Liquidity Without Volume: The Case of FINEX, Dublin

DANA R. CLYMAN
CHRISTOPHER S. ALLEN
RICHARD JAYCOBS

INTRODUCTION

Speculators, hedgers, and professional portfolio managers alike require that exchange-traded futures markets have substantial volume before they will participate. Indeed, commodity trading advisors (CTAs) often insist on volumes between 2000 to 5000 contracts per day before they will even consider trading a futures market.[1] Volume, however, is not their true concern. What these market participants actually want are liquidity assurances; that is, assurances that their orders, particularly orders to exit positions, will be executed on demand without any adverse price movement resulting from the orders themselves. Liquidity, however, is exceedingly difficult to measure directly, particularly because it is hard to distinguish between normal price movement and price movement in response to large orders. As a result, a variety of proxies are used.

The most frequently used proxy is volume, although open interest and the average size of the bid–ask spread are also monitored. Indeed, the average bid–ask spread is monitored not only because it affects the

[1]The authors thank Susan Sjo, President, Sjo, Inc., for describing common perceptions and attitudes of the CTA community, even when she disagreed with them. Conversation, May 24, 1995.

■ *Dana R. Clyman is an Assistant Professor at The Darden Graduate School of Business Administration, University of Virginia.*

■ *Christopher S. Allen is a Financial Analyst at Hewlett Packard.*

■ *Richard Jaycobs is a Director of the Computer Trading Corporation.*

The Journal of Futures Markets, Vol. 17, No. 3, 247–277 (1997)

CCC 0270-7314/97/030247-31

average price paid to execute large orders but also because it represents the average cost of a round-trip transaction for a standard-sized order, another frequently cited component of liquidity.[2] Volume, however, is often used as a proxy for this component of liquidity as well, because of the well-documented fact that low-volume markets tend to have wider bid–ask spreads than high-volume markets. George and Longstaff (1993), for example, developed a taxonomy of spreads for the S&P 100 Index options and confirmed that volume and spread were negatively related in that market. They then argued that the relationship made sense because higher trading activity both determines and is caused by tighter spreads. Not only do tight spreads attract more business, but increases in volume often reflect additional market participants, which, in turn, increases competition and further tightens spreads.

But this requirement by market participants for high volume in exchange-traded markets contrasts sharply with participants' requirements in direct-dealing markets, such as the over-the-counter markets, swaps markets, and various exotic derivatives markets. In many of these markets, volume is not only low, but in many cases cannot even be measured (because there are no formal regulatory requirements to report the statistics).

Why then are market participants willing to trade in low-volume (or even potentially unknown-volume) direct-dealing markets but not in low-volume exchange-traded markets? One possible reason is that in direct-dealing markets participants have direct access to the market makers. Each trade is executed with a known counterparty, who acts as a principal to the transaction. Because participants tend to believe that market makers provide implicit, if not explicit, assurances of liquidity on demand, they are willing to trade with them even when there is little definable volume in the security traded. Furthermore, in direct-dealing markets the market makers are usually large corporations or banks with large pools of capital and audited financial statements. Thus, participants can assess, ex ante, whether they believe their counterpart's capital pool is sufficient to sustain market-making activities.

Exchange-traded futures markets, on the other hand, have relied traditionally on a community of small, individual participants, called *locals*, to provide liquidity. As a result, there is no ex ante method to determine the size of the capital pool nor to assess the willingness of the locals, individually or collectively, to commit their capital to market-making activities in specific situations. Worse, exchange-traded markets provide

[2]The average bid–ask spread is often called the transaction cost of the market because, in theory, customers pay market makers half the spread, on average, when entering a position and half again when exiting. Indeed, many traders' businesses are built on "earning the spread."

only indirect access to these market makers. Not only is there no mechanism for trading with particular counterparties, but also it is illegal to obtain liquidity assurances from individual locals; after all, assurances of on-demand liquidity smack of prearranged trading arrangements, and the prohibition against prearranged trading is one of the cornerstones upon which the public-price-discovery mechanism of open-outcry markets is built. Thus, when trading in exchange-traded markets, participants have no choice but to turn to surrogate measures for liquidity. They cannot measure liquidity directly, they cannot assess the sufficiency of the capital pool of the participating locals, and they cannot obtain assurances of on-demand liquidity from market makers.

This of course leads to the well-known chicken-and-egg problem for launching new futures products. If volume is necessary for liquidity, and liquidity is necessary to attract participants and build volume, how do you begin? It may be exactly for this reason that there has been an explosion of successful new over-the-counter products like swaps, complex options, and exotic derivatives during a period when the new-product-introduction record of the futures industry has been poor at best. In the interbank dealing market, for example, even minor currencies are considered sufficiently liquid to trade (and do trade successfully) even though they have little or no identifiable volume. At the Chicago Mercantile Exchange, on the other hand, while approximately 23 new contracts were launched during the period, 1984–1994, over half of them had no volume in 1994. Similarly, at the Chicago Board of Trade (excluding the MidAmerica Commodity Exchange, which it purchased during the period) approximately 35 new contracts were launched, yet more than two thirds of them no longer trade.[3]

What's more, this new-product-introduction record occurred even though the exchanges have long understood that volume is a prerequisite for success because the volume-inducing perception of liquidity is necessary before most potential market participants will trade. Indeed, over the years, exchanges have made enormous efforts to build volume for new products. It is a matter of public record that exchanges have tried to generate initial volume in new markets in a host of ways including subsidizing participation through fee reductions, cash payments, low-cost trading permits or even free temporary permits, and other similar mechanisms. It is also a matter of industry folklore that some exchanges have, on occasion, pressured members (locals) to trade new contracts amongst

[3]These numbers are based on Futures Industry Association reports. They are approximate because they count major modifications of existing contracts as revisions rather than as new contracts. If counted the other way, the numbers of both new contracts and failed contracts would increase.

themselves until sufficient volume is built to attract outside participants. And to this day, one often hears the opinion expressed that the only way to launch a new exchange-traded product successfully is by enticing the locals to take the long view and endure short-term losses to build the volume needed for long-term success.[4]

Thus, the traditional focus of most liquidity-building activities at futures exchanges has been directed at locals, the traditional market makers and source of capital in futures markets. Nevertheless, there is another source of liquidity for some futures markets: the already existing liquidity of other, deeper markets that can be tapped through arbitrage. After all, whenever arbitrageurs stand ready to trade a market that is derivative to a deeper underlying market, and there are enough of them competing for the business to prevent any of them from charging excess rents, then the derivative market should be liquid regardless of volume. In fact, this is the fundamental observation that drives the theoretical pricing of all derivative securities, whether it be the Black–Scholes pricing formula for pricing options, or complex simulation models for pricing exotic derivatives. In essence, as long as it is possible to hedge a derivative position by taking offsetting positions in other markets, and as long as there is sufficient competition among the arbitrageurs, then price movements in the derivative market must be related to price movements in the underlying markets.[5] Furthermore, the prices in the two markets remain linked to one extent or another whether the arbitrage is perfect or the result of the risk-arbitrage activities of large market participants or market-makers. Thus, even when the only link is the risk-arbitrage activity of large market participants managing matched books, the liquidity of the underlying markets can be transferred to the derivative market, thereby causing it to be liquid even when its volume is low.

This idea of tapping the liquidity of other deeper underlying markets to help build the liquidity of a derivative market is not new. In 1972, when the International Monetary Market (IMM) was launching its currency-futures contracts, it initiated a new class of clearing membership:

[4]The authors would like to thank an anonymous reviewer for pointing that it is rare for a low-volume futures market to be that such a market is considered successful. After all, Exchange revenues are derived principally from transaction-based fees and from the sale of ticker information, both of which tend to be volume dependent. Moreover, it is the volume in the pit that provides Exchange members, like locals, brokers, and other market participants, the opportunity to earn a profit. Nevertheless, as the reviewer also pointed out, markets that provide participants economic access to needed exchange services like the clearinghouse, credit, government oversight, confidentiality, and mark-to-market, may still be considered successful even when volume is low.

[5]For a broader treatment of this relationship between the prices of underlying and derivative markets, see Figlewski (1987).

