Security Markets, Information and Liquidity

1. Introduction

How good are current securities markets? Often, the more liquid securities markets are put forth as near-realizations of the economist's ideal, the perfect competitive market. This is not wishful thinking. If we consider a hypothetical transaction in any commodity as consisting of the ''round trip'' purchase and resale, then we can accomplish such a pair of trades far more quickly and cheaply for a listed security than for virtually any other good in the economy. By this light, the stock markets should be held up as exemplary arrangements.

Public perception, however, is far different. The spectacular stock price drops of October 1987 and October 1989 have lead to scrutiny and criticism of the institutional arrangements by which the markets for securities are made and conducted. Much of this criticism is misplaced in the sense that no conceivable and feasible market structure can support a price that in the collective judgement is too high, even if the ''everyone'' agrees that the collective judgement is irrationally pessimistic. Nevertheless we must concede some validity to the criticism, in that trading activity and price determination are undeniably sensitive to institutional arrangements. The economic specialty known as market microstructure seeks to illuminate this connection and provide a scientific basis for preferring one form of market structure over another. This paper seeks to summarize and interpret, albeit sometimes subjectively, a selection of the current work in this area. The reader seeking more detail is referred to SCHWARTZ (1988).

Until recently, it was believed that market arrangements had little relevance for the long term properties of security returns or for investors who had horizons beyond the trading session. It is now widely recognized that transactions costs, even in relatively liquid security markets, are not negligible relative to the expected returns from holding the securities. If the spread between the bid and ask quoted prices on a $10 stock is $.25, then an investor who transacts via market orders (that is, one who sells at the bid and buys at the ask) incurs a 2.5% penalty on a round trip transaction. This is nontrivial relative to the typical annual return on a stock, and cannot be ignored even if the stock is held for longer periods.

AMIHUD/MENDELSON (1986) show that these costs are incorporated into the prices of securities. Taking the spread as a measure of the transaction cost, they find in a comprehensive analysis of New York Stock Exchange issues that shares with higher spreads also have higher expected returns (or equivalently, lower prices). In their concluding example, a hypothetical stock that moves from the highest spread group in their sample (average spread 3.2%) to the lowest spread group (average spread .486%) would realize an increase in value of about 50%. While not all firms could achieve such a dramatic revaluation, the analysis suggests that managers

intent on increasing shareholder wealth should pay attention to the market arrangements by which their stock is traded. Furthermore, from a social viewpoint, transactions costs are evidently a large component of the cost of new investment. Awareness of these facts underlies much current interest in microstructure.

A further development motivating present research into market organization is the expansion in the market process of computer and telecommunications technology. Although many market practices and customs from earlier times persist, few markets are now tied in any essential fashion to a particular geographic location. In many instances, automation has opened up the trading floor and allowed more widespread direct participation in the market. Furthermore, many market procedures which would tax the abilities of a human market-maker beyond the point of feasibility may now be considered quite practical if implemented with the appropriate technology. This technology has expanded greatly the set of alternatives, and it is not surprising that this has lead to a search for guiding principles of market design.

There can be, of course, no simple unanimous agreement about the best market organization. Individually and collectively, the demands we make on a market are many and conflicting. The statement that economics is a science of tradeoffs holds forcefully in the case of market structure. For example, if I wish to buy a security, then the best market structure is that which would present me with the lowest offer price. From the point of view of the seller, however, this may be the worst possible market structure. The example is trivial in that most of us are both buyers and sellers at different times, and so would not overly identify with either group. More subtle tradeoffs in market design abound, however, and will be considered shortly.

Depending on the nature of their market interactions, most observers hold certain beliefs about what constitutes a fair market. I provide here a brief list of possibilities:

- In a fair market, nobody should be at an informational disadvantage.
- In a fair market, the trade record should be public knowledge.
- In a fair market, a trader should be able to buy or sell without moving the price.
- In a fair market, large price swings should never arise.
- If there are more traders, prices will be better.
- A fair market should not require an active market-maker.

None of these are extreme or radical pronouncements. Each would be considered reasonable by at least some group of participants. Even in the current state of understanding, however, no statement can be said to hold without qualification. The remainder of this paper will attempt to communicate a deeper understanding of the implications of these statements.

2. The Elements of Market Theory

Most of this essay will deal with continuous markets. A market is continuous if it is open (that is, if trades can occur) at any point in the interval of time designated as the trading session. A continuous market is designed to process orders as they arrive. Although procedures vary across markets, the following scheme is fairly representative. Potential traders can observe bid and ask quotes: indicative prices at which the security is available for sale or purchase. A trader may obtain an immediate execution by entering an order to transact at the market quotes. Such an order is termed a market order. Alternatively, the trader may leave an order specifying a price and quantity for sale or purchase. This is a limit order and will only be executed if it becomes at some point the prevailing quote (that is, if it is the best bid or ask among all the other outstanding orders) and if a market order arrives to transact against it. Some markets do not allow limit orders to be entered by the public. The New York and American stock exchanges and the Swiss options and futures exchange permit public limit orders; the U.S. National Market System and the Swiss stock exchange for the most part do not.

