Overview: market structure issues in market liquidity

Maureen O'Hara[1]

The behaviour of prices and even the viability of markets depend on the ability of the trading mechanism to match the trading desires of sellers and buyers. This matching process involves the provision of market liquidity. The role of the market maker in providing liquidity is widely recognised, but liquidity can also arise from other aspects of the trading mechanism. In particular, rules and market practices governing the trading process, such as how trading orders are submitted and what trading information must be disclosed, can affect the creation of liquidity. This raises the question of whether changes in market structure can enhance the provision of liquidity. Is there a "Golconda exchange" that provides optimal liquidity?

What is microstructure?

Issues related to market liquidity are part of a broader analysis of the microstructure of markets. Market microstructure refers to the study of the process and outcomes of exchanging assets under a specific set of rules. While much of economics abstracts from the mechanics of trading, microstructure theory focuses on how specific trading mechanisms affect the price formation process.[2]

Much of the microstructure literature has focused on the *price-setting problem* confronting market intermediaries. The Walrasian auctioneer provides the simplest (and oldest) characterisation of the price-setting process. The auctioneer announces a potential trading range, and traders determine their optimal order at that price. If there are imbalances in traders' demands and supplies, a new potential price is suggested, and traders then revise any orders. No trading takes place until a market-clearing price is found. The London gold fixing loosely resembles the Walrasian framework, but most other markets differ dramatically. In particular, specific market participants play roles far removed from the passive one of the auctioneer. Demsetz (1968) was one of the first economists to analyse how the behaviour of traders affects the formation of prices. Demsetz argued that while a trader willing to wait might trade at the single price envisioned in the Walrasian framework, a trader not wanting to wait could pay a price for immediacy, ie liquidity. This results in two equilibrium prices. Moreover, since the size of the price concession needed to trade immediately depends on the number of traders, the structure of the market could affect the cost of immediacy and thus the market-clearing price.

The price-setting problem examined by Demsetz has been investigated more formally using inventory-based models. These models view the trading process as a matching problem in which the market maker - or price-setting agent - must use prices to balance supply and demand across time. There are several distinct approaches to modelling how prices are set by market makers: Garman (1976) focused on the nature of order flow; Stoll (1978) and Ho and Stoll (1981) examined the optimisation problem facing dealers; and Cohen, Maier, Schwartz and Whitcomb (1981) analysed the effects of multiple providers of immediacy. Common to each of these approaches are uncertainties in order flow, which can result in inventory problems for the market maker and execution problems for traders.

An alternative approach to modelling the behaviour of prices focuses on the *learning problem* confronting market intermediaries. Starting with Kyle (1984, 1985), Glosten and Milgrom (1985) and Easley and O'Hara (1987), market structure research has given greater attention to the effect of asymmetric information on market prices. If some traders have superior information about the underlying value of an asset, their trades could reveal what this underlying value is and so affect the behaviour of prices.

The key to extracting information from order flows is Bayesian learning. Each trader has a prior belief about the true value V of an asset. Traders observe some data, say a trade, and then calculate the probability that V equals their prior belief given that these data have been observed. This conditional

[1] Johnson Graduate School of Management, Cornell University and President-elect of the American Finance Association. Special thanks to Philip Wooldridge at the Bank for International Settlements for transcribing this presentation.

[2] For a survey of the literature, see O'Hara (1995) or Madhavan (2000). Lyons (forthcoming) provides a comprehensive review of the microstructure of foreign exchange markets.

probability incorporates the new information that traders learned from observing the data, and is hence their posterior belief about V (Graph 1). The posterior then becomes the new prior, more data are observed, and the updating process continues.

Graph 1

Bayesian learning

In information-based models, the solution to this learning problem determines the prices set by market makers. The ask price a_t equals the expected value of V given that a trader wishes to buy, and depends on the conditional probability that V is either lower ($V = \underline{V}$) or higher ($V = \overline{V}$) than the market maker's prior belief given that a trader wishes to buy. The bid price b_t is defined similarly given that a trader wishes to sell. An important characteristic of these prices is that they explicitly depend on the probability of a sale or buy (Graph 2). If uninformed traders are assumed equally likely to buy or sell whatever the information, good news ($V = \overline{V}$) will result in an excess of buy orders as informed traders decide to buy. Likewise, bad news ($V = \underline{V}$) will result in an excess of sell orders as informed traders decide to sell.