Class B Clearing Members. According to the rules of the Chicago Mercantile Exchange (CME),

> Class B Clearing Members shall be non-bank member firms affiliated with a Class A [normal] Clearing Member solely for the purpose of proprietary arbitrage with a single IMM approved bank.[6]

Class B clearing members were arbitrageurs, required by the regulations of the CME to exactly offset every trade entered into on the floor of the Exchange with an opposite trade with the member's approved bank. Thus, through the arbitrage activity of the Class B members, the IMM was tapping the liquidity of the interbank currency markets.[7,8]

Over time, the importance of the Class B memberships declined[9] and the IMM marketplace flourished. Thus, by tapping this alternative source of liquidity, the IMM found a way to resolve the chicken-and-egg problem; the liquidity was provided, the volume followed, and by all measures the contracts became successful.[10]

As did the IMM, the FINEX, the financial instruments division of the New York Cotton Exchange, is currently attempting to tap the interbank currency markets as a source of liquidity for its various currency-derivative contracts. This article focuses on the Dublin-based trading of FINEX's U.S. Dollar Index (USDX) contract—a trade-weighted geometric basket of currencies from ten countries (Germany, Japan, France, the United Kingdom, Canada, Italy, the Netherlands, Belgium, Sweden, and Switzerland).[11] This index was first published by the Federal Reserve Board in the 1970s. Because the USDX contract consists of a trade-

[6]Rules of the International Monetary Market of the Chicago Mercantile Exchange, Inc., January 1973, p. 1.

[7]Reportedly, these relations were so successful that the first arbitrageur [Class B clearing member] easily made over a million dollars in the first 30 days. This, of course, led to other members establishing themselves as Class B clearing members, bringing ever more liquidity into the market. Over time, as the banks established their own floor operations, the Class B members as a group lost their competitive edge.

[8]The authors would like to express their grateful thanks to Mark Powers, Editor, *The Journal of Futures Markets,* for pointing out the existence of the Class B memberships on the IMM—memberships which he had a hand in creating—and their historical significance for this study.

[9]Although the Class B membership rules remain on the books of the Exchange, according to Scott Brusso, Director, Currency and Interest Rate Products, only a very small handful of people still trade under these rules. Furthermore, the trading that does occur appears to be at the discretion of banks that either wish to avoid carrying futures positions on the books or showing brokerage charges.

[10]The authors would like to thank an anonymous reviewer for pointing out that there is similar evidence in stock markets supporting the idea that liquidity in one market can help in another. For an excellent and accessible treatment of liquidity relationships between markets, see O'Hara, Chap. 8.

[11]The USDX contract trades in both New York and Dublin at different times during the day. Trading on both trading floors, however, is regulated by the CFTC, and treated as one market. In other words, a contract entered into on one floor can be exited on the other.

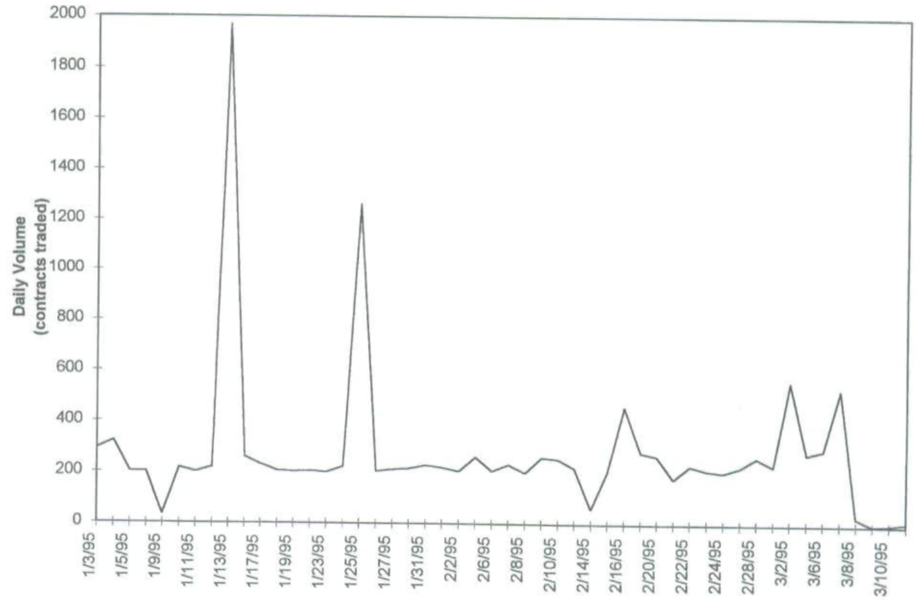

FIGURE 1
March 1995 USDX contract: Dublin volume.

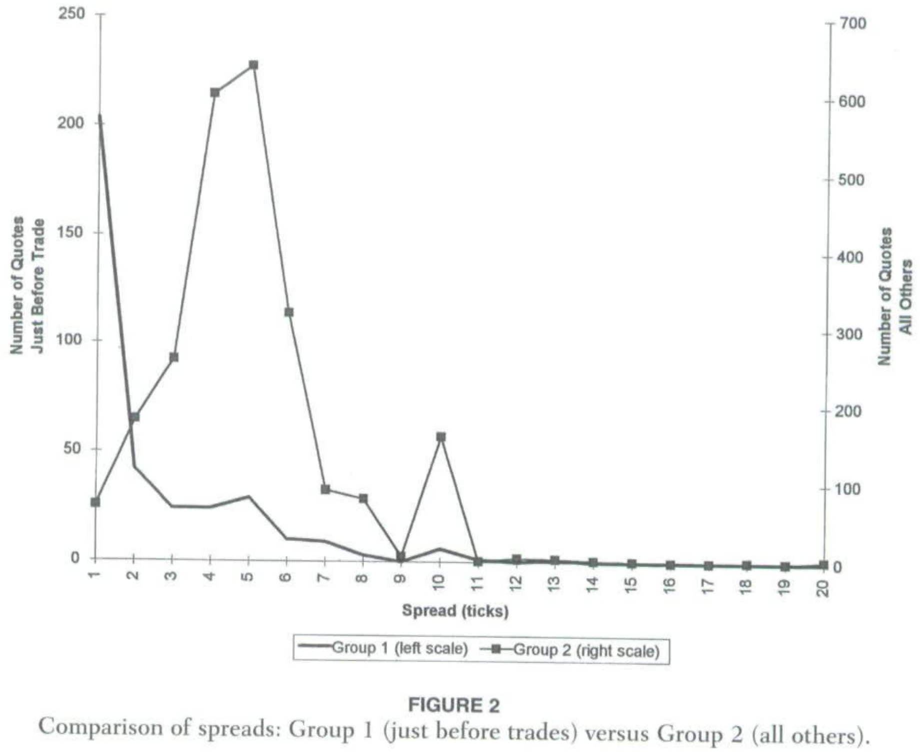

FIGURE 2
Comparison of spreads: Group 1 (just before trades) versus Group 2 (all others).

weighted, geometric basket of currencies—some large, some small—it too is derivative to the interbank currency markets. To exploit this connection, FINEX's liquidity-building efforts are focused not on the locals, but rather on facilitating arbitrage.[12]

One way that FINEX attempts to facilitate arbitrage-enhanced liquidity is through the FINEX Large Lot System (FILLS). Adopted in June 1993, this trading rule allows market participants to request bid-and-offer quotes for orders of any size (initially the rule required trades of 100 or more lots; later this was reduced to 50 or more lots) without specifying the direction (buy or sell).

> By giving banks the ability to quote for size directly into the futures markets, FINEX can offer contracts that don't meet traditional profiles of trading activity. It allows FINEX to list contracts for which the measure of liquidity will not be volume and open interest, but rather the efficiency of the bank arbitrage to the cash market.[13,14]

Indeed, the idea is to enhance the ability of the futures market to work in concert with the underlying cash market, and thereby create a market that combines the liquidity of the interbank currency market with the guarantees of an exchange-based clearing association and the integrity of a CFTC-regulated marketplace. If the FINEX is successful at facilitating the arbitrage-based connection between its markets and the underlying currency markets, then it should be possible for the FINEX markets to be liquid, even when volume is low. Nevertheless, many would find this result surprising because the existence of a low-volume, liquid futures market is highly counterintuitive under the old paradigm where liquidity comes from the active trading of locals.[15]

Furthermore, the large-lot rule should make it easy to test for liquidity. After all, the large-lot rule results in quotes (bid and ask) for orders of various sizes at which market participants are prepared to trade. Thus, by examining these quotes, one could determine the functional

[12]Interestingly, because the USDX is a geometric basket of currencies no perfect arbitrage is possible. In fact, in the absence of transaction costs, one can build a dynamic portfolio that always outperforms a geometric index. Nonetheless, there are risk-arbitrageurs ready to tap the liquidity of the interbank currency markets and compete with each other to make markets in this contract.