The principal alternative to a continuous market is a call market, sometimes called a clearing house or a batched system. In a call market, buy and sell orders specify a quantity and a price. They are not executed as they come in, but are recorded and held. At prearranged times, the market-maker determines the price at which supply equals demand, and all orders which meet or better this market clearing price are executed at this price. Call markets are used to open continuous markets, and are sometimes used in markets which are normally continuous if the volume and imbalance of buy and sell orders is such that batching would result in a more orderly market.

Throughout this discussion, two concepts are of basic importance: liquidity and information. Liquidity refers to the immediacy with which a trade can be consummated. This simple definition stresses the timeliness aspect of liquidity, but there is in addition a price dimension as well. A demander of immediacy can almost always achieve rapid trade execution if he is willing to incur especially inferior terms of trade. In a liquid market, however, immediacy is achieved at minimal cost.

Information is data which someone believes is relevant to the valuation of the security, or which is material in the formulation of trading strategies. It is sometimes useful to contrast information that originates outside the market with information produced by the market. The former includes news releases and published financial statements. The latter comprises the trade history (price, quantity and identity of participants), quotes and orders. One problem in market structure is determination of exactly how much market-produced information should be disseminated and to whom.

Information may also be characterized by whether it is public or private. Public information is available to all market participants, while private information is known to some (but not all) agents. All private information is essentially advance knowledge of public information. Private information includes "inside" information whose use in trading is, in the U.S., generally considered illegitimate: a broker who buys for his own account based on advance knowledge of a customer's large buy order. But it may also result from legitimate activities; for example, a particularly astute security analyst who believes from his interpretation of the public record that a company is an attractive takeover target.

It is also useful at the outset to identify the active participants in the market. Liquidity demanders are traders motivated by idiosyncratic and individual needs which are unrelated to the fundamental value of the security. An indexed mutual fund, for example, buys shares in a company whenever it receives an inflow of investible funds, and sells whenever it is hit with redemptions. This buying and selling is indiscriminate in that it is not motivated by any company-specific information, public or private. In contrast, information traders are those who seek to profit by use of private information, i.e., by buying or selling in advance of fuller public awareness of the news. Liquidity suppliers accommodate the demands of other traders for immediacy in execution. They may be specially designated as market-makers (dealers or specialists) by the market authority, or they may simply be individual traders who by posting limit orders are trying to achieve a better price at the risk of no execution.

These categories should be considered as arising from loose attributes rather than sharp distinctions. In the first place, they are not stable over time: an individual who might buy on private information may sell due to a pressing liquidity need. Alternatively, a trader might buy via a liquidity-demanding market order, and sell via a liquidity-supplying limit order. Furthermore, even at a single point in time, the categories are not mutually exclusive. They nevertheless constitute useful functional groupings.

3. Market Structure and Informational Efficiency

We will consider a market to be informationally efficient if the price of a security fully reflects all public information. Alternatively, a market is infor-

mationally efficient if no trading strategy based on public information will yield consistent abnormal trading profits. The force leading to informational efficiency is competition among traders: any deviation from the efficient price represents a profit opportunity for an astute trader.

Informational efficiency ranks high in social importance. The value of an efficient and observable security price goes far beyond the relatively small number of individuals who are at the moment active participants in the market. A much larger number of individuals may hold the security passively, and use the price to figure their wealth and plan their consumption. Venture capitalists deciding where to channel investible funds use relative security prices as a guide to the opportunities with the highest returns. If the prices on which these agents rely are either not observable or do not reflect the best current information, then the agents will make mistakes in the level of consumption or savings and in the types of investments they undertake.

Given this strong social value of a price that is publicly available and informationally efficient, it would seem that as a matter of policy, the market authorities should take steps to ensure that the price incorporates as much information as possible, and that all prices generated by the market should be disseminated as widely as possible. There are, however, costs to such moves. Mechanisms that promote price stability and smoothness are generally at odds with informational efficiency. If newly-arrived information is sufficiently material, a substantial revaluation may ensue, and there is no good reason to block or impede this process. With certain types of information, on the other hand, revelation may lead to public benefit but private detriment.

Consider the seemingly innocuous question of whether or not transaction prices should be immediately communicated to the rest of the market. The process by which new information is incorporated into security prices is known as price discovery. As described in the next section, it is basically a competitive process whereby traders with an informational advantage transact in such a fashion as to drive the price to its efficient value. Knowledge of transaction prices, as they are determined, clearly facilitates the price discovery process. Traders can transact more confidently, having observed recent prices of similar transactions.