Graph 2

Dealer pricing

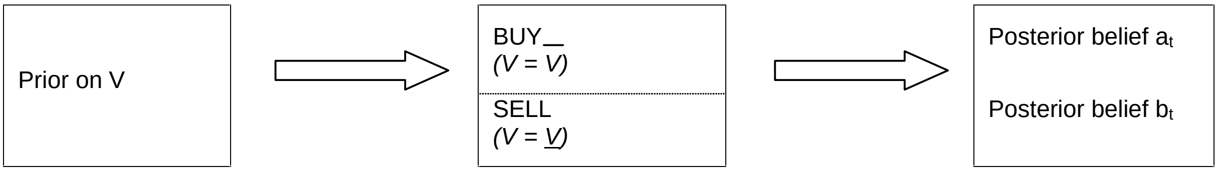

What we have learned

The information-based approach has greatly enhanced our understanding of the behaviour of markets and by extension the nature of market liquidity. Perhaps the greatest insight of this approach is how information affects quotes and spreads. Information-based models highlight the role of market parameters such as the size of the market or the ratio of large to small trades in the adjustment of prices. This in turn provides an explanation for the existence of bid-ask spreads even in competitive markets, without reference to explicit transactions or inventory costs. Inventory-based explanations of the bid-ask spread are problematic because empirical evidence of inventory effects in financial markets is weak.

Another important conclusion is that prices ultimately converge to their true, full-information value; in the limit markets are strong-form efficient.[3] This follows from the Bayesian learning process. It is not entirely clear, however, what market efficiency means in a dynamic setting. Given that some traders have superior information, prices along the adjustment path do not exhibit strong-form efficiency, and indeed there can be very great differences in the speed with which prices move toward full-information levels. Markets with greater volume, for example, adjust faster (in clock time) to information. The time between trades, in particular the tendency for transactions to cluster, also appears to affect the adjustment of prices.

The time varying process by which transactions arrive has important implications for econometric modelling of market volatility. Generalised Autoregressive Conditional Heteroscedasticity (GARCH) models and Autoregressive Conditional Duration (ACD) models have come to be widely used for analysing price and transactions data, respectively.

[3] Following the categorisations of the efficient market hypothesis used by Fama (1970), weak-form efficiency assumes that security prices fully reflect all security-market information, semi-strong form efficiency assumes that security prices fully reflect all publicly available information, and strong-form efficiency assumes that security prices fully reflect all information from public and private sources.

Finally, much has been learned about the information contained in specific trades. Different types of trades seem to have different information content. Similarly, trades in different markets seem to have different information content.

What we still do not know

For all that we have learned, there remain several puzzling issues concerning the trading process. Foremost is what determines volume. While empirical research has identified a strong link between volume and price movements, it is not obvious why this should be so. Volume may simply be a consequence of the trading process; whereas individual trades cause prices to change, volume per se may not affect prices. Or as seems more likely, volume could reveal underlying information, and thus be a component in the learning process. Pfleiderer (1984), Campbell et al (1991), Harris and Raviv (1993), Blume et al (1994), and Wang (1994) have examined this informational role.

A second set of issues revolves around what the uninformed traders are doing. It is the uninformed traders who provide the liquidity to the informed, and so understanding their behaviour can provide substantial insight and intuition into the trading process. Information-based microstructure models typically assume that uninformed traders do not act strategically. Yet, if it is profitable for informed traders to time their trades, then it must be profitable for uninformed traders to do so as well. Admati and Pfleiderer (1988, 1989), Foster and Viswanathan (1990), Seppi (1990) and Spiegel and Subrahmanyam (1992) among others have applied a game-theoretic approach to modelling the decisions of uninformed traders. A common outcome with this approach, however, is the occurrence of multiple equilibria.

Another open question is what traders can learn from other pieces of market data, such as prices. Neither sequential trade models such as Glosten and Milgrom (1985) nor batch trading models such as Kyle (1985) allow traders to learn anything from the movement of prices that is not already in their information set. But in actual asset markets the price elasticity of prices appears to be important. Technical analysis of market data is widespread in markets, with elaborate trading strategies devised to respond to the pattern of prices.