[13]Taylor, Hunt, "A Paradigm for Futures Markets," *The Journal of Foreign Exchange and Money Markets*, (Nov 1994), Vol. 1, Issue 4.

[14]In December 1995, the Chicago Mercantile Exchange announced plans to introduce a large lot rule of its own.

[15]There are also other reasons why this result may be surprising. For example, because the arbitrage is a risk arbitrage, not a riskless arbitrage, there are issues of asymmetric information and adverse selection. Standard microstructure theory would argue that market makers in both markets would make wider markets in the market where they are subject to asymmetric information and adverse selection. (See O'Hara, Chap. 8 for an excellent treatment of this subject.)

relationship between order sizes and bid–ask spread over a range of order sizes, a relationship that is ordinarily unobservable in traded markets.

Unfortunately, at least from the perspective of a researcher, the large-lot rule is rarely used. Although there is plenty of anecdotal evidence that the bid–ask spreads for large orders are small, the large-lot rule is used too infrequently to provide a means for measuring liquidity.[16] When asked why the rule is not used, market participants generally respond that it is not needed. Indeed, floor traders contend that with the arbitrageurs prepared to trade using normal open-outcry trading procedures, the markets are sufficiently liquid to handle orders of any size. Thus, even though the lack of trading through the large-lot rule may support the idea that these markets are liquid even if volume is low, the rule could not provide the means for testing whether that liquidity actually exists.

To conduct such a test, therefore, another methodology is needed, and this article pioneers one for measuring liquidity in futures markets. This new methodology takes advantage of common trading practices in futures markets and the institutional details of the trade-matching systems employed at most exchanges to identify whether a series of trades is buyer or seller initiated. It then looks at successive trades in the series to search for systematic price responses. Using this methodology, which is described in greater detail in the next section, as well as more traditional analyses of spreads, the article searches for signs of illiquidity in the Dublin-based trading of FINEX's USDX contract. Because no signs of illiquidity are found, the article challenges the common belief that high volume is necessary to support a liquid marketplace.

The article also examines and confirms the seemingly extraordinary claim by exchange officials that bid–ask spreads in this low-volume market actually narrow more prior to the execution of larger trades than smaller ones. Under the old paradigm, where liquidity is provided by poorly capitalized locals, one would expect spreads to widen as order sizes increase because of the correspondingly greater risk of holding the position. Indeed, exchange officials argue that for spreads to narrow as order sizes increase it must be true that liquidity is being borrowed from the interbank market. Unlike futures markets where liquidity is provided by locals, liquidity in the interbank market is provided by large, well-capitalized institutions that enjoy economies of scale enabling them to tighten quotes and work for smaller profit margins when order sizes are larger.[17]

[16]In fact, the rule was used only once during the three-month period examined for this article. Because no trade occurred, the order size was not captured, although, by the rules of the Exchange then in effect, the quote had to be for at least 100 contracts. The spread quoted was three ticks wide: 8840 bid, 8843 asked.

[17]This argument may be found in Taylor (1994).

In summarizing the analytical results, the first and most important finding is that no significant signs of illiquidity in FINEX's low-volume, Dublin-based, Dollar-Index market are detected. The new methodology allows regressions of price response against various measures of order size for buyer- and seller-initiated trades separately. If the market is illiquid, one would expect to find correlations between price movement and order size: positively for buy series and negatively for sell series. While one cannot demonstrate liquidity using this methodology, one would expect to detect illiquidity if it existed. In none of these regressions are the slopes significantly different from zero, suggesting no price response to order pressure on either the buy or sell side. This would be surprising if the market is small, isolated, and illiquid. However, it is not at all surprising for a small market linked by arbitrage to a much larger, far deeper market. After all, if the interbank market is providing the liquidity in the pit, then it would be reasonable to suppose that the arbitrage bounds would be similar for all arbitrageurs in the pit. Under this paradigm, therefore, it is not surprising that a broker could break up and execute a series of trades at similar prices with different members of the pit.[18]

The second finding is that spreads narrow prior to trading. The distribution of spreads quoted prior to a trade is significantly tighter than either the distribution of all spreads quoted on the exchange or the distribution of spreads quoted just prior to the close of trading. Since it is these other spreads that are most frequently monitored, while the spreads actually in existence just prior to trading are the ones that really matter, any analysis of this market based on the more easily monitored spreads would overstate the transaction costs of the market.[19]

Third, the narrowing process is relatively symmetrical. Although brokers executing buy orders do pay a greater portion of the spread than the brokers executing sell orders receive (confirming the oft-stated idea that liquidity providers are in the business of earning the spread), the difference is not large. Indeed, on average, buyer-initiated trades execute about 59% of the way between the bid and the ask; whereas, seller-initiated trades execute about 49% of the way between the bid and the ask. Furthermore, the narrowing process itself is also fairly symmetrical. When one examines how spreads tighten prior to a trade by tracking the series

[18]The authors would like to thank an anonymous reviewer for pointing out this argument for the reasonableness of the result under the arbitrage paradigm.

[19]As one reviewer pointed out, this pattern is also found ". . . in other low-volume markets, such as inactive option series and deferred months of active futures. Reported bid–ask spreads [in low-volume markets] are usually wide, used mainly to provide an indication only, until someone expresses a genuine trading interest." This is exactly the reason given by traders in this market for this phenomenon.

of bid–ask spreads that precede trading, one finds that they tighten in essentially the same way regardless of whether the trade ultimately executed is the first in a series of buy trades or the first in a series of sell trades. In fact, the average locations of the trade relative to the earlier, prenarrowed bid–ask spread is about 53% of the way from the bid to the ask for buyer-initiated trades and 49% of the way for seller-initiated trades.

Finally, as pointed out by both Exchange management and traders on the floor, not only are spreads narrower prior to trading than at other times, but among the spreads in existence prior to trades, they are narrowest prior to the largest trades.

Taken together, these analyses suggest a liquid market where the source of liquidity is not just the active trading of locals, but the risk-arbitrage induced link to the interbank currency markets.

The next section presents a description of the contract examined, the data and data sources, and the methodology. The following section reports the empirical results, beginning with analyses of how spreads behave in the Dollar-Index market and concluding with analyses of price movement in response to buy and sell orders of various sizes. The final section offers conclusions and summary comments.

DATA, SOURCES, AND METHODOLOGY

The data consist of all trades in the March 1995 USDX contract executed on the FINEX trading floor in Dublin, Ireland, during the first quarter of 1995. Trading in Dublin occurs during the hours of 3:00 AM to 8:00 AM eastern standard time, while the New York trading floor is closed.[20]

The data set includes 50 trading days from the first trading day in January through the middle of March. During that period, the contract, which is equal to $1,000 times the U.S. dollar index, ranged in price from 83 to 90, notional values between $83,000 and $90,000. The minimum allowable price movement in the contract was one tick, equal to .01 times a USDX point, or $10.

Figure 1 presents the daily volume of the USDX contract in Dublin during the study period, which averaged 278 contracts per day. While trading ranged from a low of about 30 contracts to a high of almost 2000, this average included five trading days where volume was close to or exceeded 500 contracts. Median volume, perhaps a more informative mea-

[20]Trading on the New York floor runs from 8:05 AM to 3:00 PM. To accommodate Pacific-rim markets, trading resumes in small room off the main trading floor in New York during the hours of 7:00 PM to 10:00 PM.

sure, was about 220 contracts. Closing spreads ranged from 1 tick ($10) to 13 ticks. The average was 4.86 ticks. By comparison, volume in New York averaged 1656 contracts a day with a median volume of 1348 contracts.

Data for the study are from two primary sources. The first source is the daily *Broker Reconciliation Report*. This report includes all data necessary for brokers to reconcile their trading records with the exchange's official records. The report lists all transactions executed on the floor and includes, for both sides of the transaction, broker and clearing-member numbers, order-type category, and customer account number, as well as commodity code, expiration month, price, quantity, time of execution (to the nearest minute), and a computer-generated transaction number.[21] The subset of the *Broker Reconciliation Report* provided was stripped of any data that might permit reconstruction of customer or clearing-member positions.