The actual parties to the trade may favor concealment of the terms, however. There are many instances of this. A buyer who is attempting to accumulate quietly a sizable position would prefer not to signal his activities with a report of the transaction. A dealer who has obliged a customer by acquiring a large block that he intends to resell would certainly elect to conceal the transaction. In fact, virtually every scrap of information about a transaction that public traders would find useful can be made available only at cost to the original transactors. There is thus an obvious divergence between public and private welfare.

By custom most markets report prices and quantities, but do not reveal trader identities (at least not immediately). There have been, however, instances of experimenting with delayed reporting of some transactions. (On the London Stock Exchange, current rules permit a delay of a day in the reporting of large trades.) The choice of what to report and when may not be a decision over which the market authority has complete discretion. The ability of a particular market center to enforce participation according to its rules varies considerably. If the incentives for concealing trade information are sufficiently high, alternative market centers may arise. This may lead to fragmentation of the market. In extreme instances, no reliable price may be publicly reported, and the security may simultaneously trade at different prices [1].

Fragmentation is always a significant danger. New communications technology is generally viewed as a force inimical to fragmentation, as it facilitates and reduces the cost of informational transfers. On the other hand, this technology also makes it easier for orders to be routed to alternative markets, which would increase the risk of fragmentation.

4. Public and Private Information

The issue of private information is one of the most active in current microstructure research. Granting for the moment that well-informed traders (holders of private information) exist, how do they affect the market process? It is clear that they will attempt to profit by buying on favorable inside information and selling on unfavorable information.

What may be less obvious is the response that this provokes on the part of other market participants. GLOSTEN/MILGROM (1985) consider a market consisting of the three types of agents described above. A market-maker posts bid and ask quotes: prices at which he is willing to buy and sell. Liquidity traders need to transact for reasons unrelated to the value of the security. Informed traders possess superior information. The first point is that the market-maker will lose in a transaction with an informed trader for the simple reason that the informed trader will transact only if the trade is to his advantage and the market-maker's loss. The market-maker may attempt to protect himself by setting a wide spread (a very low bid price and a very high ask price), but as long as the final decision to trade rests with the trader, a market-maker loss is certain.

COPELAND/GALAI (1983) take the point of view that the market-maker extends to the market at large a pair of options of indefinite maturity: an option to buy (a call) at the ask price and an option to sell (a put) at the bid price. Ignoring any premium received for writing these options, this kind of spread loses money if (at "maturity") either option is exercised. The market-maker receives no explicit premia, however, and so will lose whenever an option is exercised, that is, if and when an informed individual decides to trade.

If the informational asymmetries are sufficiently extreme, no one will voluntarily post quotes. For example, if the market-maker believes that all traders have information superior to his own, then he expects to lose any time an incoming trader hits his quotes. Faced with the impossibility of making a profit or breaking even, the market-maker will withdraw. This situation is one of market failure. Market failure does not imply that no one wishes to trade, only that trade will not be feasible for rational individuals under the prevailing market arrangement.

Clearly, a market composed solely of informational traders cannot remain open. If the market-maker is to be willing to post quotes, then there must remain a counter-vailing force to the costs imposed on the market-maker by the informed traders. This is generally assumed to come from the liquidity traders. If there are only liquidity traders in the market, the random arrival of buy orders at the ask price and sell orders at the bid price enables the market-maker to capture the spread between the bid and ask on every share that is turned over. This is a source of profit.

If the market-maker is confronting a trading population consisting of both informed and liquidity traders, then he would prefer to set the spread to service the liquidity traders and to discriminate against the informed individuals by refusing to trade. But if he cannot positively identify whether a trader is of the informed or liquidity type, then he may set one spread for all transactions. On average he will profit from the liquidity traders and lose to the informed traders.

The net effect here is, of course, a transfer of wealth from the uninformed to the informed traders, and so the aims of the two parties are naturally inimical. This conflict gets played out in trading strategy and market design in the following ways. Liquidity traders have a strong incentive to differentiate themselves as "know-nothings". Many mutual funds, portfolio insurers, index arbitrageurs and passive index investors fall in this category. They do not engage in active security analysis and believe that they could obtain better prices if their trading strategies could credibly communicate their lack of information. Such "separating" strategies might involve publicizing their identity or their willingness to trade. For example, Mark Rubinstein has proposed that program traders employ "sunshine trading": preannouncement of trading intentions. The possibility of such strategies is determined in large part by the market rules.

But informed traders too would prefer to be identified as liquidity traders, for the reason that no one will trade with them if they are correctly identified. To this end they disguise their activities by splitting up their trades over time, and in extreme cases by trading through dummy or nominee entities. Market regulation clearly imposes limits on these strategies as well.

