

Quote Disclosure and Price Discovery in Multiple-Dealer Financial Markets

Mark D. Flood

Concordia University (Montreal) and UNC-Charlotte

Ronald Huisman

Erasmus University

Kees G. Koedijk

University of Maastricht and Erasmus University

Ronald J. Mahieu

Erasmus University

We examine the effects of price disclosure on market performance in a continuous experimental multiple-dealer market in which seven professional market makers trade a single security. The dealers trade with one another and with computerized informed and liquidity traders. Our key comparison is between fully public price queues (pretrade transparent market) and bilateral quoting (pretrade opaque). We find that opening spreads are wider and trading volume is lower in the opaque markets due to higher search costs there. More importantly, however, higher search costs also induce more aggressive pricing strategies, so that price discovery is much faster in the opaque markets.

This article addresses the institutional design of financial markets. It concentrates on the degree of public quote disclosure in the market and its effect on the price discovery process. We are most interested in transactable

The authors would like to thank Shawn Lamaster and the Economic Science Laboratory at the University of Arizona for providing the netbios communications libraries (and related advice) for the experimental software; Rob Scott for his work coding a preliminary version of the program; Otto Perdeck of the Center for Research in Experimental Economics and Political Decision-Making (CREED) at the University of Amsterdam for his professional efforts in coding the final version; Arthur Schramm and CREED for making their laboratory facilities available for the experiments; ABN Amro Bank, de Generale Bank, ING-Barings Bank, Optiver, Rabobank, and their respective securities dealers for their participation; and Auke Schilder and Ingrid van Sundert for their assistance during the experiments. The authors would also like to thank several anonymous referees, Robert Battalio, Robert Bloomfield, Peter Bossaerts, Michael Brennan, Piet Eichholtz, Lawrence Glosten (the editor), Linda Goldberg, Adam Land, Martin Lettau, Rich Lyons, Greg Lypny, Jim Moser, Maureen O'Hara, Carol Osler, Franz Palm, Ailsa Röell, Gideon Saar, Peter Schotman, Seymour Smidt, Casper de Vries, and workshop participants at the California Institute of Technology, Concordia University, Cornell University, EFA 1996 meetings, FMA 1996 meetings, Georgia State University, the Federal Reserve Banks of New York and San Francisco, Maastricht University, Tilburg University, and the University of California, Berkeley for their helpful comments on a previous version of this and a related article. All remaining errors pertain to the authors. This research was supported in part by a grant from the Social Sciences and Humanities Research Council of Canada, and by an FRDP grant from Concordia University. Address correspondence to Mark D. Flood, Department of Finance, Belk College of Business Administration, University of North Carolina at Charlotte, 9201 University City Blvd., Charlotte, NC 28223, or e-mail: mdfflood@email.uncc.edu.

The Review of Financial Studies Spring 1999 Vol. 12, No. 1, pp. 37–59
© 1999 The Review of Financial Studies 0893-9454/99/\$1.50

prices (i.e., live quotes) rather than indicative prices or posttrade transaction prices. Specifically we ask how the degree of public quote disclosure of transactable prices affects important market characteristics, such as transaction costs and the dynamic adjustment to strong-form efficient prices (price discovery).

Worldwide, many financial markets coexist while differing crucially in the way market participants receive price information. For example, electronic limit order books, such as those on the Paris Bourse, are very price transparent, since brokers can obtain information on quotes and order sizes from all limit orders on the book. On the other hand, the London SEAQ display (and, to a lesser extent, the NASDAQ) provide participants with quotes, but these are largely indicative, since most market makers will improve upon their quoted bids and asks in telephone negotiations, offering within-the-quotes prices for most transactions. Even close substitutes, such as foreign currency futures and forward contracts, trade in markets with very different degrees of quote disclosure. One explanation for these coexisting quote disclosure regimes in financial markets is clearly found in the type of securities being traded. For example, futures trading is typically centralized (while forward trading usually is not), at least in part because futures contracts are standardized. However, it is frequently the case that the same security trades in different microstructures, often simultaneously, and sometimes even in the same room. For example, the NYSE runs a call market at the morning opening and continuous specialist trading throughout the day, along with a separate upstairs market for block trades. The foreign exchange market operates simultaneously as a decentralized interbank direct market and a more centralized brokered market.

It is generally assumed that price-transparent microstructures better allow traders to extract information from outstanding quotes, leading (rather intuitively) to prices impounding a maximum of available information. For example, O'Hara (1995, p. 270) notes that, "the transparency of prices allows traders better ability to extract price information from the market price, a process that surely abets the goal of equilibrium price discovery."¹ From this point of view, it is not obvious why markets with different levels of pretrade transparency coexist (following Pagano and Röell (1996), we say a market is fully pretrade transparent if at all times participants can clearly see all outstanding quotes). In this article we measure the effects of pretrade transparency on the pricing and inventory strategies of market makers, in terms of liquidity (spreads and volume) and the price discovery process.

¹ There is a vast literature on market transparency, including empirical and experimental studies of price discovery. Surveys of the relevant microstructure literature are available in Admati (1991), Yu (1993), and O'Hara (1995, especially ch. 9). Stoll (1992) and Harris (1993) survey issues related to institutional design. Goodhart and O'Hara (1994) and Guillaume et al. (1994) cover recent empirical results. Davis and Holt (1993, especially chs. 3 and 5) and Duxbury (1995) survey the experimental literature.

We find a trade-off between liquidity and price efficiency that may partly explain the coexistence of microstructures with different levels of pretrade transparency.

Our study attempts to advance the literature in several dimensions. First, we provide a direct comparison of alternative microstructural arrangements. Such direct comparisons are relatively rare, as theoretical studies require restrictive assumptions about dealer behavior and information processing, while traditional empirical studies are subject to strong *ceteris paribus* caveats in such a comparison across regimes. The experimental methodology directly addresses both of these limitations: by using human subjects in a controlled environment, the need for both behavioral and *ceteris paribus* assumptions is sharply reduced. Second, we focus exclusively on multiple-dealer markets. Most microstructural analyses concentrate on lone dealers, in spite of the importance of multiple-dealer markets, such as the interbank foreign exchange and money markets or the over-the-counter stock, bond, and derivatives markets. We argue that interactions among multiple dealers can produce interesting phenomena that do not arise in single-dealer environments. Lastly, we focus on the interaction between inventory management and pricing in a dealer's strategy. Such a focus is greatly facilitated by the methodology, which produces a large amount of simultaneous data on prices, inventories, and dealer information sets.