The second data source is the daily *Time and Sales Report*. This report shows every quoted bid and ask captured by the price reporters at the head of the trading ring, as well as the price of every executed trade. The data for this report are captured real-time by reporters located on a platform at the head of the trading ring, and include a time stamp showing the time of each entry. These data are the source of the live feeds providing real-time prices to the electronic networks and, hence, to trading rooms around the world.

Using the *Broker Reconciliation Report* and common trading practices in open-outcry futures markets, a new methodology is developed for determining whether trades are buyer or seller initiated. In open-outcry trading, instead of executing large customer orders in one single trade, brokers often parcel out large orders into a series of smaller trades. For example, rather than execute a single transaction with one broker for 300 contracts (lots), a broker is more likely to execute a series of successive transactions with several brokers for fewer lots, for example, 100, 75, 50, 75.[22]

Because of this practice, it is possible to identify series of successive trades, one rapidly following the other, where one broker is on the same

[21]U.S. exchanges are required by CFTC regulation (section 1.35) to maintain these data for audit purposes for a period of five years (section 1.31).

[22]The trade-matching protocols used by the Exchange split single transactions into transactions reported on multiple lines of the Broker Reconciliation Report for a number of reasons. For example, if two brokers execute a 20-lot trade but one of them assigns half the contracts to one account and half to another account, then the report would show two trades for 10 lots each. By using the transaction number included in the Broker Reconciliation Report, however, which would be the same for both of those trades, it is possible to differentiate between legitimate trade series and individual trades reported on multiple lines because of the trade-matching protocols.

side (buy or sell) of each trade in the series. Whenever this is the case, it is assumed that all of the trades in the series are the result of an order from the side on which that broker traded. That is, if the common broker is always a buyer (seller), then it is assumed that the trades in the series are initiated because of a buy (sell) order, and the trades are classified as buyer (seller) initiated.

In particular, a trade series is defined as a collection of successive transactions, each within two minutes of the one before, all with a single broker on the same side (buy or sell) of every trade in the series.[23] When compiling these data, crossed trades are ignored because the rules enabling crosses limit their effect on market price in an open-outcry market.[24]

During the period covered by this article, there are 75 buy series consisting of 179 trades for 4692 contracts, and 73 sell series consisting of 168 trades for 4098 contracts. These series are not unique, however, because of the small number of brokers trading in Dublin. After series that appear in whole or in part as both buy and sell series are eliminated, there are 48 unique buy series consisting of 122 trades for 2791 contracts, and 44 unique sell series consisting of 106 trades for 2346 contracts, none of which overlapp.[25]

This mechanism for classifying trades as buyer (seller) initiated contrasts sharply with the mechanisms used by others. The rule used by Vijh (1990), for example, in his study of CBOE equity options, assumed that any trade occurring above the bid–ask midpoint was buyer initiated, and any trade occurring below the midpoint was seller initiated. By using that measure, Vijh made the assumption that the trade initiator must pay more than half the spread to induce the other side to trade. This assumption is inconsistent with the trading behavior observed in the USDX market, where over one third of the trades launching buy series were executed below the midpoint of the spread and close to half of the trades launching sell series were executed above the midpoint.

[23]The choice of a maximum of two minutes between trades is somewhat arbitrary. A one-minute rule would break some series and eliminate others; a three-minute rule would add only a few additional series.

[24]Crosses occur when a single broker matches buy and sell orders from his or her own personal deck of customer orders. Crosses are allowed only after the broker has announced both the intention to trade and the price to the other brokers in the ring and has offered to either buy or sell at that price with anyone else. If no broker responds, the crossing broker may execute the trade alone; thus crossed transactions do not involve any other brokers and occur within the bid and ask of the currently quoted market.

[25]The astute reader might note that 27 buy series (75–48) are eliminated because of overlap with sell series, while 29 sell series (73–44) are eliminated because of overlap with the buy series. The reason these numbers do not have to be the same is that the series do not line up perfectly. For example, a buy series of four transactions may line up with two sell series (because of two different selling brokers) of two transactions each.

The combination of identifying series of trades and labeling them as either buyer or seller initiated facilitates breaking the question of measuring liquidity into the two parts recommended by Massimb and Phelps (1994). They characterized liquidity as a combination of a market's ability to provide immediate execution of standard-sized incoming market orders (often called immediacy) and to execute large orders without large changes in price (often called market depth or price resiliency).

This article, therefore, examines liquidity in two parts. First, the effective bid–ask spread is examined to measure the effective price of transacting standard-sized orders in the market. Second, the market's price response to buy and sell orders of different sizes is examined to assess the price resiliency of the market.

Because it is generally difficult to determine baseline prices, most liquidity studies focus on only the first of these two components, the effective bid–ask spread. While this article also conducts the standard analyses of spread width, because of the ability to classify some trades as either buyer or seller initiated, it also conducts a broader set of analyses. In particular, it examines the widths of spreads quoted just prior to buyer- or seller-initiated trades and compares them with each other and with closing spreads and the general population of spreads.

More importantly, however, the methodology for classifying series of trades as buyer- or seller-initiated also facilitates a unique analysis of price resiliency. Because the methodology identifies not just individual trades but series of trades as buyer- or seller-initiated, it is possible to examine how far the price moves as each successive trade is executed. Thus, using the previous example, although it is not possible to say anything about how far the price moved before the first 100 lots traded, it is possible to explore how much further the price moves as each successive trade in the series (in the example for 75, 50, and 75 lots) is executed.

Indeed, the price response to the order pressure is examined in two different ways. On the one hand, the analysis groups together the follow-on trades and looks at the difference between the price at which the first hundred lots trade and the average price at which the following 200 lots (75, 50, and 75) trade. Thus, this grouping provides a measure of the price response to execute the follow-on order for 200 contracts. On the other hand, the price changes are viewed successively; that is, the price change from the first to second trade in the series is associated with an order for 75 lots. Similarly, the price change between the second and third trade is viewed as the price change associated with an order for 50 lots, and so on.

This is the most unique component of the methodology because it provides a mechanism for measuring resiliency. Because, by definition, a single broker is on one side of all trades in a series, one can determine whether the series of trades represents buyer or seller pressure. Thus, as long as the successive transactions in the series are executed in relatively short order, data can be collected on price movement due to buyer or seller pressure from orders.

In particular, once the trade series is identified and classified as buyer or seller initiated, it is possible to regress price movement against order size. If the market is illiquid, one would expect to find a positive correlation between price movement and order size for buy series and a negative correlation for sell series. While one cannot demonstrate liquidity using this methodology, one would expect to be able to detect illiquidity if it existed. In none of these regressions are the slopes significantly different from zero.

TESTS OF LIQUIDITY

This section is divided into two major subsections. The first subsection explores the transaction costs of the market by examining spreads. It begins with an analysis of the differing populations of quoted spreads and demonstrates that the spreads quoted just prior to trading are narrower than both spreads normally quoted and spreads quoted at the end of the day. The discussion on spreads continues with an analysis of the bargaining process and shows that the wider initial spread narrows to the narrower final spread in essentially the same way whether the upcoming trade is the result of a buy or sell order. It concludes by confirming the Exchange's assertion that spreads are narrowest prior to the largest trades, lending support to the argument that the primary source of liquidity in the USDX contract is arbitrage, not the trading of locals.

The second subsection examines price resiliency. Here a number of regressions are performed to explore the market's price response to buy and sell pressure represented by buy and sell orders of various sizes. If the market is illiquid, one would expect to detect a greater upward price movement in response to larger buy orders than to smaller ones, and a greater downward price movement in response to larger sell orders than to smaller ones. None of the regressions indicates significant price response to orders.

Analysis of Spreads

Differing Populations of Quoted Spreads

Although traders bid to buy and offer to sell throughout the day, trading is not continual, and during the more "inactive" periods the market is very slow.[26] Moreover, during these inactive periods, not only are quotes infrequent, but the spreads quoted tend to be wider. (Indeed, during inactive periods, if asked for a quote, unless the floor community believes there is serious interest in trading, the quotes offered are frequently off-the-cuff and ten ticks wide.)[27] To confirm the observation that spreads quoted just prior to a trade are different from the others, three populations of spreads are examined: spreads quoted just prior to a trade, all other spreads, and spreads quoted at the end of the day.