As a practical matter then, market structure and rules establish limits on the separation of informed and liquidity traders. Ultimately, this makes separation a political issue, and in this connection the concept of "fairness" must be addressed. To return to one of the opening questions, does fairness require that everyone possesses the same information? Alternatively, should we encourage market structures that favor revelation of the amount of information underlying a trade?

Although at first glance such an outcome would be considered equitable by many, it is by no means obvious that the majority would find it desirable once the full implications were understood. The problem is that if informed traders are driven from the market, there will remain no incentives to produce information. The security price, therefore, will be much less informative. The high social value of a visible and informative price is enjoyed more by uninformed (liquidity) traders than by the informed. Collectively, therefore, the transfer of wealth from the liquidity to the informed traders is compensation for the information produced and enjoyed.

In an article provocatively titled "On the Impossibility of Informationally Efficient Markets", GROSSMAN/STIGLITZ (1980) explore these and other issues. The tradeoff between the cost borne by liquidity traders and the value of the information produced by informed traders is not amenable to easy quantification, however. Other aspects of the private information problem have been examined by KYLE (1985) (the trading patterns of a monopolistic insider), ADMATI/PFLEIDERER (1988, 1989) and FOSTER/VISWANATHAN (1987) (insider information and intraday trading activity).

5. Liquidity, or, why do trades move prices?

In the introductory remarks, liquidity was defined as the immediacy with which a trade could be consummated. The classic attributes of liquidity are considered to be depth, breadth and resiliency. A market possesses depth if there are many pending sell orders at prices immediately above the prevailing price and many pending buy orders at prices below. In addition, the market is broad if these orders are not only numerous but large. Finally, resiliency is the property that small price movements due to transitory supply and demand imbalances are quickly reversed by new incoming orders. The key qualification here is "transitory". Liquidity is sometimes spoken of as the affording the ability to buy and sell large amounts of a security without moving the price very much. As the following discussion demonstrates, however, the response of prices to trades may be closely connected to the extent of private information.

At first glance, the observation that buy orders tend to make prices move up, and sell orders tend to make prices move down is so commonplace as to not require further comment. But consider for a moment two alternative mechanisms.

1. A buy order causes prices to go up merely because the order takes all the available supply leaving only higher priced offers. It will take time for sellers to return to the market.
2. A buy order causes prices to go up because everyone in the market thinks that an informed trader is accumulating the stock.

The key feature of explanation 1 is transience. A more complete paraphrase is that price movements in response to trades are due to supply and demand imbalances that average out over time. Within the trading session, there are by chance momentary excesses of buy orders at some times and excesses of sell orders at other times. Under this explanation, price movements reverse and deviation from the true underlying price tend to die out. Explanation 2, on the other hand, entails a more permanent

revaluation of the stock. The market at large may not know the nature of the superior information, but if the trade is reported it constitutes a public signal. All traders realize that although the trade might have originated with an uninformed liquidity trader, it may also have been entered by an informed trader transacting at a profit. Is someone buying because they believe a merger announcement is imminent? Is someone selling because they have superior information on earnings announcements? In such cases, the trade itself is useful information and it is permanently incorporated into the stock price, to be reversed only by new information of opposing sign.

With the recent availability of new databases, empirical work has focused on determination of the permanent price impact. This is not as simple as pairing up a trade and the immediate price change: the price change may result from public information that is unrelated to the trade; the impact of the trade may not be immediate; and finally, trades have a tendency to persist in direction (buy orders tend to follow buy orders, and similarly for sell orders).

Nevertheless certain statistical procedures are robust to these problems. HASBROUCK (1990a) implements such an analysis for a sample of issues traded on the New York Stock Exchange and investigates the effect of trades on a representative price defined as the average of the bid and ask quotes. There are several key findings. First, the full impact of a trade on the quotes is not felt immediately, but with a protracted lag of up to five transactions. Second, the adjustment is unidirectional. Although studies of large (block) trades have documented a rebound or partial price reversal subsequent to the trade, the vast majority of all trades do not on average cause the quotes to overreact. HASBROUCK also notes that large trades cause the spread to widen momentarily, and that trades which occur when the spread is large have a larger impact on price than those which occur in the face of a narrow spread. Finally, preliminary work suggests an asymmetry between buy and sell orders: a buy order appears to move the price more (i.e., conveys more information) than a sell order of equal size. This is consistent, of course, with the observation that large (block) sales are more common, and therefore less informative, than block purchases.

In summary then, the persistent impact of a trade on the price of a security reflects the private information the market infers from the trade, which is in turn a consequence of the extent of the informational asymmetry in the market. Liquidity more properly refers to the absence of transient effects of the trade. In a liquid market, the price response to a trade is timely and complete, without a pronounced lagged adjustment or overreaction. This distinction is useful when we turn to the measurement of liquidity.