Despite the numerous studies of posttrade transparency [defined as the amount of transaction information available to market makers, as per Pagano and Röell (1996)], the effects of (pretrade) quote disclosure on market performance have received much less attention.² Madhavan (1992) distinguishes between "order-driven" and "quote-driven" markets, comparing their relative price efficiency. In quote-driven markets, investors trade against outstanding prices, while in order-driven markets, participants must submit orders first, after which prices are determined. The quote-driven markets are necessarily more pretrade transparent than order-driven markets. Madhavan finds that a quote-driven market is more price efficient than an order-driven market. Bollerslev and Domowitz (1993) compare different "book lengths" for an electronic limit order book (patterned on the CME/Reuters Globex trading system) in a computer simulation. Their general conclusion is that increasing book length—the number of limit orders held on the book and broadcast to participants—improves price discovery (with the exception of trading in unit lots, for which a technical artifact of the order-posting algorithm slows price discovery).

The experimental approach to financial markets is a growing field. Examples are the studies of Schnitzlein (1996) and Lamoureux and Schnitzlein

² Recent examples of the former include Bloomfield and O'Hara (1999), Gemmill (1996), Pagano and Röell (1996), Röell (1996) and Flood et al. (1997).

(1997), who designed an auction market based on the theoretical framework of the Kyle (1985) model. In this article, on the other hand, the experimental environment can be seen as the continuous, multiple-dealer version of the Glosten and Milgrom (1985) model, which is a quote-driven market. This means that traders both set quotes and trade and are confronted with informed and liquidity motivated investors who view the quoted bid and ask and decide whether to trade one unit at a time or not. Our experimental subjects are not university students, but professional securities traders from five Dutch banks, a fact that we feel enhances the reliability of our results. The continuous-time environment of the experiment produces large amounts of data. Also, subjects must continuously allocate their scarce time among a number of different activities, such as modifying quotes, initiating trades, and analyzing their information sets. Such a setting, in which traders must make instantaneous decisions about where to focus their energies, is designed to mimic the rough-and-tumble of real-world securities markets. Furthermore, our subjects, realistically, not only post quotes but also initiate trades on their own account. As dealers are more numerous here (seven quoting dealers) than in most previous experimental studies, our market generates significant interdealer trading.

This experimental environment is closely related to Flood et al. (1997), and it is comparable with Bloomfield and O'Hara (1999). Flood et al. (1997) use almost the same setting as is used here, and concentrate on the effects of posttrade transparency on market performance. Bloomfield and O'Hara (1999) study the effects of both posttrade and pretrade transparency in a noncontinuous sequential trade setting with two market makers, and they find no significant effects of quote disclosure on price efficiency or spreads.

As discussed, we compare directly the effects of two market structures on the pricing and trading strategy of professional market makers, and therefore on the price efficiency of those markets. The markets differ only in the way quotes are disclosed. In the transparent market, all quotes are disclosed publicly and immediately, which means that all market makers have all outstanding quotes presented on their private trading screens. In the opaque market, no quotes are publicly disclosed, and the market makers must "call" one another for price quotes. We find a clear trade-off between liquidity and market efficiency due to the degree of public quote disclosure. In markets where all quotes are disclosed publicly, opening spreads are smaller, and volume is higher, but this market is much poorer in terms of price efficiency. Dealer prices are less responsive to new information, and pricing errors decline less rapidly over time than in the opaque microstructure. We explain the price efficiency result as dealers' optimal speculative response to reduced search costs in the transparent market.

The plan of the article is as follows. In section 1, we discuss the experimental methodology. In section 2, we analyze the experimental data, testing

several hypotheses, and offering an interpretation of the results. Section 3 concludes.

1. Experimental Design and Implementation

1.1 Experiment overview

In this article we test the effect of two quote disclosure regimes on market performance, especially the price discovery process. Our tests involve a computerized experimental securities market in which seven professional dealers act as market makers, with two computerized traders operating as non-market-making customers.³ The experiment occurs over eight trading rounds of five minutes each. Each round is independent, in the sense that the parameter settings are unique for each round; information about the security's value derived from one round is not relevant in another.

Each market maker has his or her own computer trading screen and a keyboard. A sample computer screen appears in Figure 1. Over the rounds we vary the amount of price information on the traders' computer screens. In the first four rounds (the opaque rounds), no quotes are publicly disclosed and market makers can obtain price information only by "calling" other market makers. In rounds 5 through 8 (transparent), all quotes are publicly disclosed at all times in the price queues in the center of the screen.

The microstructure of our experimental environment can be seen as a continuous, multiple-dealer version of the pure dealership market of Glosten and Milgrom (1985). In their model, a specialist sets bid-ask quotes for investors, who observe the quoted bid and ask prices and then decide whether to trade (one unit per transaction). The specialist is free to adjust prices at any time. The investors represent informed and liquidity traders and do not compete with the specialist, since they do not set limit orders in the market. The Glosten and Milgrom market is thus quote driven. Our experimental design differs from the Kyle (1985) framework in this sense.⁴ Instead, our experimental design is most similar to the quote-driven experimental markets used by Bloomfield (1996), Bloomfield and O'Hara (1999), and Flood et al. (1996, 1997). Flood et al. (1996), however, fail to incorporate liquidity traders, making their results less realistic than those reported here. Compared with the experimental design in Bloomfield (1996) and Bloomfield and O'Hara (1999), there are three significant differences. First, market makers in the Bloomfield (1996) and Bloomfield and O'Hara (1999) stud-

³ The experimental setup used here is similar to that used in Flood et al. (1996) and Flood et al. (1997). We use the terms "market makers," "traders" and "dealers" interchangeably to refer to the seven human subjects. Similarly, we use the terms "robots" and "customers" interchangeably to refer to the two computerized non-market-making customers.

⁴ See Madhavan (1992) for an overview of the differences between quote-driven and order-driven markets. With some modifications, the Kyle model is the underlying framework of the experimental markets used by Schnitzlein (1996) and Lamoureux and Schnitzlein (1997).

Your Position		Price Quotes			Transactions			
Money	: 790.00	ID	Sell	Buy	ID	#	buy sel	units price
LIFE Shares:	2					1	*	B 1 100.00
Your quote :						2	E *	1 110.00
bid: 90.00 ask: 110.00						3	*	J 1 110.00
Approximate wealth						4	*	R 1 110.00
based on last price:								
1010.00 esquires								
Time remaining: 158 sec.								