Finally, microstructure theory has not yet convincingly addressed how the existence of more than one liquidity provider in more than one market setting affects the price adjustment process. Much of the literature assumes the existence of a single market-clearing agent. However, alternative mechanisms could arise that divert order flow away from the specialist. Multi-market linkages introduce complex and often conflicting effects on market liquidity and trading behaviour. Indeed, it is not even obvious whether a segmented market equilibrium is sustainable. Current models of liquidity, for example, suggest that securities markets may have an inherent disposition toward being natural monopolies. Further research in this area is particularly important given the rapid increase in the number of electronic exchanges in recent years.

Market structures

Markets are currently structured in a myriad of ways, and new market-clearing mechanisms are arising with surprising frequency. All trading in a particular security can be directed to a single specialist, who is expected to make a market in that security. The New York Stock Exchange (NYSE) is the best known example of such a market structure (Table 1). Alternatively, dealers can compete for trades, buying and selling securities for their own account. Traditionally dealers competed in a central location, such as the London Stock Exchange or NASDAQ, but competition need not be centralised. Bonds, for example, trade primarily through bilateral negotiations between dealers and customers. A still third trading mechanism is the automatic matching of orders through an electronic broker. Today the majority of trading in the global foreign exchange market takes place over electronic exchanges such as Reuters and Electronic Broking System (EBS).

Table 1
Market structures

	Specialist	Dealer	Electronic
Equity	New York Stock Exchange	NASDAQ London Stock Exchange	Stock Exchange of Hong Kong Instinet Paris Bourse
Bond		Bond dealers	Tradenet EUREX
Foreign exchange		FX brokers	Reuters EBS

Actual markets do not conform to simple structures. Indeed, they typically involve more than one structure. What is important, therefore, is not the operation of any specific trading mechanism, but rather the rules by which trades occur. These rules dictate what can be traded, who can trade, when and how orders can be submitted, who may see or handle the order, and how orders are processed. The rules determine how market structures work, and thus how prices are formed.

Since rules can affect the behaviour of prices, liquidity might also naturally depend on how a market is structured. Indeed, liquidity concerns may dictate the structure of the market. Drawing on the extensive body of research investigating the interaction between market structure and liquidity, the remainder of this paper focuses on two critical issues in the creation of liquidity: the impact of limit orders, and the effects of transparency.

Limit orders

A wide variety of order types are found in securities markets. The most familiar type is a market order to buy or sell one round lot at the prevailing price. Other orders, such as “market-at-close”, “fill-or-kill” and “immediate-or-cancel” allow traders to control the timing, quantity or execution of their trades. By far the most common alternative type of order is a limit order specifying a price and a quantity at which a trade is to transact. Limit orders specify a price either above the current ask or below the current bid and await the movement of prices to become active. If the market is rising, the upward price movement triggers limit orders to sell; if the market is falling, the downward movement triggers limit orders to buy. Limit orders thus provide liquidity to the market.

Limit order traders receive a better price than they would have if they had submitted a market order, but face the risk of non-execution and a winner’s curse problem. Whereas a market order executes with certainty, limit orders await the movement of prices to become active, ie a limit order is held in a “book” until either a matching order is entered or the order is cancelled. Moreover, because once posted their prices do not respond to the arrival of new information, limit orders are more likely to be executed when they are mispriced. Foucault (1999) finds that in deciding whether to submit a market order or post a limit order, traders’ main consideration is the volatility of an asset. In a volatile market, the probability of mispricing an asset is higher, and so limit order traders quote relatively wide bid-ask spreads. This raises the cost of market order trading, thereby increasing the incentive to use limit orders rather than market orders. But as a result of fewer market orders, the execution risk associated with limit orders increases.

Order size may also influence investors’ choice between market and limit orders. Seppi (1997) concludes that small retail and large institutional investors prefer hybrid markets such as the NYSE, where specialists compete with limit orders to execute market orders.[4] Mid-size investors, on the other hand, might prefer pure limit order markets such as electronic exchanges. According to Seppi, specialists will undercut limit order prices at the margin. Such undercutting lowers the probability that limit orders will execute, thus resulting in reduced depth in the book. Evidence in Sofianos (1995) of a

[4] In hybrid markets, the ability of limit orders to compete with market makers depends on priority rules. Limit orders to sell at prices at or below the price at which the specialist proposes to sell, or limit orders to buy at or above the specialist’s bid price, typically have priority for execution.