6. Supplying liquidity: call markets

The preceding discussion characterizes illiquidity as the existence of transient price movements due to the momentary impact of the order on the market. Clearly if traders arrive infrequently or if their orders are small (that is, if the market lacks depth and breadth), then the momentary impact of an arriving order on the security price is likely to be large. Alternatively, the deviations between actual transaction prices and the underlying fair value of the security are likely to be decreased as the number of active market participants increases. To increase liquidity, then, this suggests that the market be structured so as to increase the number of participants.

One way of doing this is to use a call market (defined above) in which all incoming orders are held and the market is cleared at regular intervals. This will result in fewer price determinations than in a continuous market, since the transactions (market clearings) will take place infrequently. In principle, a call should result in better prices, since each market clearing involves a relatively large volume of orders. Citing this and other advantages, COHEN/SCHWARTZ (1988) make a strong case for the superiority of an electronic call market. The available empirical evidence, however, refutes these

intuitions. The New York Stock Exchange employs a call at the opening, and continuous trading at all other times. AMIHUD/MENDELSON (1987) compare the volatility of open-to-open daily returns and close-to-close returns. Both sets of returns derive from the same underlying information flow, so in view of the supposed liquidity advantages of call markets we would expect the open-to-open returns to be less volatile. Surprisingly, AMIHUD/MENDELSON find the opposite, and this finding is confirmed by STOLL and WHALEY (1989). It is still unclear as to how the expectations of call market superiority can be reconciled with the empirical evidence. The empirical findings may reflect idiosyncrasies of the opening flow of liquidity orders, but a more definitive resolution awaits further research.

One aspect of call market implementation that may affect the quality of the resulting prices is whether or not the clearing procedure reports indicative prices and permits modification of orders. In what economists term a Walrasian auction, the auctioneer announces the market clearing price, then permits buyers and sellers to modify their bids, and computes a new market clearing price. This process is repeated until no further changes in price or orders occurs, at which time the market is cleared. The interim indicative prices convey and summarize information, of course, and this may lead to a more stable final price. In the Paris Bourse a la criée system, this process is conducted verbally among the floor traders. Under the NYSE's Opening Automated Reporting System (not in place during the time covered by the two empirical studies cited above), indicative quotes may be disseminated prior to the clearing. COHEN/SCHWARTZ note that in the present state of computer and telecommunications technology there are no practical barriers to implementing this process even with geographically dispersed traders.

7. Supplying liquidity: dealers

Liquidity is supplied to a continuous market by agents who are willing to expose bid or ask quotes to other potential traders. Quotes are indicative prices. In the absence of actual transaction prices, quotes can perform the same socially valuable role of information summary. While quotes are a distinct public good, however, their exposure constitutes a leak of information and reveals something about the agent's propensity to trade. An agent who announces that he is willing to trade a stated amount at a stated price is placing himself at a competitive disadvantage. He is extending an option to the market in general and to informed traders in particular. What are the compensations for doing this, and how are these compensations affected by market structure?

In most markets, quotes may be posted by dealers or by limit-order traders. Although there is a sense in which anyone who posts a limit order is operating as a dealer, the term is usually reserved for traders who generally maintain a continuous presence in the market. In this section, we consider dealers; limit-order traders are considered in the next.

In designing a market it might be hoped that dealers would voluntarily enter as passive or neutral providers of liquidity. The ideal dealer would always stand ready to buy the security at some amount c below its fair value and selling it at c above its fair value, where the margin c would enable the dealer to recoup on average the costs of maintaining a market presence. It has long been recognized, however, that despite the apparent reasonableness of this requirement, it exposes the dealer to substantial inventory risk. GARMAN (1976) pointed out that random arrival of buyers and sellers would lead to large dealer positions in the security, even over short periods of time. AMIHUD/MENDELSON (1980) and O'HARA and OLDFIELD (1986) have analyzed this problem, and find that to restore the desired inventory position, the dealer must at certain times adjust the quotes in such a way as to elicit a balance of buy and sell orders. Of necessity then, a market-maker cannot always be a passive provider of liquidity, but must at times actively influence the price.

To make matters worse from the prospective dealer's viewpoint, many institutions impose affirma-

tive obligations on the market-maker. They are often required to make a market at virtually all times, to keep spreads narrow, and maintain price continuity. In addition, they may be restricted in the nature of trades they can make for their own account. All of these requirements directly benefit the trading public and impose very real costs on a dealer.

The offsetting benefits to being a dealer range from the obvious to the obscure. It is clear that the dealer captures the bid-ask spread from the liquidity traders, and also receives (in many markets) a small commission on all trades. The market may permit the dealer access to information not publicly available, such as knowledge of pending orders. As pointed out above, any private information is valuable. The exchange may also accord the dealer discretionary authority in regulating trading activity, resolving disputes and allocating quantities among traders. Traders may respond to this authority by favoring the dealer in trades.