PageUp/PageDown = scroll

Figure 1
Trading screen

This is an example of a trading screen in the market without quote disclosure. Each trader has his or her own trading screen. The window on the upper left presents the cash balance (790 esquires), the inventory of this dealer (2 LIFE shares long), the current outstanding quote of this trader (90 - 110), and the approximate profit of this trader based on the price of the last transaction in which he or she was involved (1,010 esquires). The middle window on the left side shows the time remaining in this round. The black window in the center of the screen is where a quote appears when this trader calls another trader. Under the heading "ID Sell Buy ID," ID denotes the identity of the dealer presenting the quote; Sell denotes the quoted bid (at which this trader can sell); the Buy column contains the quoted ask (at which this trader can buy the share). In the markets where all quotes are disclosed publicly, all bids are presented below Sell ranked from highest to lowest, and all asks are presented below the heading Buy ranked from lowest to highest. Information of past transactions appears at the right of the trading screen. By default, it displays the details of the last 20 transactions. The trader can scroll through the list with the PageUp and PageDown keys. For all transactions, the identities of the buying (under heading buy) and the selling (sel) trader are displayed, along with the number of shares involved in the trade and the price at which the trade cleared. For example, the first row indicates that this trader (his identity is shown as an asterisk *) bought one LIFE share for 100 esquires from trader B. The ten traders' identities are denoted by letters ranging from "A" through "J". The robots are denoted with the letter "R". The fourth transaction is thus an example of a trade in which a robot sold to this trader one share for 110 esquires.

ies set quotes but do not initiate trades. In contrast, we allow for interdealer trading, which can be a salient characteristic of multiple-dealer markets.⁵ Second, our market is continuous instead of sequential, in the sense that market makers may trade and revise their quotes at any time; continuous

⁵ For example, in the foreign exchange market, Perraudin and Vitale (1996) report that the magnitude of interdealer trading is roughly 80% of total volume. In our experimental markets, interdealer volume ranges from roughly 85% to 95% of total volume.

trading has the notable benefit that it yields large amounts of data. Third, we use professional security traders as subjects rather than students.

1.2 Dealers' objectives and the true price of the asset

At the start of each trading round, each market maker is given an initial endowment of 1,000 esquires (a fictional numéraire currency). All traders are instructed to maximize their end-of-round wealth by trading for a single security. End-of-round wealth is expressed in esquires; it equals the cash balance at the end of the round plus the end-of-round inventory valued at the "true price" of the asset. The true price of the asset or the "underlying" value can be seen as an ex post liquidation value as used in, among many others, the Glosten and Milgrom (1985) and Kyle (1985) models. This true price differs across rounds to achieve independence of the rounds.

Before the start of the experiment, the market makers were told that the true price would be uniformly distributed between a minimum of 1 esquire and a maximum of 200 esquires; this is the only a priori information traders have about the true price at the start of every round. The dealers can obtain information about the true price through transactions with robots, as discussed below. Dealers are not instructed about possible trading strategies. Each dealer trades according to his or her own expectations and predilections. Traders can profit (or lose) over the round by buying and selling (i.e., jobbing) and/or by building a long or short inventory of the security that is to be converted to cash at the true price at the end of trading (i.e., speculating on the true price). There are no short-sale restrictions or penalties. There is no upper or lower limit on a dealer's allowable securities inventory. During each round, traders can trade at any time they desire in unit quantities.

The true prices are chosen as follows. First, recall that at the start of the experiment, we told the market makers the true price would lie between 1 and 200. Therefore a uniform prior expectation of the true price would be approximately 100 esquires. In the experiments we wanted to have instances of both upward and downward price discovery. Therefore we drew four prices randomly from the [1,100) interval and four prices from the (100,200] interval. Within these intervals we would like to have prices both far away and close to the uniform prior expectation of 100. Therefore we chose two prices uniformly from each of the following four intervals: [1,50), [50,100), (100,150], and (150,200]. The true prices are reported in Table 1.

1.3 Trading and robot behavior

Each dealer is obliged to enter a quote—both bid and ask prices—within 10 seconds of the start of each round (otherwise, a penalty of 10 esquires/second accrues).⁶ The dealers can change their quote at any time during the

⁶ Ex post, this penalty was assessed on only one occasion: dealer 3 in round 8 was eight seconds late in submitting his quote.

Table 1
Summary statistics

	No quote disclosure (opaque markets)				Public quote disclosure (transparent markets)			
	1	2	3	4	5	6	7	8
Round settings								
True price	69	163	134	15	22	75	118	185
Results								
Quotes set	80	72	64	66	76	83	80	88
Trades								
- Total	461	473	533	549	1,300	1,557	1,841	1,423
- Interdealer	387	399	459	475	1,226	1,483	1,769	1,349
- Robot	74	74	74	74	74	74	72	74
Uninformed (%)	57	38	47	46	50	46	51	51
Off-market (%)	36.2	23.6	21.1	26.7	0.0	0.0	0.0	0.0
Average end-of-round capital in esquires^a								
Market makers								
1	-486	1,881	834	-5,071	6,127	3,656	4,633	301
2	29	-6,869	-816	-3,409	2,459	1,536	-319	-2,745
3	-35	2,686	-5,894	6,855	-5,070	-21,166	-13,041	2,436
4	-282	-1,794	487	-2,265	15,074	6,012	119	-14,364
5	-376	466	169	882	-75,664	-8,302	787	13,651
6	234	-1,410	424	-2,954	12,436	6,311	5,565	-13,683
7	571	3,481	3,224	3,757	41,689	10,656	1,386	12,495
Average	-49	-223	-225	-315	-421	-185	-124	-273
Robots								
Informed	155	752	745	972	1,549	550	328	707
Uninformed	18	28	42	131	-74	99	108	248

^a Initial capital of 1,000 esquires is excluded. None of the average profits of the dealers and uninformed robots are significant. The average profits of the informed robots are all significant. Note that robot profits are stated as the average over the two robots.

round, but cannot otherwise withdraw their quotes. Therefore, at any time all traders have a quote outstanding in the market. The maximum individual spread is limited to 30 esquires.⁷

In the transparent market, all bids and asks outstanding are presented on the trading screen of every market maker. The bids (asks) are ranked from the highest (lowest) on top through the lowest (highest) at the bottom, so that the market spread is always obtainable at the top of the queues. To buy securities, the trader presses the "B" key; the trader buys automatically against the lowest ask available. To sell, the trader presses the "S" key; the trader sells automatically against the highest bid. As in Glosten and Milgrom (1985), we normalize the trade size to exactly 1 share per transaction; there is, however, no limit on the number of transactions. After the transaction is effected, the trade details are presented in the private transaction history window of both counterparties.