From the *Time and Sales Report,* 2776 bid–ask spreads are identified during the investigation period for the March USDX contract that meet the following criteria: the bid and ask are quoted within five seconds of each other and no trade occurs between the bid and ask. Any bid or ask that is quoted without its complement's being quoted within five seconds is dropped from the data set.[28]

The data are then segmented into three groups. Group 1 includes the last spread quoted prior to a trade, as long as the last quote is within two minutes of the trade.[29] Group 2 includes all other spreads in the data set. Group 3 consists of the final quote of each day (because no trades ever occur after the last quote, Group 3 is a subset of Group 2).

Table I presents the number of quotes in each group and the average spread in ticks. As is obvious from the table, the average spread just prior to a trade is significantly narrower than the average spread in either of the other two groups (numbers in parentheses represent t-statistics for the comparison with Group 1).

Figures 2 and 3 present graphs comparing the distributions of Group 1 to Group 2 spreads and Group 1 to Group 3 spreads, respectively. As is apparent from the table and exhibits, not only are the means of the Group 1 spreads significantly different from those of the other groups, but so are the distributions themselves (with chi-squares of 3787 and 1016, respectively). Group 1 spreads (just prior to a trade) are nearly

[26]Indeed, even in New York where the volume is far greater than in Dublin, there are periods during the day when the brokers and locals in the pit sit on the floor, read newspapers, and wait for something to happen.

[27]Many thanks to an anonymous reviewer who pointed out that it is not uncommon in other low-volume markets, like inactive option series and deferred months of active futures, for effective bid–ask spreads to narrow just prior to trading. As in the USDX market, reported spreads in low-volume markets are often used mainly to provide an indication only until there is genuine trading interest.

[28]As with the other time intervals in this article, the choice of five seconds was somewhat arbitrary. The statistical result presented, however, is striking enough that it is hard to imagine a different interval having much effect. Note that we also eliminated quotes where either side (bid or ask) was later revised by a floor reporter.

[29]The interval, although again arbitrary, is chosen to be consistent with the interval used in defining trade series.

TABLE I

Comparison of the Average Spread by Group

Group	*Number of Quotes*	*Average Spread (in Ticks)*
Group 1 (just before a trade):	353	2.35
Group 2 (all other quotes):	2423	4.92
		(20.79)
Group 3 (day-end quote):	50	4.86
		(6.60)

always one tick and seldom more than three to five ticks, which contrasts sharply with the distribution of spreads in the other two groups.

An immediate consequence of this analysis is that conclusions concerning the liquidity of the USDX market in Dublin based on examinations of overall spreads or end-of-day spreads are misleading. Although those are the most easily monitored spreads, the true spreads faced by traders in this market are far tighter. Thus, the effective bid–ask spread component of liquidity in this market is far better than the more cursory analyses of spreads would suggest.

Another conclusion that one might draw from the preceding analysis is that trading in the USDX is characterized by an aggressive bargaining process that causes spreads to narrow whenever real interest in trading exists. To begin to understand this bargaining process, tests are conducted of whether the frequency of quotes is correlated with the narrowing of spreads. To do so, the size of spread (*spreadwidth*) is regressed against the number of minutes since the last quoted spread (*min*); this analysis is repeated while controlling for whether the spread is the last one quoted prior to a trade using the dummy variable (*lasttradedummy*), which is set equal to 1 if the spread is the last one quoted prior to a trade. These regression results are presented in Table II.

Intuitively, these results make sense. As trader interest increases, one would expect the time between quotes to shorten and the spread to narrow. Moreover, because the tightest spreads are quoted just prior to a trade and because experienced traders understand this dynamic, trades do not occur until the spread has tightened. This confirms, therefore, that any evaluations concerning the average transaction costs based on spread data in the USDX marketplace should be based on spreads quoted just prior to trading; evaluations based on average spreads or end-of-day spreads significantly overstate the costs of transacting in this market.

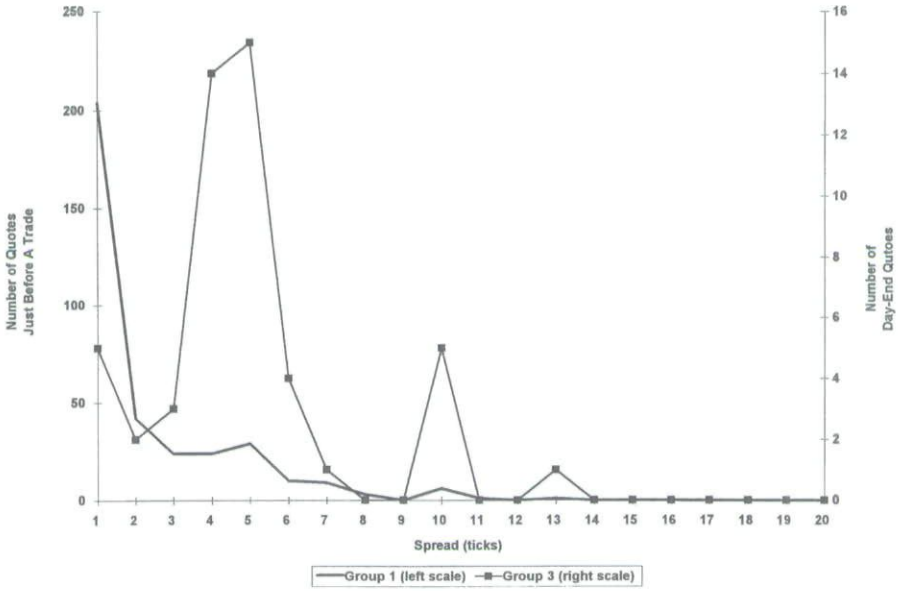

FIGURE 3

Comparison of spreads: Group 1 (just before trades) versus Group 3 (day-end).

TABLE II

Spreads Narrow Prior to Trading

	Model 1 *Spreadwidth* $= a + b_1(min)$ $(n = 2{,}723)$	
Variable	*Coefficient*	*T-Statistic*
Intercept	4.0371[a]	(67.47)
min	0.1029[a]	(14.27)

	Model 2 *Spreadwidth* $= a + b_1(min) + b_2(lasttradedummy)$ $(n = 2{,}723)$	
Variable	*Coefficient*	*T-Statistic*
Intercept	4.4385[a]	(72.44)
min	0.0814[a]	(11.72)
lasttradedummy	−2.3006	(−17.41)

Location Analyses

Using the new methodology for classifying trades as buyer or seller initiated, these analyses of spreads can be extended in several new ways. In particular, it is possible to explore further the question of whether the narrower final spread in place just prior to a trade accurately reflects the true transaction costs of the market. This can be done in three ways.

First, one can examine whether (or to what extent) brokers executing customer orders to buy pay more of the spread than the brokers executing customer orders to sell receive. While one would expect to see some difference (after all, common wisdom dictates that liquidity providers are in the business of earning the spread), it would be hard to label a market as liquid if buyers' brokers regularly transacted at the asking price, while sellers' brokers regularly received only the bid. Second, one can explore the question of whether the narrowing process itself is biased against the initiator of the transaction. Third, one can ascertain whether the sizes of the buy and sell orders affect the spread.

To investigate the first question, the location of a transaction relative to the spread is defined to be the percentage distance from the bid to the ask. Thus, a 0 (or 0%) represents an execution at the bid, a .5 (or 50%) represents an execution at the midpoint between the bid and ask, and a 1 (or 100%) indicates an execution at the ask. Next, for each buyer- and seller-initiated trade series, the final spread in place just prior to the first trade of each series is identified using the *Time and Sales Report,* and the location of the transaction is calculated. Following the same timing rules as before, this examination yields active spreads prior to the first trade for 64 of the 75 buy series and for 59 of the 73 sell series. Of these series, 41 of the buy series and 36 of the sell series are in the set of nonoverlapping series.

