, D., Subrahmanyam, A., 1998. Investor psychology and security market under- and overreactions. Journal of Finance 53, 1839–1885.

de Jong, F., Nijman, Th., Röell, A., 1995. A comparison of the cost of trading French shares on the Paris Bourse and on SEAQ International. European Economic Review 39, 1277–1301.

Diamond, D., Verrecchia, R., 1981. Information aggregation in a noisy rational expectations economy. Journal of Financial Economics 9, 221–235.

Easley, D., Ledyard, J., 1993. Theories of price formation and exchange in double oral auctions. In: Friedman, D., Rust, J. (Eds.), The Double Auction Market. Addison-Wesley, Reading, MA, pp. 63–97.

Flood, M., Huisman, R., Koedijk, K., Mahieu, R., 1999. Quote disclosure and price discovery in multiple-dealer financial markets. Review of Financial Studies 12, 37–59.

Flood, M., Huisman, R., Koedijk, K., Mahieu, R., Röell, A., 1997. Post-trade transparency in multiple dealer financial markets. Working paper, Concordia University, Limburg Institute of Financial Economics, Erasmus University Rotterdam and Université Libre de Bruxelles, March.
Flood, M., Huisman, R., Koedijk, K., Dijik van, M., van Leeuwen, I., 1998. The more you see, the less you get: price-competing insiders under different trading mechanisms. Working paper, University of North Carolina and Limburg Institute of Financial Economics, May.
Friedman, D., 1984. On the efficiency of experimental double auction markets. American Economic Review 74, 60–72.
Friedman, D., 1991. A simple testable model of double auction markets. Journal of Economic Behavior and Organization 15, 47–70.
Friedman, D., 1993a. How trading institutions affect financial market performance: some laboratory evidence. Economic Inquiry 21, 410–435.
Friedman, D., 1993b. Privileged traders and asset market efficiency: a laboratory study. Journal of Financial and Quantitative Analysis 28, 515–534.
Glosten, L., 1994. Is the electronic open limit order book inevitable? Journal of Finance 49, 1127–1161.
Glosten, L., Milgrom, P., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. Journal of Financial Economics 14, 71–100.
Greene, J., Watts, S., 1996. Price discovery on the NYSE and the NASDAQ: the case of overnight and daytime news releases. Financial Management 25, 19–42.
Haller, A., Stoll, H., 1989. Market structure and transaction costs: implied spreads in the german stock market. Journal of Banking and Finance 13, 697–708.
Hellwig, M., 1980. On the aggregation of information in competitive markets. Journal of Economic Theory 22, 477–498.
Ho, Th., Schwartz, R., Whitcomb, D., 1985. The trading decision and market clearing under transaction price uncertainty. Journal of Finance 40, 21–42.
Huang, R., Stoll, H., 1996. Dealer versus auction markets: a paired comparison of execution costs on NASDAQ and the NYSE. Journal of Financial Economics 41, 313–357.
Keim, D., Madhavan, A., 1996. Transaction costs and investment style: an inter-exchange analysis of institutional equity trades. Working paper, August.
Krahnen, J., Weber, M., 1999. Does information aggregation depend on market structure? Market makers vs. double auction. Zeitschrift für Wirtschafts- und Sozialwissenschaften 119, 1–22.
Krahnen, J., Weber, M., 2000. Marketmaking in the laboratory: does competition matter? Working paper, April.
Kyle, A., 1985. Continuous auctions and insider trading. Econometrica 53, 1315–1335.
Kyle, A., 1989. Informed speculation with imperfect competition. Review of Economic Studies 56, 317–356.
Lamoureux, Ch., Schnitzlein, Ch., 1996. When it's not the only game in town (the effect of bilateral search on the quality of a dealer market). Journal of Finance 52, 683–712.
Lee, Ch., 1993. Market integration and price execution for NYSE-listed securities. Journal of Finance 48, 1009–1038.
Madhavan, A., 1992. Trading mechanisms in securities markets. Journal of Finance 47, 607–641.
Mendelson, H., 1982. Market behavior in a clearinghouse. Econometrica 50, 1505–1524.
Mendelson, H., 1987. Consolidation, fragmentation, and market performance. Journal of Financial and Quantitative Analysis 22, 189–207.
Neal, R., 1992. A comparison of transaction costs between competitive market maker and specialist market structures. Journal of Business 65, 317–334.
Odean, T., 1998. Volume, volatility, price and profit when all traders are above average. Journal of Finance 53, 1887–1934.
O'Hara, M., 1995. Market Microstructure Theory. Basil Blackwell, Oxford.

Pagano, M., Röell, A., 1990. Trading systems in European stock exchanges: current performance and policiy options. Economic Policy 10, 63–115.
Pagano, M., Röell, A, 1992. Auction and dealership markets? - what is the difference? European Economic Review 36, 613–623.
Pagano, M., Röell, A., 1996. Transparency and liquidity: a comparison of auction and dealer markets with informed trading. Journal of Finance 51, 579–611.
Plott, Ch., Sunder, S., 1982. Efficiency of experimental security markets with insider information: an application of rational expectations models. Journal of Political Economy 90, 663–698.
Roll, R., 1984. A simple implicit measure of the bid/ask spread in an efficient market. Journal of Finance 39, 1127–1139.
Rustichini, A., Satterthwaite, M., Williams, St., 1994. Convergence to efficiency in a simple market with incomplete information. Econometrica 62, 1041–1063.
Satterthwaite, M., Williams, St., 1993. The Bayesian theory of the k-double auction. In: Friedman, D., Rust, J. (Eds.), The Double Auction Market. Addison-Wesley, Reading, MA, pp. 99–123.
Schnitzlein, Ch., 1996. Call and continuous trading mechanisms under asymmetric information: an experimental investigation. Journal of Finance 51, 613–636.
Smith, V., Suchanek, G., Williams, A., 1988. Bubbles, crashes, and endogeneous expectations in experimental spot asset markets. Econometrica 56, 1119–1151.
Stoll, H., 1993. Equity trading costs in-the-large. Journal of Portfolio Management 19 (4), 41–50.
Stoll, H., Whaley, R., 1990. Stock market structure and volatility. Review of Financial Studies 3, 37–71.
Van Boening, M., Williams, A., LaMaster, S., 1993. Price bubbles and crashes in experimental call markets. Economics Letters 41, 179–185.
Wilson, R., 1987. On equilibria of bid–ask markets. In: Feiwel, G. (Ed.), Arrow and the Ascent of Economic Theory, Essays in Honour of Kenneth G. Arrow. Macmillan Press, Basingstoke.