In the opaque market, information about the quotes of other market makers is, by definition, not publicly available. Instead, in order to obtain direct

⁷ This restriction was imposed to prevent dealers from effectively exiting the market by quoting infinite spreads. In practice, it was almost never binding.

price information and to trade, traders have to call one another bilaterally. This is done by pressing the “C” key, and then entering the identity of the trader whose quote the caller wishes to observe.⁸ The quote of the dealer called then appears on the caller’s trading screen for a maximum of seven seconds, during which time the caller can initiate a trade. To buy assets, the trader presses the “B” key, to sell the “S” key. After this the caller enters the number of shares he or she wants to trade. Again, we normalized the trade size to exactly 1 share per transaction (however, traders can initiate several one-share trades during the seven seconds that the quote is active). Note that the caller hits the quote of another market maker; the latter is not allowed to refuse the trade. When a dealer is called, he or she is not immediately informed about the call (the quoted price is simply taken from the nonpublic price queues). However, as soon as the caller initiates a transaction, the trade details are presented in the quoting dealer’s private transaction history window.

As discussed previously, two types of subjects are involved in the experiment. In addition to the seven (human) market makers, who both set quotes and trade, there are two computerized customers (robots) that trade in the market. The robots do not set quotes and cannot be called by the market makers. The two robots represent external customers of the exchange and proxy for both informed traders and liquidity or noise traders. The two robots are indistinguishable to the market makers. Each robot is programmed to trade every 7 seconds against the best prices in the market (on average, there is a robot trade every 3.5 seconds). Whether a robot trade is to be informed or uninformed is determined at random, just prior to each robot transaction; the probability that the trade is informed equals .5.⁹ The market makers are told this probability, and therefore know that, on average, half of the robot trades are informed.

If a robot initiates an informed trade (with probability .5), it buys (sells) if the lowest ask price (highest bid) outstanding at that time is below (above) the true price of the security. An informed robot does not trade if the best price equals the true price. Note that the informed robot tries to maximize its profit only at the trade level; over the whole round it is restricted to trading only at multiples of seven seconds. Since the informed robot trades are a function of the true price, these transactions convey information about the true price of the asset. If a robot initiates an uninformed trade, it is determined with probability .5 whether the robot sells or buys; if it sells (buys) it does so against the (highest bid price) lowest ask price available in both market structures.

⁸ The identity of traders is denoted with letters; {A, B . . . , G} refer to the seven market makers, while “R” refers to the robots. Robots do not quote prices, however.

⁹ This setup of the noise structure parallels the setup in Röell (1996). Note that the robots here always enjoy pretrade transparency, since they always transact at the best price available in the market in both the transparent and opaque markets.

When a market maker is involved in a robot transaction, he or she does not receive any special notification. That is, as for the interdealer trades, the transaction information appears in his or her private transaction history window with identity “R” for robot. Note that the trader does not know whether the robot trade was informed. Given their knowledge of the probability that a robot is informed, traders must filter the relevant price information from these robot transactions. If a robot trade is informed and the robot sells (buys), it is a clear signal that the bid price (ask price) of the trader is above (below) the true price of the asset.

In both the transparent and opaque microstructures, the transaction history window contains all interdealer and robot trades in which the particular trader is involved; we thus have a low-level posttrade transparent market. For all transactions listed in the transaction history window, the identities of the buyer and seller are presented, along with the transaction size (which always equals one share in our setup) and the price at which the transaction cleared. There is no delay in transaction disclosure.

1.4 Payoff to the subjects

At the end of each round, each trader is informed about the true price of the asset and his or her end-of-round profit. Esquire amounts are translated into Dutch guilders as follows. The dealer with the greatest end-of-round wealth receives 7 guilders, the second-best dealer receives 6 guilders, etc.¹⁰ The traders are informed about this payment arrangement before the start of the experiment.

1.5 Data

The data in this article were collected from the experiment described above, which was conducted at the Center for Research in Experimental Economics and Political Decision Making (CREED) at the University of Amsterdam on April 29, 1997. The subjects were seven professional traders from de Generale Bank, Optiver, and Rabobank. All subjects traded as market makers in eight independent rounds; i.e., information about asset values obtained in one round is irrelevant in other rounds.

Table 1 presents the settings and summary statistics. We started with four rounds without public quote disclosure, followed by four rounds in which all quotes were publicly disclosed. This order of microstructures might have caused learning effects over the rounds; i.e., skills formed in the first market might have enhanced trading behavior in the last four rounds. However, this

¹⁰ At the time of the experiments, the guilder-dollar exchange rate was roughly 1 USD = 1.90 NLG. We also gave the market makers a fixed amount equal to 125 guilders for showing up. It should be noted, however, that the dealers were strongly motivated by professional pride in their trading ability, a factor which should be roughly proportional to esquire trading profits and which may dominate their guilder payoffs as a behavioral incentive.

article mimics the experiments from Flood et al. (1996), with qualitatively very similar results. As noted, the difference with those experiments is that we include liquidity motivated traders, where Flood et al. (1996) did not. Furthermore, we reverse the sequence of the market structures relative to Flood et al. (1996), who started with public price disclosure and ended with opaque markets. Since the results we present in this article are qualitatively similar to theirs, we feel comfortable about the robustness of the results.

From Table 1 we find, on average, 76 quote settings and 430 interdealer transactions in the opaque markets, and 82 quote settings and 1,457 interdealer transactions in the transparent markets.¹¹ There is an average of 74 robot transactions initiated per round in both market types. Table 1 also shows that, on average, dealers made small losses (although the zero-sum nature of the experiment means that if some market makers make large gains, others generally make large losses). On average, the market makers make slightly negative profits, largely because they must trade with informed robots, to whom they always lose. In principle, market makers might recover these adverse-selection losses by earning the spread, especially as information revelation reduces the magnitude of adverse-selection losses. However, on net, the adverse-selection losses are not fully recovered. Note that market makers' average losses do not differ systematically across the transparent and opaque market settings—these are determined mostly by informed robots' profitability; however, the dispersion of profits is higher in the transparent market, reflecting more intensive interdealer trading there.

2. Results

In this section we test three hypotheses concerning the effects of quote disclosure on market performance. The tests are based on the experimental data drawn from seven professional traders over eight rounds, as discussed in the preceding section. In brief, we find a marked trade-off between liquidity and price efficiency. We conclude that both the higher liquidity and slower price discovery of the more transparent market are the result of speculating dealers exploiting the lower search costs in this market. It is worth noting that the results reported here correspond qualitatively with those found in two pilot experiments.