On average, buyer-initiated trades are executed approximately 59% of the way between the bid and the ask, whereas seller-initiated trades are executed approximately 49% of the way between the bid and the ask. A difference, but hardly striking (the t-statistic for the difference in means is 0.82). Figure 4 presents a histogram showing the percentage of trades executing below the bid, at the bid, between the bid and midpoint of the spread, at the midpoint, between the midpoint and the ask, at the ask, and above the ask, for the buyer- and seller-initiated trades examined (the chi-square is 0.61).[30]

[30]As is evident from Figure 4, a system that classifies trades as buyer or seller initiated based on where the trade is executed relative to the midpoint of the spread would be inconsistent with observed trading behavior in this market.

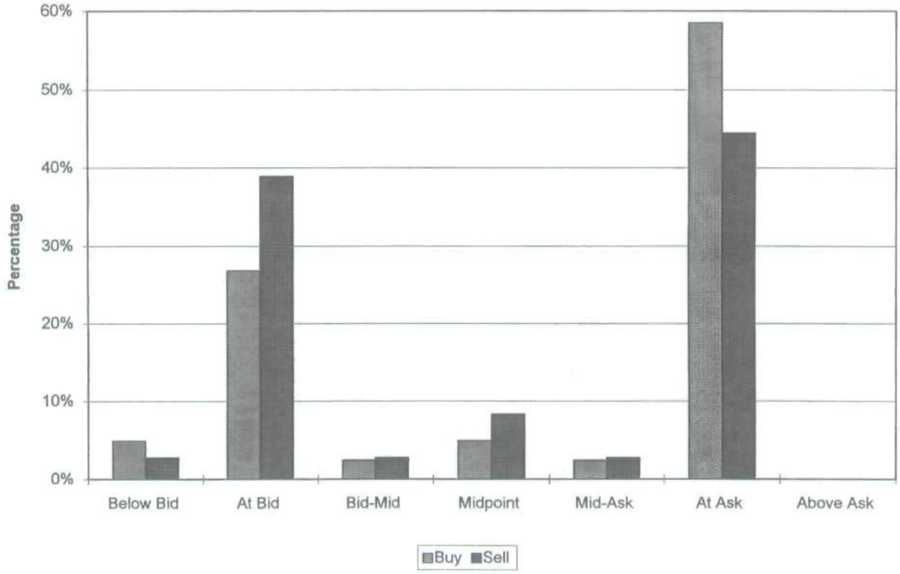

FIGURE 4
Location distributions within final spread—nonoverlapping data (41 buy series, 36 sell series).

This examination, however, also raises an interesting secondary question: If there is no significant bias relative to the final spreads, is the narrowing process itself biased? After all, it is possible that although the average location of the execution relative to the final spread is relatively indistinguishable for buyer- and seller-initiated trades, there may still be a difference between the groups resulting from the process by which the earlier, wider spreads narrow.

In an illiquid market, for example, one might expect narrowing in response to buyer pressure to result in upward movement of the bid toward the ask, while for seller-initiated transactions the narrowing might occur through downward movement of the asking price toward the bid. Whereas in a liquid market (where, by definition, price does not respond to buy or sell pressure), one would expect to see a narrowing process that treats buyer- and seller-initiated trades symmetrically.

Figures 5 and 6 present examples of these alternative narrowing processes. In Figure 1, for a hypothetical illiquid market, the bid moves upward as the spread narrows for an upcoming buyer-initiated trade. One could also imagine the reverse for an upcoming seller-initiated trade. In Figure 6, for a hypothetical liquid market, the narrowing is more symmetrical.

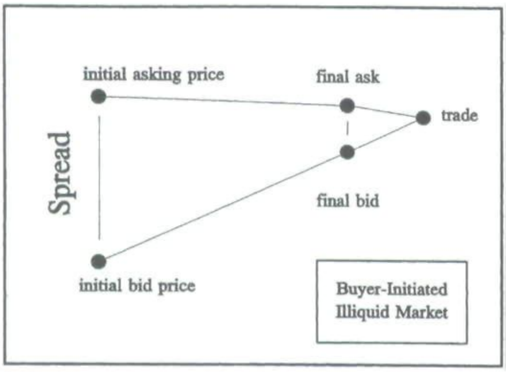

FIGURE 5
Hypothetical narrowing process:
Illiquid market.

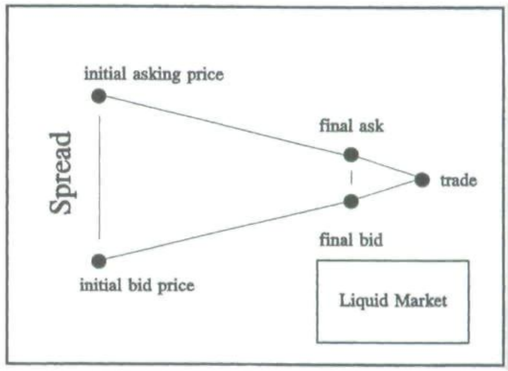

FIGURE 6
Hypothetical narrowing process:
Liquid market.

This question can be explored using a location analysis similar to the one performed above, this time relative to an "initial" spread in place prior to the beginning of the bargaining process. To identify the initial spread, a series of spreads are identified working backward from the trade. To be included in the series, every spread in the series must occur within three minutes of the transaction, and the time between successive spreads can not exceed 45 seconds. Using these rules, it is possible to identify initial spreads for each of the buyer- and seller-initiated transaction series that have "final" spreads. While the timing rules are somewhat arbitrary, these spreads represent reasonable estimates of the prior wider spreads and provide a good base line from which to identify the process by which the spread narrows.

Using these initial spreads, one can repeat the location analysis to identify any differences in value conceded by either side en route to meeting at the eventual transaction price. The difference in mean location is even less pronounced for the initial spreads: on average, buyer-initiated trades are executed approximately 53% of the way between the bid and the ask, while seller-initiated trades are executed approximately 49% of the way between the bid and the ask (the t-statistic for the difference in means is 0.40). Figure 7 presents the comparable histogram relative to the initial spread in existence prior to the beginning of the bargaining process (the chi-square of 2.60 does not suggest a significant difference). These data, the average location values relative to both the final spread and the initial spread, along with t-statistics and chi-squares testing the differences in the buyer and seller distributions, are summarized in Table III.

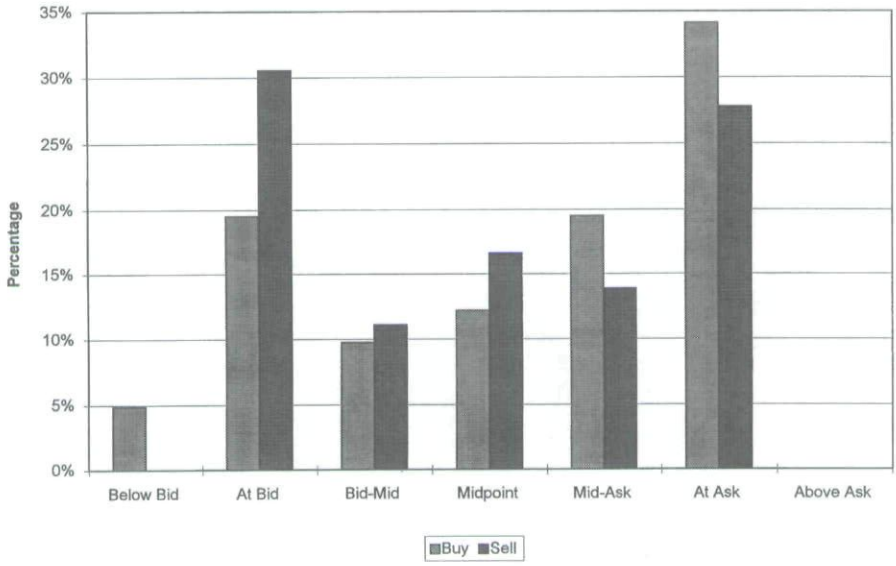

FIGURE 7

Location distributions within initial spread—Nonoverlapping data (41 buy series, 36 sell series).

TABLE III

Average Location Comparisons for Nonoverlapping Buy and Sell Trade Series Relative to Final and Initial Spreads

Group	*Buy Series*	*Sell Series*	*T-Statistic*	*Chi Square*
Final spreads	0.592	0.492	0.824	0.610
Initial spreads	0.530	0.488	0.401	2.602
	$n = 41$	$n = 36$		

Figures 4 and 7 reveal a somewhat greater propensity for buyers to pay the ask and for sellers to receive the bid. While these differences are not statistically significant, they are certainly believable. What is most surprising is that the differences between buyer- and seller-initiated trades are not more pronounced with buyer-initiated trades more concentrated at the ask, and seller-initiated trades more concentrated at the bid. After all, as seen in Figures 2 and 3, the distributions of spread widths just prior to trade execution is quite narrow, with a mode of only one tick.