U-shaped relationship between specialists' total revenue and trade size suggests that specialists do indeed provide relatively more liquidity to small and large trades.

The composition of order flows is a dynamic process, with investors' preferred order type changing in response to developments over time. Goldstein and Kavajecz (2000) examine the behaviour of liquidity providers on the NYSE during periods of extreme volatility. They find that following a precipitous drop in equity prices, traders abandoned limit orders in favour of floor brokers. In particular, whereas specialists maintained narrow spreads and normal depth, liquidity drained out of the limit order book. Similarly, in foreign exchange markets, trading tends to move from electronic order-matching markets to dealer markets during periods of market stress. Such dynamics raise the question of whether dealer markets handle information more efficiently than pure limit order markets.

Another issue relating to limit orders is whether they can provide enough liquidity for every type of trade. The experience of limit order markets suggests not. For example, on the NYSE, the Toronto Stock Exchange and other exchanges with features of limit order markets, a substantial proportion of block trades - trades of 10,000 shares or more - are submitted to block traders or "upstairs market makers", who form a syndicate of buyers to take the other side of the trade. One reason for using block traders rather than limit orders is that large transactions might be interpreted as signalling new information, and so move prices against the seller. Limit order systems are constantly evolving as new technologies are developed, and indeed OptiMark designed an electronic trading system that was supposed to minimise the impact that large orders had on price. OptiMark's system ensured that orders remained anonymous until executed in full and was initially lauded as presaging the transformation of institutional trading. Despite the system's advantages, however, it was poorly received by brokers and OptiMark ran into financial difficulties in mid-2000.

Finally, there is the question of how much information about the limit order book is optimal. On the NYSE and a number of other exchanges, orders held in the specialist's book are not common knowledge, although the specialist may choose to allow traders to view the book. By contrast, on electronic exchanges the order limit book is usually transparent. Madhavan and Panchapagesan (2000) find that on the NYSE the ability to observe the evolution of the book conveys valuable information to the specialist. In particular, specialists use information from the order book to set a more efficient opening price than the price that would prevail if all orders - both market and limit orders - were considered. Coppejans and Domowitz (1999) examine a pure limit order market and conclude that the trading process is influenced only by the flow of orders, not the stock of orders on the book. The book is not irrelevant; flows, after all, are changes in stocks. But in a market with an open book, the book per se does not appear to contain information on the value of the asset being traded. While helping us to understand how price formation occurs in actual markets, the results of these empirical studies do not imply that one particular market structure provides for more efficient price discovery than another. The experimental methods discussed below offer more meaningful insight into such hypothetical questions.

Transparency

As the information-based microstructure models demonstrated, the information available in the trading process can affect the trading strategies of market participants. It thus follows that the market equilibrium depends on the degree of transparency, ie the ability of market participants to observe the information in the trading process. Consider the previous discussion of the limit order book. If the book were known only to the market maker (as on the NYSE), then the market maker, as well as the informed and uninformed traders, would behave differently than if the book were common knowledge (as in the market examined by Coppejans and Domowitz).

The openness of the book is but one of many differences in the degree of transparency across markets. The breadth of trade data reported and even the timeliness of the reported data can also differ tremendously. Some markets such as bond dealers provide only pre-trade information, meaning that quote data are made available but not transactions data. Other markets require post-trade transparency, ensuring that the price and quantity of trades are observable. The NYSE and NASDAQ, for example, are required to report immediately all quotes and trades. At the other extreme, trades handled "off board" - trades executed outside of the United States after US markets close - need not even be acknowledged.

Differences in transparency may play a significant role in the creation of liquidity. As a factor in traders' strategic decisions, transparency can influence their willingness to participate in the trading process. In the United Kingdom, for example, the Financial Services Authority allows the reporting of large trades to be delayed for a period of time because it believes that immediate disclosure would expose market makers to undue risk as they unwound their positions and so discourage them from providing liquidity. Transparency is also a crucial consideration in the competition among markets for trading volume, and thus in the prospects for further fragmentation of liquidity.