2.1 Volume

Regarding liquidity, we first note the sharp discrepancy in trading volume between the transparent and opaque markets. The average trading volume in the markets with no public quote disclosure was 504, while the average for the transparent markets was roughly three times higher at 1,530. We at-

¹¹ This higher number of transactions in the market where all quotes are publicly disclosed is largely due to the trading mechanism, since prices are more difficult to obtain in the markets without quote disclosure, which increases the time needed to initiate a transaction.

tribute this difference directly to transparency: dealers spent relatively less time sampling prices (and correspondingly more time engaging in transactions) due to reduced search costs in the transparent market. Note that there is both an informational component (knowing more prices sooner) and a logistical component (less time is spent canvassing the market and identifying the counterparty for a transaction) to this reduction in search costs. Both components play a role in the increased volume.

The fact that search costs are a significant factor in the opaque market is revealed in the proportion of executed interdealer trades that fail to find the best currently available price (we refer to these as “off-market prices”). We presume that a dealer would not intentionally trade at an off-market price. Thus, in the transparent market, all trades automatically go to the dealer whose price is at the top of the relevant queue, implying no off-market trades there. In the opaque market, on the other hand, the market spread is not automatically hit, and off-market trades do occur. The proportion of trades occurring at off-market prices is in the 20% to 30% range for the opaque market (see Table 1). The flip side of the fact that off-market prices can attract trades is the fact that the best price sometimes fails to get the trade in the opaque market. As a result, dealers tend to reprice more aggressively in the opaque market, a fact examined more closely below.

A noteworthy characteristic of these markets is the large proportion of trading that is interdealer. This fact is related to speculative position-taking by the dealers. The basic speculative force at work in our experiments is obvious—dealers should buy when they see robots buy, sell when they sell (even lacking a robot trade, a dealer may nonetheless infer something about the robots’ behavior by observing his own transactions with other dealers). However, implementing this basic speculative tactic is complicated by the zero-sum property that imposes pronounced patterns among participant inventories. That is, any large position taken by an individual dealer must be absorbed somewhere else in the system; since robot positions are constrained by their predefined transaction rate, this implies that the bulk of any dealer position must be balanced by a countermanding position with another dealer. Thus a dealer’s inventory may not reflect his or her speculative intentions. This is apparent in Table 1, which shows individual dealers’ profitability in each round. While individual dealers’ profits or losses can be enormous (e.g., consider round 5), the cross-sectional average is always quite small. Indeed, it seems that the high level of interdealer trade results largely from a struggle among the dealers to keep their positions on the right side of the market, a fight that not all of them can win.

2.2 Spreads

In financial markets, the costs of market making and information asymmetry among market participants are seen as the primary motivations for the bid-ask spreads. The theoretical implications of information asymmetries for the

bid-ask spread are quite sensitive to the microstructure under consideration. There is a broad consensus that a dealer's spread will widen to compensate the dealer for losses to informed traders. At the same time, certain dynamic strategies—for example, Madhavan (1995)—call for the dealer to narrow spreads in early trading in order to attract informed order flow, so that the information so learned can be exploited in subsequent trading. On the other hand, Leach and Madhavan (1993) compare dynamic dealer pricing strategies in monopoly-dealer and multiple-dealer markets. They conclude that a monopoly specialist will widen, rather than narrow, spreads to better filter information from their trades: wider spreads are more likely to drive away noise traders than informed traders. Competing dealers do not engage in price experimentation, since the signal-to-noise ratio is only improved by trading at wider spreads, but this is impossible since widening one's spread drives order flow to the other dealers.

Madhavan's (1995) analysis would suggest that our dealers should narrow spreads to attract informed order flow and widen them afterward. Leach and Madhavan's (1993) model would suggest the opposite. Notably, however, our market is very close to the setup described in Leach and Madhavan (1993), with the important differences that we allow for interdealer trading, whereas they do not, and their market is fully pre- and posttrade transparent, whereas ours is not. Interdealer trading is important, because our dealers can exploit information from "purchased" robot order flow against their fellow market makers, whereas Leach and Madhavan's cannot. To summarize, then, we propose the following null hypothesis:

Hypothesis H1₀. *Dealer spreads are narrowed to attract informed (i.e., robot) trades.*

We test this hypothesis by examining the response of dealer spreads to robot transactions. If the spread was narrowed specifically to attract informed order flow, then the spread should widen after the event, since subsequent robot trades should be significantly less informative than the first such trade. We define $s(i, j, r)$ as the log spread size for the i th quote set by trader j in round r , and $\Delta s(i, j, r)$ as the change in spreads between quotes i and $(i - 1)$. We limit the sample of events to quote changes that occurred within 10 seconds after a robot trade. The rationale for restricting the event window is to reduce the likelihood that contaminating events would bias the results.¹² We thus estimate the following regression:

$$\Delta s(i, j, r) = \alpha + \alpha_j + \beta I(r \in QD) + \epsilon(i, j, r), \quad (1)$$

¹² To check the robustness of this procedure for the possible introduction of contaminating factors in the event window, we reestimated the regression for 20-, 30-, and 60-second windows. The results are qualitatively similar. In particular, all parameter estimates have the same signs and order of magnitude. However, the absolute size of the estimated coefficients and the associated R^2 's decreases gradually as the event window grows, suggesting that the larger window allows more noise, but not contaminating events.

Table 2
Spread-size responses

	Intercept (α)	Market structure (β)
Whole trading period (start at 0 through 300)		
10-second event window	0.635 (0.046)	-0.110 (0.089)
Half trading period (start at 0 through 150)		
10-second event window	0.508 (0.078)	-0.025 (0.155)

This table contains the estimates of the coefficients from the fixed individual effects panel model [Equation (1)]. Here the individual spread adjustment (changes in log spread) following robot trades is regressed on a constant (α), dealer-specific dummies (α_j), and a market structure indicator (β) that equals 1 for the transparent markets and 0 otherwise. The sample is restricted to spread changes occurring within 10 seconds after a dealer's prices are hit by a robot trade. α thus represents the widening of the spread after a robot trade.

Standard errors in parentheses.

For the whole trading period, $R^2 = 0.173$ and $n = 204$.

For the half trading period, $R^2 = 0.183$ and $n = 107$.

where α_j are trader-specific constants (i.e., a fixed individual effects panel model), $I(\cdot)$ is an indicator function that equals 1 if the condition in parentheses is true and 0 elsewhere, QD is the set of all rounds with public quotes disclosure, and ϵ is the usual error term. According to $H1_0$, α should be positive and significant. Estimates of Equation (1) appear in Table 2.