When available, reports of dealer profits provide a reasonable guide to the overall competitive position of the dealer vis a vis other traders. It is very difficult to determine the private cost or value to the dealer, however, of the separate obligations and privileges associated with market making, let alone their social cost or value. For a dealer to remain in business, his costs must be covered, but it is difficult to discern who is ultimately bearing these costs and in what fashion. These questions are politically sensitive. Whatever the merits of a market system that places dealers in a privileged position may be, those who must transact but cannot closely observe the execution process may perceive themselves to be at a disadvantage.

8. Supplying liquidity: limit-order traders

Public limit-order traders are those who opt for limit orders in preference to market orders. The factors bearing on the choice between limit and market order may be illustrated under the simplifying assumption that the bid and ask quotes are placed symmetrically about the "fair" value of the security. An investor buying with a market order pays the ask price, incurring a cost of one-half the bid-ask spread. An investor who places a limit order to buy at the current bid price receives the bid price (a discount from fair value of one-half the spread) if the order is executed. Execution is not, of course, certain. If the price moves up, there will be no execution, and if the trader still wishes to buy, he may have to enter a subsequent order at terms worse than those a market order would have originally achieved. COHEN et al. (1981) model the tradeoff between price and execution certainty and find that a public limit order trader is compensated by avoiding the bid-ask spread. In an active market, the liquidity provided by such traders or the "trading crowd" may be substantial.

Increasing the flow of limit orders is a particular priority in markets which either do not formally designate dealers or else afford them few special privileges. In such markets, it is expected that public traders will supply the bulk of the liquidity, and that there must be special incentives to ensure sufficient volume. This is particularly true of automated systems.

To understand the nature of these incentives, consider a trader who intends to sell a large amount of a security. Submission of the entire amount as a limit order would reveal to the market the existence of a large overhanging block. In principle the trader could submit a sequence of smaller limit orders, waiting for one to execute before submitting the next. This would not protect, however, the trader's priority in the queue: each new order would ordinarily be sequenced after other orders which had arrived in the interim.

The solution to this dilemma in the Globex trading system planned for after-hours trading by the Chicago Mercantile Exchange is to allow the trader to designate primary and secondary quantities. Secondary quantities are not revealed to the market, but maintain some time-priority protection. (A summary of the features of several automated trading systems is given by DOMOWITZ, 1989). Similarly, in the electronic call market proposed by CO-

HEN/SCHWARTZ, traders can submit disclosed or undisclosed orders. Undisclosed orders are reflected in the clearing process, but lose priority to disclosed orders if quantity rationing is necessary at the clearing price. As a third example, the interdealer market in U.S. Treasury debt grants the limit order trader the right of first refusal for all additional trades at his initial posted price. All of these designs encourage traders to reveal prices, while not necessarily revealing the full extent of their demands.

A further cost of exposing a limit order to the market is the cost of monitoring the market and modifying the order to accommodate new developments. These developments may include executions of pending orders for other securities previously submitted by the trader or price changes in securities the trader believes to be "bellwether" or leading indicators. Allowing complex contingent orders reduces the monitoring cost by relieving the trader of labor involved in physically watching the market. Such orders are feasible in computer-based systems, and essentially allow a computer to mimic the activities of a human floor broker acting in the trader's behalf.

Finally, the timing and placement of limit orders is clearly affected by price discreteness: the smallest unit in which a price can be quoted, i.e., the minimum "tick". HARRIS (1989) has recently analyzed price clustering behavior on the New York Stock Exchange. This study finds statistical evidence for natural levels of discreteness, and also provides a methodology for determining when the minimum tick is too confining.

9. Measuring liquidity

Microstructure analysis has not yet arrived at the point where we can assert a priori that a particular market structure will afford the most liquidity for a given security. In the current state of the art, market structure is best determined by judicious experimentation among the institutionally practical alternatives, and this is likely to remain true in the foreseeable future. While polls of market participants can provide valuable feedback about what does and does not enhance the quality of the market, it is also useful, to possess objective and quantitative measures of liquidity. In this area, we seek statistics that will measure the frequency and magnitude of transient discrepancies between "fair security values" and actual transaction prices. This section describes the state of the art.

9.1 The bid-ask spread

The spread between the bid and ask quotes is often taken as a measure of market liquidity. It has the advantage of being relatively easy to obtain and interpret. The rationale for using the spread is as follows. The spread is viewed as arising solely from the need of dealers to cover the direct costs of transactions. As such, they set bid and ask prices symmetrically about the fair value of the security. A market-order buyer then pays a price that is higher than the fair value by half the spread, and a market-order seller receives a price lower by half the spread. These deviations from the average of the bid and ask prices constitute the transient trade-induced deviations.