Third, the question of whether the location within the spread is dependent on the size of the order is addressed. After all, if the market is

TABLE IV

Regression of Location on Trade Size

Variable	*Location* = *a* + *b*(*tradesize*) *Model 4-1*	*Model 4-2*
Buy Series (n = 41)		
Intercept	0.617[a]	0.553[a]
	(6.406)	(6.391)
ordersize	−0.001	−0.001
	(−0.669)	(−0.069)
Sell Series (n = 36)		
Intercept	0.460[a]	0.510[a]
	(4.370)	(5.539)
ordersize	−0.000	−0.001
	(−0.456)	(−0.037)

[a]Significant at the 0.05 level.
Model 4-1: *Location* is the location within the final spread.
Model 4-2: *Location* is the location within the initial spread preceding the narrowing process.
In both models, *Tradesize* is the size of the first trade in the series.

only illiquid at larger order quantities, and buyers (sellers) pay (receive) the ask (bid) for big orders, it would hardly matter what they pay (receive) for small orders. To explore this question, the location of the execution within the spread is regressed against the size of the trade. These results are presented in Table IV. As is evident from Table IV, there is no relationship between order size and the location of the execution relative to either the final or initial spread for either buyer- or seller-initiated trades. Thus, none of the location analyses suggest illiquidity in this market.

The Larger the Trade, The Narrower the Spread

To complete the analyses of spreads, the extraordinary claim made by both floor traders and Exchange management that the larger the trade the tighter the spread is explored. As noted earlier, such a market dynamic would be more consistent with a market whose source of liquidity is arbitrage rather than the active trading of locals. To explore this question, one can look again at the nonoverlapping buy-and-sell-series data and regress the size of the spread before the first trade in the series against two measures of the series order size: the size of the first trade and the total volume traded. As before, for the size of the spread, one can use either the final spread or the initial spread in place prior to the narrowing process. Furthermore, because an examination of the residuals suggests that one buy order for 500 lots in the first trade may be skewing the results (the next highest initial trade is for only 100 lots and there are several such trades), the regressions can be run with and without that data point.

The results are summarized in Table V. Because the test concerns whether spread width decreases as order size increases, a one-tailed test is appropriate. As is evident from the table, there is indeed a statistically significant relationship between the size of the final spread and the size of both the first trade and the volume of the entire series. On the other hand, as one would hope, there does not appear to be much of a relationship between order size and the size of the spread prior to the narrowing process. Naturally, because the two measures used for order size (size of the first trade and total volume in the series) are highly correlated, it is not surprising that results using the alternative choices of order size are very similar.

It also merits noting that regardless of statistical significance, the sign of the *ordersize* coefficient is always negative.

Generalizability of Spread Analyses

The prior analyses focused on buyer- or seller-initiated trade series. Obviously, these trades were chosen because the existence of the series enabled the determination of whether the trade represented buy or sell pressure. Nevertheless, one could question whether the results are generalizable to the market as a whole.

In theory, prior to a trade, only the initiating broker knows how many contracts are involved. A broker offering to sell 3 contracts may be selling only 3, or may be selling 20. To other market participants, therefore, the true size being offered is indeterminate. Traders in the trading rings, however, earn their living by identifying and exploiting telltales that indicate whether such offers are for stand-alone trades or represent greater buying or selling pressure. That is, after all, exactly what trading in the ring is all about.

Thus, the greater the ability of other market participants to determine whether these transactions represent the beginnings of series, rather than stand-alone trades, the more one-sided the results should be. After all, knowing that there is additional buying (selling) pressure would provide an advantage. Thus, if the results are not directly generalizable, they are probably conservative.

Price Resiliency

To address the question of price resiliency (the extent to which the price moves in response to orders of different sizes) one can again begin with

TABLE V

Regression of Spread Width on Order Size

Variable	*SpreadWidth* = *a* + *b*(*ordersize*) Model 5-1	Model 5-2	Model 5-3	Model 5-4
Buy Series Only (n = 41)				
Intercept	2.302[a]	2.414[a]	3.612[a]	3.681[a]
	(6.524)	(6.465)	(9.991)	(9.461)
ordersize	−0.004	−0.004[b]	−0.005	−0.004
	(−0.970)	(−1.300)	(−1.235)	(−1.240)
without outlier (n = 40)				
Intercept	2.841[a]	3.038[a]	3.692[a]	3.693[a]
	(5.918)	(6.126)	(7.206)	(6.848)
ordersize(−0.032[a])	−0.032[a]	−0.018[a]	−0.010	−0.004
	(−1.806)	(−2.198)	(−0.507)	(−0.468)
Sell Series Only (n = 36)				
Intercept	2.139[a]	1.990[a]	3.259[a]	3.186[a]
	(8.535)	(8.645)	(7.526)	(8.242)
ordersize	−0.017[a]	−0.004[a]	0.016	−0.005[b]
	(−2.203)	(−1.719)	(−1.207)	(−1.298)
Buy and Sell Series (n = 77)				
Intercept	2.119[a]	2.213[a]	3.354[a]	3.442[a]
	(9.841)	(9.803)	(12.918)	(12.918)
ordersize	−0.005[b]	−0.004[a]	−0.006[b]	−0.004[a]
	(−1.459)	(−1.944)	(−1.460)	(−1.756)
without outlier (n = 76)				
Intercept	2.473[a]	2.327[a]	2.327[a]	3.477[a]
	(9.292)	(9.570)	(9.570)	(11.692)
ordersize	−0.023[a]	−0.006[a]	−0.006[a]	−0.005[b]
	(−2.567)	(−2.233)	(−2.233)	(−1.414)

[a]Significant at the 0.05 level; [b]Significant at the 0.10 level.
Model 5-1: *Spreadwidth* is the final spread; *Ordersize* is the size of the first trade in the series.
Model 5-2: *Spreadwidth* is the final spread; *Ordersize* is the volume of the entire series.
Model 5-3: *Spreadwidth* is the initial spread; *Ordersize* is the size of the first trade in the series.
Model 5-4: *Spreadwidth* is the initial spread; *Ordersize* is the volume of the entire series.

the buyer- and seller-initiated series of trades. Because it is unnecessary to associate a final spread with the first trade, the entire set of nonoverlapping series data (48 buy and 44 sell) can be used.

The underlying concept behind the analysis is that if the market is illiquid then one would expect to find a pronounced price response to sustained buying or selling pressure. That is, not only should buying pressure move the price upward while selling pressure should move the price downward, but also higher volumes of sustained buying pressure should move prices upward farther than lower volumes, and higher volumes of sustained selling pressure should move prices downward farther than lower volumes.

As discussed in the methodology, order size is measured in two ways. First, all trades in a series after the first are grouped together. In this "grouped" model, the order size is given as the total volume of the series less the size of the first trade. In this instance, the price change is defined as the difference between the weighted average price of the successive trades and the price of the first trade.

Second, order size is also measured successively. That is, each trade in the series is treated individually. In this "successive" model, the order size is the size of each trade (after the first), and the price change is defined as the difference in price between the trade being examined and the prior one.

Figures 8 to 11 present scatter plots showing the average price change and series volume for each of these data sets, as follows: Figures 8 and 9 present the grouped model for buy data and sell data, respectively, and Figures 10 and 11 present the successive model for buy and sell data, respectively.

As these exhibits illustrate, volumes are concentrated in the 1-to-50 lot range. Nevertheless, it is important to remember that average daily volume for the entire market is fewer than 300 contracts per day. Given that frame of reference, the volumes tested are relatively large.

The regression results are presented in Table VI. If the market is illiquid, one would expect to find nonzero slopes. In fact, one would expect significantly positive slopes for the buy-series regressions and significantly negative slopes for the sell-series regressions. The results in Table VI show that this is not the case, even when using one-tailed tests for significance.