The results confirm the null hypothesis, that dealers widen their spreads after trading with robots. The α estimates are all positive and significant. The implication is that dealers indeed attempt to purchase informed order flow, at least until they achieve a potentially informative robot trade. The parameter estimates are not strikingly different when estimated only over the first 150 seconds of trading. The β estimates are all negative but insignificant, implying this behavior does not differ strongly across transparency regimes. The negative sign on β indicates that spread adjustments are smaller in the transparent market; this is consistent with the evidence presented below of more aggressive pricing strategies in the opaque market.

We also consider the broader evolution of spreads throughout the trading round. In the literature on posttrade transparency, the general consensus is that opening spreads in transparent markets are wider, but that the difference in spreads between transparent and opaque markets disappears over time.¹³ Madhavan (1995) rationalizes this by stating that market makers compete for order flow to obtain information; they therefore quote more competitive

¹³ See Bloomfield and O'Hara (1999) and Flood et al. (1997) for experimental results. Röell (1996) analyzes theoretically the effects of transparency on spreads. Gemmill (1996) presents an empirical investigation of the London Stock Exchange, but finds mixed results.

spreads in (posttrade) opaque markets. This competition is reduced in transparent markets, since information is disclosed publicly, and quoted spreads therefore widen.

Thus the disclosure of transaction information has substantial effects on spread sizes, but there is less evidence on the effects of pretrade quote disclosure on spreads. Bloomfield and O'Hara (1999) address this latter issue, but find no effect of quote disclosure on (opening) spreads. However, since spreads arise to cover the operational and asymmetric information costs of market making, one would expect different quote disclosure regimes to have different impacts on spreads. For instance, the higher search costs imposed in a market with a low level of public price disclosure should result in higher spreads. (It is worth emphasizing that the literature predicts that pretrade transparency and posttrade transparency will have opposite effects on spreads. Increasing pretrade transparency narrows opening spreads, whereas an increase in posttrade transparency leads to wider opening spreads.)

In addition, disclosing quotes publicly reduces the asymmetric information problem, and should therefore lead to smaller spreads. This creates a second-order effect. As trading evolves over the trading period, more information is shared by the market makers and the informational differences between market structures diminish. This would imply that the differences in spreads, due to information asymmetry, would decline over the trading period. Combining both effects, we test the following hypothesis:

Hypothesis H2₀. *Dealer spreads are wider in a market without public price information than in a market in which all quotes are publicly disclosed; furthermore, the difference in spread size declines over time.*

To examine the spread differences between microstructures, we split each trading round in 30 intervals of 10 seconds each. In Figure 2 we present the average spread size over all market makers and all rounds for all intervals in both the market with public quote disclosure and without public information on quotes. As Madhavan (1992) predicts for quote-driven markets, the average dealer spread shrinks over the trading period in both market types, which we attribute to declining information asymmetries as more transactions are settled. Furthermore, Figure 2 indicates clear differences in spreads between both market types, as stated under the null hypothesis. The opening spreads in the rounds without quote disclosure are significantly wider than in the rounds with public information on quotes.¹⁴ Over the trading period,

¹⁴ Significance is at a 10% confidence level in approximately the first 120 seconds. We compared the average spreads of each market maker over both market structures. At each interval we therefore have two samples consisting of seven matching spreads, that is, one for each market structure. Then we conducted a two-sample sign test [see Bain and Engelhardt (1987) for more details]. This test compares the medians of both samples, with the advantage that we do not need to assume normality of the spreads. We do, however, need to assume independence, but this is likely given that we are comparing two samples from different market structures.



Figure 2

Average dealer spreads over time

This figure presents the time-series patterns of the average dealer spreads, in esquires, over 30 intervals of 10 seconds each. Each line represents the average of the spreads over rounds with the same market structure, defined as the difference between the quoted ask and bid price, over all quotes outstanding during each interval.

the market-maker spreads narrow in the markets without quote disclosure on quotes such that eventually no differences in spread size are apparent. Both aspects of the null hypothesis are supported by the data.

2.3 Price efficiency

Intuitively, one expects that the more price information is available, the more information will be impounded in dealers' quotes. O'Hara (1995, p. 270), for example, makes this case, as noted above in the introduction. In other words, price transparency improves price discovery. In their experiments, Bloomfield and O'Hara (1999) do not find a significant effect of quote disclosure on pricing errors. These errors, defined as the absolute difference between the true price of the asset being traded and the midpoint of the market spread, decline over the trading period, but Bloomfield and O'Hara find no significant evidence that they decline more rapidly in markets where all quotes are publicly disclosed. On the other hand, Flood et al. (1996) show that the price discovery process is significantly slower in markets where all quotes are publicly disclosed than in markets without quote disclosure. The




Figure 3

Average price errors over time

This figure presents the time-series patterns of the price error, averaged over all dealers and rounds with the same market structure at each second during the trading round.

experimental market they used is similar to the market discussed here, but without noise traders; their market makers were confronted instead only with informed customers. Nonetheless, we state a null hypothesis to conform to the intuitive consensus:

Hypothesis H3₀. *Public disclosure of dealer quotes enhances price efficiency.*

To examine the effects of quote disclosure on prices, we first consider the time path of dealer pricing errors, defined as the absolute difference between the midpoint of a quote and the true price of the asset. In Figure 3 we plot the average price error over all dealers and rounds with the same market structure for each second. We discard the first 10 seconds, since this is the timespan we allowed the market makers to enter their initial quotes. The price errors decline over time and we observe that these price errors decline more rapidly in the market without public quote disclosure.

Figure 3 clearly suggests differences in price efficiency between the two market structures. To examine this more formally, let $P(j, t, r)$ be the

price error obtained from the quote of market maker j in round r at time t (measured in seconds). We test whether the pricing errors decline more rapidly in a transparent market by regressing the pricing errors on a time trend variable, t , whereby the effect of quote disclosure is split over two slope coefficients.

$$P(j, t, r) = \alpha + \alpha_j + \beta_1 I(r \in QD) + \beta_2 t I(r \in QD) + \beta_3 t I(r \in NQD) + \epsilon(j, t, r), \quad (2)$$

where $I(\cdot)$ is an indicator function that equals 1 if the condition in parentheses is true and 0 elsewhere, t is the time trend, QD is the set of all rounds in which quotes are publicly disclosed, NQD is the set of rounds in which quotes are not disclosed, and ϵ is an i.i.d. error term. Note that we pooled the data over all eight rounds and the seven market makers in the regression.¹⁵ We added specific constants α_j (individual effects) to control for the specific characteristics of trader j . The effect of quote disclosure on the price errors is captured by both the initial level effect—measured by the market structure dummy (β_1)—and by the difference between the slope coefficients for the markets where all quotes are disclosed (β_2) and where no quotes are disclosed (β_3). Table 3 contains the estimates of these slope coefficients.