Bloomfield and O'Hara (1999, 2000) use laboratory experiments to address some of these issues. Their experiments include multiple dealers operating under varying degrees of transparency and traders with differing trade motivations. A key finding is that low-transparency dealers are more likely to set the highest bid and the lowest ask (inside quotes) in early rounds of trading, in order to capture more order flow (Graph 3). The information learned from the order flow allows low-transparency dealers to quote narrower spreads than their more transparent competitors and to avoid money-losing trades. This informational advantage declines with repeated rounds of trading because low-transparency dealers reveal their information through their choices of quotes. Moreover, as trade progresses and individual dealers learn from trade outcomes, spreads for all dealers decline (Graph 3).

Trading gains follow a pattern similar to spreads. Wide spreads in early rounds result in large gains because traders in need of liquidity are forced to buy at high prices and to sell at low prices (Graph 4). Gains then decline in concert with the decline in spreads. Notably, neither high- nor low-transparency dealers earn money at outside quotes, ie bids that are lower than the highest bid and asks that are higher than the lowest ask. Even though trades at outside quotes are executed at more favourable prices, dealer profits are eliminated by the higher likelihood of transacting with an informed trader. At inside quotes, the proportion of total trades coming from informed traders is approximately 10%, but at outside quotes, the proportion rises to 70% (Graph 4). Liquidity traders' preference to transact at the best available quote results in this higher degree of adverse selection at outside quotes.

Interestingly, traders do not behave strategically in these experiments. Concern about the possible impact of a trade in one round on prices in future rounds might be expected to lead traders to pay a premium to conceal their trades by trading with low-transparency dealers. Recall that this concern was one of the motivations behind the design of the OptiMark trading system. As Graph 4 shows, however, informed traders are as equally likely to transact with low-transparency dealers as with high-transparency dealers. In this setting, transparency seems to have a greater impact on dealer behaviour than on trader behaviour.

Graph 3

Quote setting behaviour

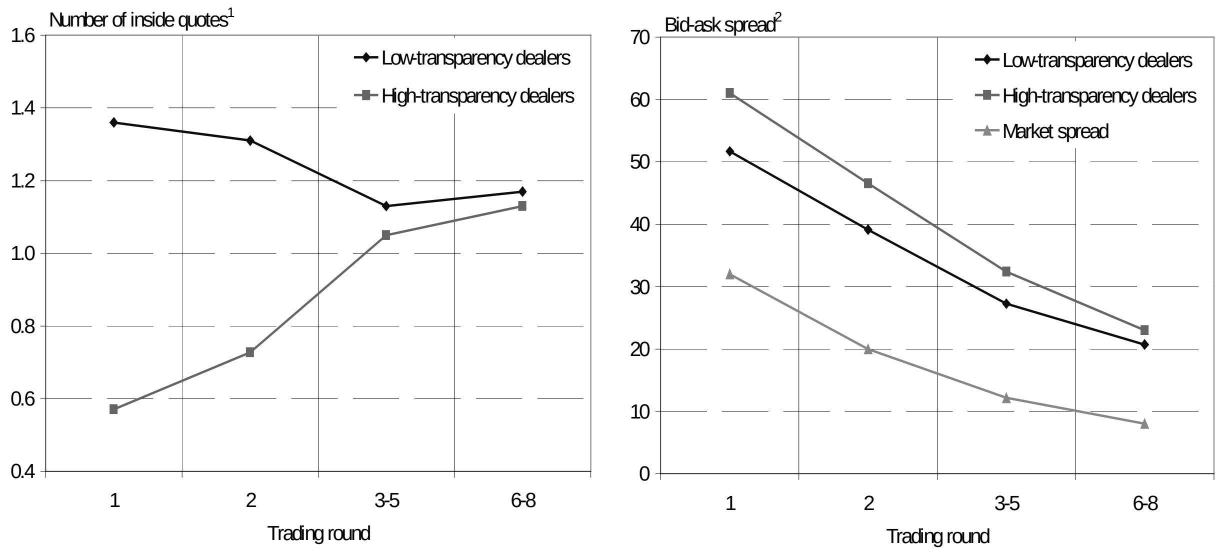

[1] Average number of inside quotes (the highest bid or the lowest ask) set by dealers. [2] Average bid-ask spread quoted for each security. The market spread is defined as the lowest ask minus the highest bid.