Although appealing in its simplicity, as a measure of liquidity, the bid-ask spread suffers from many disadvantages. Posted quotes are typically valid only for relatively small amounts of the security. The spread can therefore measure liquidity only for those who trade small amounts of the security at the posted quotes. Large trades that transact outside the posted quotes and negotiated trades that transact inside the posted quotes, both of which are frequent and important occurrences, are not reflected.

9.2 Volume-based liquidity ratios

One of the shortcomings attributed to the bid-ask spread was that it was relevant only to the relatively small trades that would take place at the posted quotes, i.e., it was insensitive to trade size. Various

liquidity ratios attempt to remedy this defect by relating price changes to trading volume. One might, for example, consider the ratio of the size of a day's price change to the day's trading volume. Since this will vary considerably from day to day, a more stable measure is obtained by taking the average of this ratio over a number of trading days.

While the liquidity ratio is not without intuitive appeal, it is a misleading measure in that it is contaminated by effects that have nothing to do with the usual notions of liquidity. The main problem is that the price change used in the numerator is in large part the result of non-trade-related public information. Suppose a company makes a significant news announcement. There will be a large price change that may or may not be accompanied by large trading volume. In addition, liquidity ratios share a common problem with the bid-ask spread in that they do not distinguish between trade impacts that are transient (and indicators of illiquidity) and those that are permanent. These and related points are discussed by GROSSMAN/MILLER (1988).

9.3 Random-walk based measures of liquidity

In this section, we consider measures of liquidity that are based on the principle that in the absence of market imperfections security prices should follow a random walk. Since the random-walk model of security prices is occasionally criticized as an extremism of academic artifice, it should be emphasized that these measures use it only as a benchmark, not as a description of how actual security prices behave.

The salient characteristic of a random walk is the absence of predictable patterns. The relevance of this property for assessing liquidity is that transient price changes induced by trades lead to predictable patterns. In an illiquid market, for example, part of the price increase following a large buy order might on average be reversed over time. On average, then, price increases are followed by price decreases (of smaller magnitude), and this constitutes a useful prediction rule. Most market imperfections that are believed to cause illiquidities can be shown to engender similar sorts of patterns. From a statistical viewpoint, there are many ways to measure "how close" a security price path resembles a random walk. The problem is obtaining a single measure that captures diverse effects over various time frames.

A random walk possesses the property that the volatility (more properly, the variance) grows in proportion to the time interval. That is, the variance of security returns computed over a two-hour horizon should be twice the variance of the security return computed over a one-hour horizon. In practice, however, market illiquidities that induce transient short-run volatility may lead to a one-hour variance that is disproportionately high relative to the two-hour horizon. HASBROUCK/SCHWARTZ (1988) exploit this effect in constructing ratios of long-run return variance to short-run return variance ("market efficiency coefficients") to assess relative short- and long-run volatility.

Variance ratios are easy to compute from observed transaction prices. Since they are based on actual transaction prices, variance ratios reflect trades of all sizes, as well as those occurring inside or outside of the bid-ask spread. Furthermore, since public information is reflected in both the long-run return variance (the numerator of the ratio) and the short-run return variance (the denominator), it effectively cancels. The main disadvantage of variance ratios is that they are sensitive in some circumstances to the horizons over which the variances are computed. If the illiquidities induce opposing return patterns over different horizons, a single variance ratio may incorrectly imply high liquidity.

Recent work in macroeconomics has lead to statistical techniques for isolating the random-walk component of a time series. These techniques hold great promise in the development of liquidity measures that assess quite closely how far transaction prices deviate from the underlying fair (random-walk) values. HASBROUCK (1990b) describes a preliminary computation and refinement of this work is continuing.

10. Summary

This essay has attempted to clarify some of the tradeoffs inherent in market design by bringing in ideas from current academic research. Here, I summarize these tradeoffs, and append some concluding thoughts.

The social value of a published transaction record vs. the private value of confidentiality
Observable transaction prices appear to be of great importance in the price discovery process, and any attempt to interfere with their dissemination impairs the speed with which information is incorporated into prices. The benefit that may accrue to the parties of a confidential transaction appears small by comparison.

Private information: the loss to uninformed traders vs. the social value of the information produced and disseminated
Market realities do not support either extreme here. The social value of information produced through superior analytical abilities is in many instances substantial enough to justify a private return. Private information illegitimately appropriated in the context of a fiduciary relationship carries no such value. On the other hand, a misplaced principle of egalitarianism is sometimes used to justify the notion that all traders should be equal in all respects. This illusion should not be encouraged. Traders who neither perform security analysis nor delegate their activities to those who do must accept the reality that they will lose on average to those who have at considerable expense and effort produced useful information. The principle of fair and equal access to markets (a desirable goal) should not be confused with the principle of equality of outcomes in trading decisions.