From Table 3 we first observe that all parameter estimates are significant except for the market structure dummy. The constant term equals 53, which equals the average absolute true price deviation from 100, which is the initial market maker's expectation of the true price when no price information has entered the market. The market structure dummy suggests that the initial level of the price errors is not significantly different between both market structures. Since these initial price errors are equal for both market structures, we use the differences in the speed with which the price errors are corrected, captured by the slope coefficients of the linear trends in Equation (1), as a measure of relative price efficiency. The estimates for both slope coefficients β_2 and β_3 are negative and significant, indicating that the pricing errors decline over time, reflecting the accumulation of information over time as trading proceeds in both markets. The difference between the two time trend coefficients equals 0.04, indicating that pricing errors decline more rapidly (the slope is more negative) in markets where quotes are not publicly disclosed. Using price errors as a measure of efficiency, we conclude that markets without pretrade information are more efficient than markets with public prices, thus rejecting Hypothesis H2₀. On the other

¹⁵ Note that if price discovery (defined here as quoted prices within three esquires of the true price) had been achieved in only some of the rounds, then a test with pooled data would be biased, since market-makers' pricing strategy would likely change in the vicinity of the true price. However, the dealers never achieved full price discovery. Furthermore, note that robot trades can only indicate the direction of the true price (higher or lower) to the dealer, so that variation in absolute levels of the true prices across rounds should not bias the test either.

Table 3
Price efficiency

	Intercept (α)	Market structure (β_1)	Public quote disclosure (β_2)	No public disclosure (β_3)
Estimated slope coefficient	53.169 (1.926)	0.528 (3.426)	-0.098 (0.010)	-0.134 (0.007)
Difference in slopes ($\beta_2 - \beta_3$)			0.037 (0.012)	

This table contains the estimates of the slope coefficients in the fixed individual effects model [Equation (2)]. Here the time path of price errors of individual market makers is regressed on a constant, on a market structure dummy (β_1) that equals 1 if the market has public quote disclosure and 0 elsewhere, and on two slope coefficients (β_2 and β_3) for the time trend in the rounds with and without public disclosure of quotes, respectively. The individuals in the panel are the market makers, so that we correct for cross-sectional differences between all market makers.

The data are pooled over seven market makers and eight rounds; $n = 16,800$.

For details on estimating fixed individual effects panel models, see Baltagi (1995).

Robust standard errors for a general variance-covariance matrix are presented in parentheses [see Arellano (1987)].

$R^2 = 0.149$.

hand, this result is in line with Flood et al. (1996), and not inconsistent with Bloomfield and O'Hara (1999) who find no significant relationship between quote disclosure and pricing errors.

To confirm this rather striking result, we examine average dealer price responses to new information. If prices indeed impound new information, then this fact should be evident in dealer quote revisions following information arrivals. We examine the magnitude of such quote revisions using the most clear-cut informational innovations for a dealer, namely robot trades. For each dealer j in round r , we collected all quote revisions, $P(t)$, where P is the new midpoint of the dealer spread and t ranges from 1 (the initial quote posting) up to the total number quote revisions by trader j . From this we calculate all log-price changes $\Delta P(t) = \ln P(t) - \ln P(t-1)$ ($t = 2, \dots$). Using the number of robot trades in the period $(t-1, t]$, we construct a dummy variable, $R(j, t, r)$, that equals 1 if a robot bought from trader j in period $(t-1, t]$, -1 if a robot sold to j , and 0 otherwise. Furthermore, we define indicator variables— $I(r \in QD)$ and $I(r \in NQD)$ —as before to represent the transparency level in the market. Finally, we pool all rounds and dealers in a panel dataset, controlling for dealer-specific (e.g., skill) effects by allowing for an individual-specific intercept, α_j , and testing for the differences in price responses between the two market structures in a fashion similar to Equation (1):

$$\Delta P(j, t, r) = \alpha + \alpha_j + \beta_1 I(r \in QD) + \beta_2 I(r \in QD) R(j, t, r)$$

Table 4
Price responses

	Intercept (α)	Market structure (β_1)	Public quote disclosure (β_2)	No public disclosure (β_3)
Estimated slope coefficient	-0.015 (0.004)	0.032 (0.007)	0.063 (0.009)	0.119 (0.013)
Difference in slopes ($\beta_2 - \beta_3$)			-0.056 (0.016)	

This table contains the estimates of the slope coefficients in the price response fixed individual effects panel model [Equation (3)]. Here the individual price adjustment (log price returns) following robot trades is regressed on a constant (α), on a market structure dummy (β_1) that equals 1 if the market has public quote disclosure and 0 otherwise, a robot trade dummy that equals +1 for a robot purchase and -1 for a robot sale, times an indicator variable for the transparent markets (β_2), or times an indicator for the opaque markets (β_3). β_2 and β_3 thus represent the average price increase (decrease), measured as log price relatives, following a robot purchase (sale). The individuals in the panel are the market makers, so that we correct for cross-sectional differences between all market makers.

The data are pooled over seven market makers and eight rounds; $n = 553$. For details on estimating fixed individual effects panel models, see Baltagi (1995).

Robust standard errors for a general variance-covariance matrix are presented in parentheses [see Arellano (1987)].

$R^2 = 0.227$.

$$+ \beta_3 I(r \in NQD) R(j, t, r) + \epsilon(j, t, r) \quad (3)$$

The results appear in Table 4. All parameter estimates are significant. The intercept α is negative, whereas the market-structure dummy (β_1) is positive. The sum of both is a function of the level of the true prices relative to the initial price expectation of 100, since these parameters reflect the average price adjustment. Regarding the price response coefficients (β_2 and β_3), they are both positive and significant as expected, indicating that dealers do impound robot trade information appropriately in both microstructures, raising prices when robots buy, and lowering prices when robots sell. More importantly, the difference between the transparent and opaque markets, $\beta_2 - \beta_3$, is negative and significant, implying that dealers in the opaque market are more aggressive in responding to new information. This is entirely consistent with the results in Table 4, again leading to a rejection of the null hypothesis.

3. Summary and Conclusions

While the issue of posttrade transparency has received considerable attention in the literature, the degree of pretrade transparency is also of central importance. Securities markets worldwide employ a bewildering variety of degrees of public quote disclosure, with implications that are not well

understood. We consider the impact of quote disclosure in a continuous multiple-dealer market in an experimental setting. Seven professional market makers set quotes and trade with other market makers and customers that represent both informed and liquidity motivated traders. Over multiple rounds we vary the amount of quote information that is publicly available between two extremes: a market where all quotes are publicly disclosed and a market with no public quote disclosure. Our study represents an advance, in that we provide a direct comparison of alternative transparency arrangements, with a focus on multiple-dealer interactions. An experimental methodology is used to reduce the need for restrictive behavioral and *ceteris paribus* assumptions. A continuous experimental market also generates large amounts of data and permits us to examine more closely the relationship between the price and inventory components of a dealer's strategy.