Source: Bloomfield and O'Hara (2000).

Graph 4

Dealer gains and adverse selection

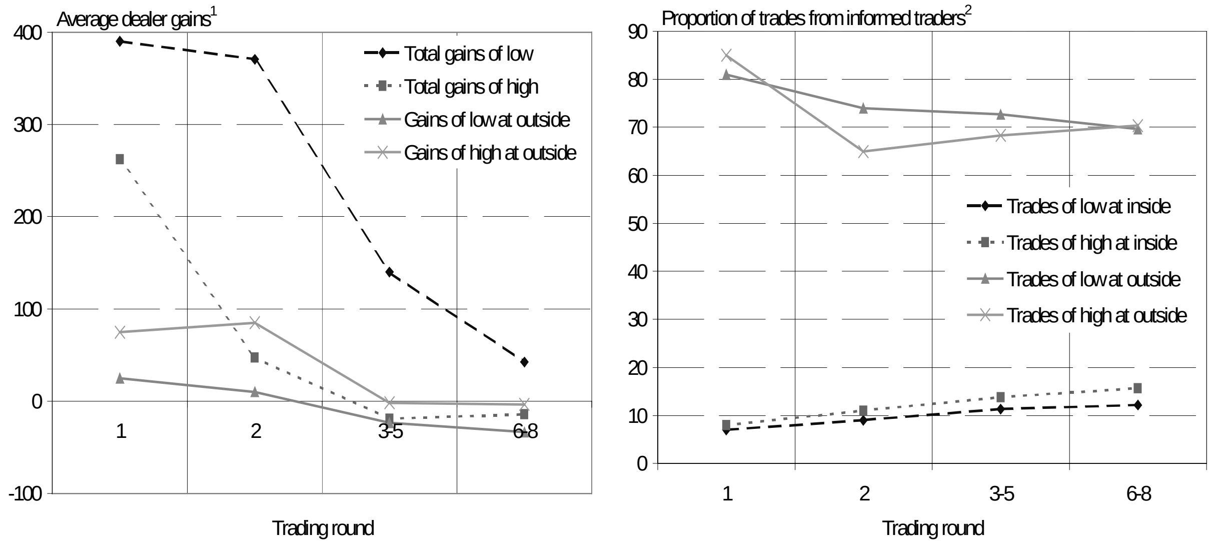

[1] Average gains of low- and high-transparency dealers over all trades (total gains) and for trades at outside quotes (gains at outside). [2] Proportion of total trades coming from informed traders. Total trades are segregated by dealer type (low or high transparency) and quote form (inside or outside quote).

Source: Bloomfield and O'Hara (2000).

In conclusion, economic experiments provide disquieting evidence that transparent markets may be less liquid than markets with weaker reporting requirements. Transparency reduces the information content of specific trades and so reduces dealers' incentive to compete for orders. As a result, bid-ask spreads in transparent markets tend to be wider than those in less transparent markets. This accords with the experience in actual markets. Spreads on Instinet, for example, are frequently narrower than those on NASDAQ. The "Golconda exchange" may be less transparent than some of the markets that currently dominate global trading.

References

Admati, Anat and Paul Pfleiderer (1988): "A Theory of Intraday Patterns: Volume and Price Variability", *The Review of Financial Studies*, 1, pp 3-40.

——— (1989): "Divide and Conquer: A Theory of Intraday and Day-of-the-Week Mean Effects", *The Review of Financial Studies*, 2, pp 189-224.

Bloomfield, Robert and Maureen O'Hara (1999): "Market Transparency: Who Wins and Who Loses?", *The Review of Financial Studies*, 12(1), pp 5-35.

——— (2000): "Can Transparent Markets Survive?", *Journal of Financial Economics*, 55, pp 425-59.

Blume, L, David Easley and Maureen O'Hara (1994): "Market Statistics and Technical Analysis: The Role of Volume", *Journal of Finance*, 49, 153-82.