The rights and obligations of dealers
Consistency requires that if the dealers in the market are to be encumbered by regulations regarding (for example) continued market presence or price stabilization, then they must also be accorded privileges in access to information or the market. While there may be common agreement that continued market presence and price stabilization are good things, it must also be acknowledged that there will be a cost imposed on the other traders.

By how much and in what fashion should suppliers of liquidity be compensated?
Whenever possible, market rules should favor public suppliers of liquidity and minimize the barriers to providing a dealer function. Markets that impose noneconomic barriers to suppliers of liquidity run the risk that this liquidity will be supplied in an entirely different venue: a new market that is competing with and fragmented from the original.

As to the fashion in which these tradeoffs are addressed, and the manner in which the inevitable compromises are reached, this essay concludes with recommendations for experimentation, flexibility and pragmatism. As we move toward an era in which market rules and procedures are implemented as software in an electronic system, it is well to remember that software can be easily modified. Academic research has not lead to definitive pronouncements on market structure, nor is it likely to in the near future. Judicious experimentation remains the best way for arriving at the ideal of a fair and orderly market.

Footnotes

[1] Aside from informational inefficiencies, this may also increase transactions costs.

References

ADMATI, A.R. and P. PFLEIDERER (1988): "A Theory of Intraday Patterns: Volume and Price Variability", Review of Financial Studies, 1, pp. 3-40.

ADMATI, A.R. and P. PFLEIDERER (1989): "Divide and Conquer: A Theory of Intraday and Day-of-the-Week Mean Effects", Review of Financial Studies, 2, pp. 189-223.

AMIHUD, Y. and H. MENDELSON (1980): "Dealership Markets: Market-Making with Inventory", Journal of Financial Economics, March.

AMIHUD, Y. and H. MENDELSON (1986): "Asset Pricing and the Bid-Ask Spread", Journal of Financial Economics, 17, pp. 223-249.

AMIHUD, Y. and H. MENDELSON (1987): "Trading Mechanisms and Stock Returns: An Empirical Investigation", Journal of Finance, 42, pp. 533-553.

COHEN, K.J., S. MAIER, R. SCHWARTZ and D. WHITCOMB (1981): "Transaction Costs, Order Placement Strategy and the Existence of the Bid-Ask Spread", Journal of Political Economy, April.

COHEN, K.J. and R.A. SCHWARTZ (1988): "Realizing the Potential of an Electronic Trading System", in Information Technology and Security Markets Under Stress, New York University.

COPELAND, T. and D. GALAI (1983): "Information Effects on the Bid-Ask Spread", Journal of Finance, 35, pp. 235-48.

DOMOWITZ, I. (1989): "The Mechanics of Automated Trade Execution Systems", working paper, Department of Economics, Northwestern University.

FOSTER, D.F. and S. VISWANATHAN (1987): "Intraday Variations in Volumes, Spreads and Variances I: Theory", working paper, Fuqua School of Business, Duke University.

GARMAN, M. (1976): "Market Microstructure", Journal of Financial Economics, June.

GLOSTEN, L.R. and P. MILGROM (1985): "Bid, Ask and Transaction Prices in a Market-Maker Market with Heterogeneously Informed Traders", Journal of Financial Economics, 14, pp. 71-100.

GROSSMAN, S.J. and J.E. STIGLITZ (1980): "On the Impossibility of Informationally Efficient Markets", American Economic Review, June.

GROSSMAN, S.J. and M.H. MILLER (1988): "Liquidity and Market Structure", Journal of Finance, 43, pp. 671-637.

HARRIS, L. (1989): "Stock Price Clustering, Discreteness and Bid/Ask Spreads", New York Stock Exchange Working Paper.

HASBROUCK, J. (1990a): "Measuring the Information Content of Stock Trades", forthcoming, Journal of Finance.

HASBROUCK, J. (1990b): "Assessing the Quality of a Security Market", working paper, Stern School, New York University.

HASBROUCK, J. and R.A. SCHWARTZ (1988): "An Assessment of Stock Exchange and Over-the-Counter Markets", Journal of Portfolio Management, Spring.

KYLE, A.S. (1985): "Continuous Auctions and Insider Trading", Econometrica, 53, pp. 1315-1336.

O'HARA, M. and G.S. OLDFIELD (1986): "The Microeconomics of Market-Making," Journal of Financial and Quantitative Analysis, 21, pp. 361-376.

SCHWARTZ, R.A. (1988): "Equity Markets: Structure, Trading and Performance", Harper and Row, New York.

STOLL, H.R. and R.E. WHALEY (1989): "Stock Market Structure and Volatility", working paper, Owen Graduate School of Management, Vanderbilt University.