We can summarize our results as follows. Pretrade transparency significantly reduces search costs, thus alleviating some uncertainty and facilitating trade. As a result, market liquidity, measured by spreads and volume, is greater in the transparent market: opening spreads are smaller and inter-dealer trading volume is much higher. Dealers learn about the underlying price both directly, by trading with (potentially) informed robots, and indirectly, by observing changes in the prices quoted by other dealers. Dealers actively attempt to attract robot trades by quoting relatively tight spreads until a robot trade is achieved. Negligible search costs in the transparent market imply that speculating dealers see clearly how to shade quotes to avoid being picked off in a trade on the "wrong" side of the market and how to keep their price competitive on the "right" side. The important and counterintuitive upshot is that rational speculating dealers use less aggressive price adjustments in the transparent market, thus slowing price discovery. Conversely, in the opaque market, search costs prevent accurate observation of competing prices, so that finely tuned quote shading is discouraged, and price adjustments are significantly more aggressive. The overall result is a distinct trade-off between liquidity and price efficiency.

Our explanation for this trade-off revolves around search costs. Pretrade transparency reduces search costs, thus alleviating some uncertainty and facilitating trade. Market liquidity, measured by spreads and volume, is therefore greater. At the same time, however, these reduced search costs imply that dealers can (optimally) more narrowly fine-tune their price improvements. The upshot is that rational speculating dealers use less aggressive price adjustments, thus slowing price discovery. In light of these results, it seems plausible that some combination of both microstructures could offset the disadvantages of each individually. We note with interest that curious combinations of microstructures do exist in practice (direct versus brokered foreign exchange trading or upstairs versus floor trading of common stocks). Further research is warranted to gain insights into the performance of such combined market structures.

References

- Admati, A. R., 1991, "The Informational Role of Prices: A Review Essay," *Journal of Monetary Economics*, 28, 347–360.
- Arellano, M., 1987, "Computing Robust Standard Errors for Within-Group Estimators," *Oxford Bulletin of Economics and Statistics*, 49, 431–434.
- Bain, L. J., and M. Engelhardt, 1987, *Introduction to Probability and Mathematical Statistics*, Duxbury Press, Boston.
- Baltagi, B. H., 1995, *Econometric Analysis of Panel Data*, Wiley, New York.
- Bloomfield, R., 1996, "Quotes, Prices, and Estimates in a Laboratory Market," *Journal of Finance*, 51, 1791–1808.
- Bloomfield, R., and M. O'Hara, 1999, "Market Transparency: Who Wins and Who Loses?" *Review of Financial Studies*, 12, 5–35.
- Bollerslev, T., and I. Domowitz, 1993, "Some Effects of Restricting the Electronic Order Book in an Automated Trade Execution System," in D. Friedman and J. Rust (eds.), *The Double Auction Market: Institutions, Theories, and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Vol. XIV, Addison-Wesley, New York, 221–252.
- Davis, D. D., and C. A. Holt, 1993, *Experimental Economics*, Princeton University Press, Princeton, N.J.
- Duxbury, D., 1995, "Experimental Asset Markets within Finance," *Journal of Economic Surveys*, 9, 331–371.
- Flood, M. D., R. Huisman, K. Koedijk, and R. Mahieu, 1996, "Price Discovery in Multiple Dealer Financial Markets: The Effect of Pre-Trade Transparency," working paper, Limburg Institute for Financial Economics, Maastricht University.
- Flood, M. D., R. Huisman, K. Koedijk, and A. Röell, 1997, "Post-Trade Transparency in Multiple Dealer Financial Markets," working paper, Limburg Institute for Financial Economics, Maastricht University.
- Gemmill, G., 1996, "Transparency and Liquidity: A Study of Block Trades on the London Stock Exchange under Different Publication Rules," *Journal of Finance*, 51, 1765–1790.
- Glosten, L. R., and P. R. Milgrom, 1985, "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 14, 71–100.
- Goodhart, C. A. E., and M. O'Hara, 1994, "High Frequency Data in Financial Markets: Issues and Applications," working paper presented at the First International Conference on High-Frequency Data in Finance, Olsen and Associates, Zurich, March 1994.
- Guillaume, D. M., et al., 1994, "From the Bird's Eye to the Microscope: A Survey of New Stylized Facts of the Intra-Daily Foreign Exchange Markets," working paper, Olsen and Associates, Zurich.
- Harris, L., 1993, "Consolidation, Fragmentation, Segmentation and Regulation," *Financial Markets, Institutions, and Instruments*, 2, 1–28.
- Kyle, A. S., 1985, "Continuous Auctions and Insider Trading," *Econometrica*, 53, 1315–1335.
- Lamoureux, C. G., and C. R. Schnitzlein, 1997, "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.
- Leach, J. C., and A. Madhavan, 1993, "Price Experimentation and Security Market Structure," *Review of Financial Studies*, 6, 375–404.
- Madhavan, A., 1992, "Trading Mechanisms in Securities Markets," *Journal of Finance*, 52, 607–641.

Madhavan, A., 1995, "Consolidation, Fragmentation, and the Disclosure of Trading Information," *Review of Financial Studies*, 8, 579–603.

O'Hara, M., 1995, *Microstructure Theory*, Basil Blackwell, Cambridge, Mass.

Pagano, M., and A. Röell, 1996, "Transparency and Liquidity: A Comparison of Auction and Dealer Markets with Informed Trading," *Journal of Finance*, 51, 579–611.

Perraudin, W., and P. Vitale, 1996, "Interdealer Trade and Information Flows in a Decentralized Foreign Exchange Market," in J. Frankel, G. Gallie, and A. Giovannini (eds.), *The Microstructure of the Foreign Exchange Market*, University of Chicago Press, Chicago, 183–201.

Röell, A., 1996, "Stock Market Transparency," working paper,

Schnitzlein, C. R., 1996, "Call and Continuous Trading Mechanisms Under Asymmetric Information: An Experimental Investigation," *Journal of Finance*, 52, 613–636.

Stoll, H. R., 1992, "Principles of Trading Market Structure," *Journal of Financial Services Research*, 6, 75–107.

Yu, G. G., 1993, "Information Revelation and Aggregation in Financial Markets," *Financial Markets, Institutions and Instruments*, 2, 29–54.