Campbell, J, S Grossman and J Wang (1991): "Trading Volume and Serial Correlation in Stock Returns", working paper, Princeton University.

Cohen, Kalman, Steven Maier, Robert Schwartz and David Whitcomb (1981): "Transaction Costs, Order Placement Strategy and Existence of the Bid-Ask Spread", *Journal of Political Economy*, 89, 287-305.

Coppejans, Mark and Ian Domowitz (1999): "Screen Information, Trader Activity and Bid-Ask Spreads in a Limit Order Market", working paper, Pennsylvania State University.

Demsetz, Harold (1968): "The Cost of Transacting", *Quarterly Journal of Economics*, 82, pp 33-53.

Easley, David and Maureen O'Hara (1987): "Price, Trade Size and Information in Securities Markets", *Journal of Financial Economics*, 19, pp 69-90.

Fama, Eugene (1970): “Efficient Capital Markets: A Review of Theory and Empirical Work”, *Journal of Finance*, 25(2), pp 383-417.

Foster, F and S Viswanathan (1990): “A Theory of the Intraday Variations in Volume, Variance and Trading Costs in Securities Markets”, *The Review of Financial Studies*, 3, pp 593-624.

Foucault, Thierry (1999): “Order Flow Composition and Trading Costs in a Dynamic Limit Order Market”, *Journal of Financial Markets*, 2, pp 99-134.

Garman, Mark (1976): “Market Microstructure”, *Journal of Financial Economics*, 3, pp 257-75.

Glosten, Lawrence and Paul Milgrom (1985): “Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Agents”, *Journal of Financial Economics*, 13, pp 71-100.

Goldstein, Michael and Kenneth Kavajecz (2000): “Liquidity Provision During Circuit Breakers and Extreme Market Movements”, working paper 2000-02, New York Stock Exchange.

Harris, M and A Raviv (1993): “Differences of Opinion Make a Horse Race”, *The Review of Financial Studies*, 6(3), pp 473-506.

Ho, Thomas and Hans Stoll (1981): “Optimal Dealer Pricing Under Transactions and Return Uncertainty”, *Journal of Financial Economics*, 9, pp 47-73.

Kyle, Albert (1984): “Market Structure, Information, Futures Markets and Price Formation” in *International Agricultural Trade: Advanced Readings in Price Formation, Market Structure and Price Instability*, eds G Story, A Schmitz and A Sarris, Westview Press, Boulder, CO.

——— (1985): “Continuous Auctions and Insider Trading”, *Econometrica*, 53, pp 1315-36.

Lyons, Richard (forthcoming): *The Microstructure Approach to Exchange Rates*, MIT Press, Cambridge, MA.

Madhavan, Ananth (2000): “Market Microstructure: A Survey”, *Journal of Financial Markets,* 3, pp 205-58.

Madhavan, Ananth and Venkatesh Panchapagesan (2000): “Price Discovery in Auction Markets: A Look Inside the Black Box”, *The Review of Financial Studies*, 13(3), pp 627-58.

O’Hara, Maureen (1995): *Market Microstructure Theory*, Blackwell, Cambridge, MA.

Pfleiderer, Paul (1984) “The Volume of Trade and the Variability of Prices: A Framework for Analysis in Noisy Rational Expectations Equilibria”, working paper, Stanford University.

Seppi, Duane (1990): “Equilibrium Block Trading and Asymmetric Information”, *Journal of Finance*, 45, pp 73-94.

——— (1997): “Liquidity Provision With Limit Orders and a Strategic Specialist”, *The Review of Financial Studies*, 10(1), pp 103-50.

Sofianos, George (1995): “Specialist Gross Trading Revenues at the New York Stock Exchange”, working paper 95-01, New York Stock Exchange.

Spiegel, Matthew and Avanidhar Subrahmanyam (1992): “Informed Speculation and Hedging in a Noncompetitive Securities Market”, *The Review of Financial Studies*, 5, pp 307-30.

Stoll, Hans (1978): “The Supply of Dealer Services in Securities Markets”, *Journal of Finance*, 33, pp 1133-51.

Wang, J (1994): “A Model of Competitive Stock Trading Volume”, *Journal of Political Economy*, 102, pp 127-168.