While it is impossible to state formally that the market is liquid—all one can say formally is that because the slope coefficients are not significantly different from zero the hypothesis that the markets are liquid cannot be rejected—given the lack of statistical significance there does not appear to be a strong relationship between volume and price. In other words, prices in this market seem to be unaffected over the volumes tested, which suggests a liquid market.

Furthermore, it is interesting to note that the buy-series transactions have a negative slope, while the sell-series transactions tend to the positive side, exactly the opposite of what one would expect to see in an illiquid market. Again, while not significant, these results suggest that larger orders have *less* price impact than smaller orders, which is consistent with the argument that it is easier and less expensive to arbitrage

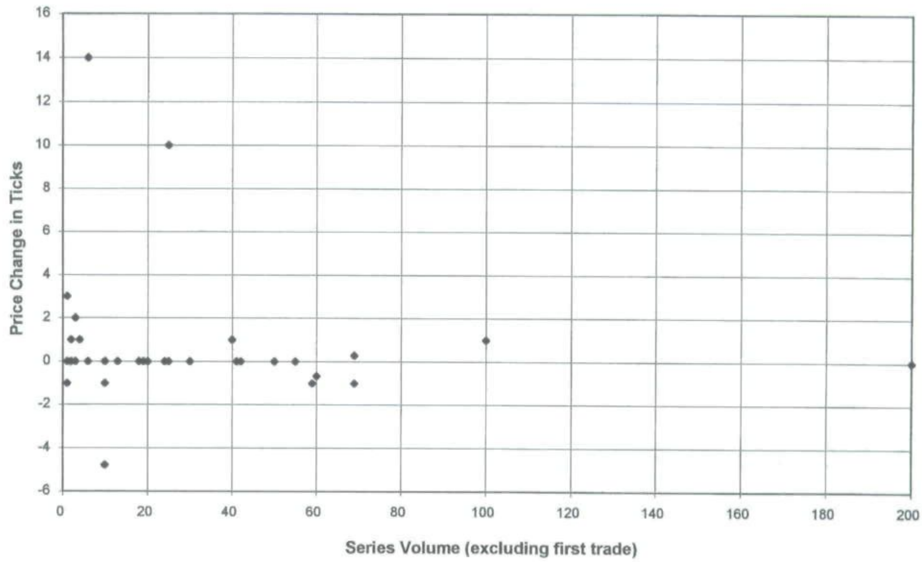

FIGURE 8
Buy series nonoverlapping group data (48 observations).

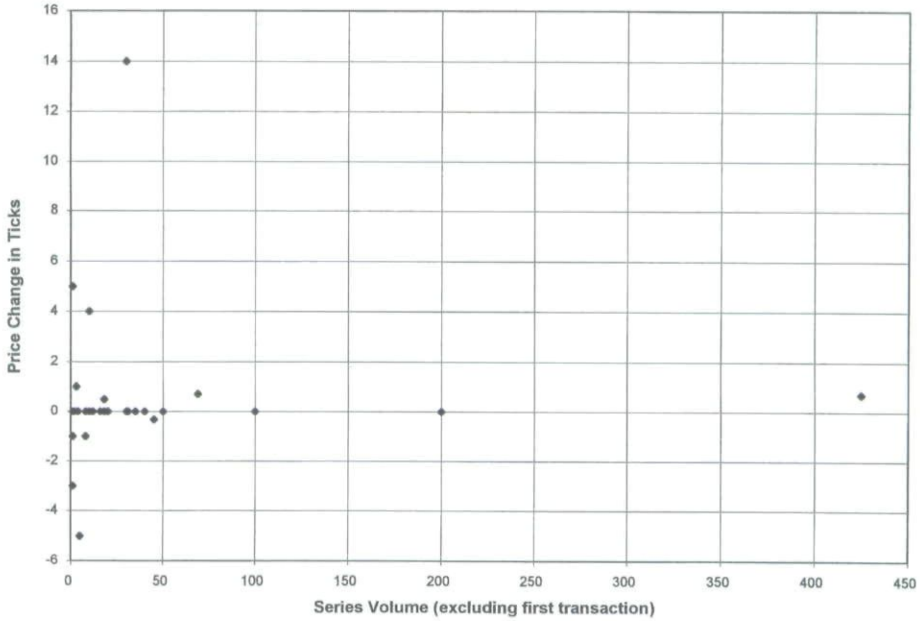

FIGURE 9
Sell series nonoverlapping group data (44 observations).

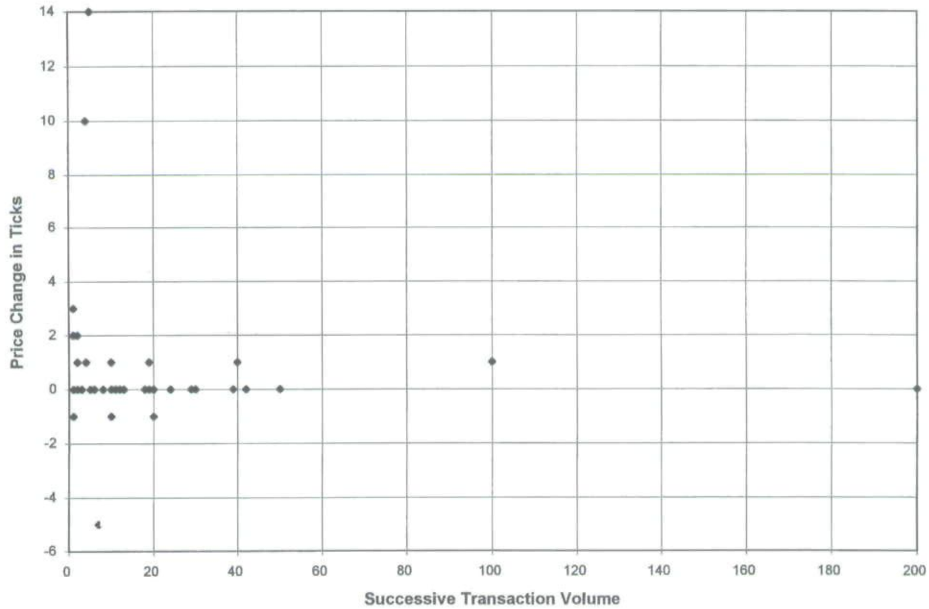

FIGURE 10
Buy series nonoverlapping successive data (74 observations without first trade).

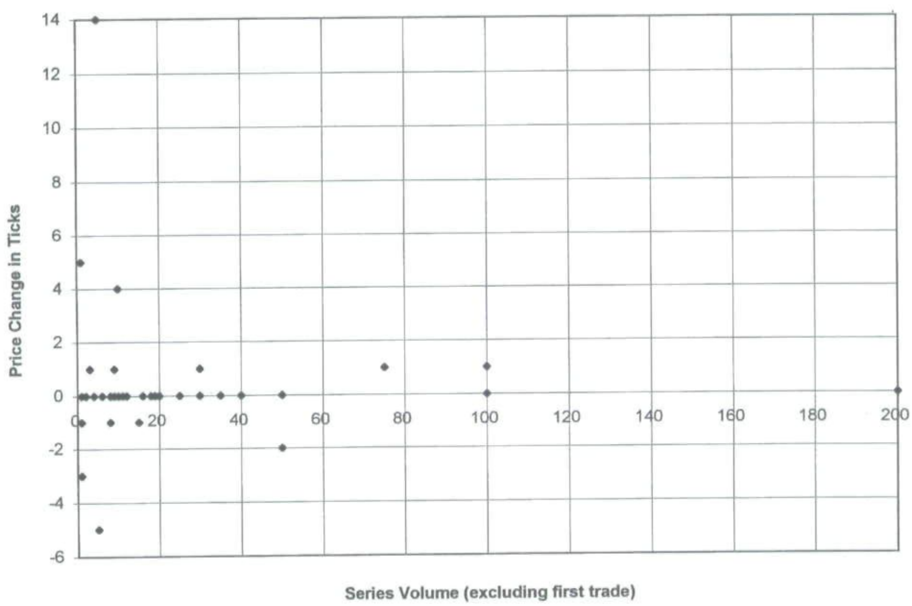

FIGURE 11
Sell series nonoverlapping successive data (62 observations without first trade).

TABLE VI

Regression of Price Change on Volume

$\Delta pricechange = a + b(ordervolume)$

Variable	*Model 6-1 Grouped*	*Model 6-1b Grouped w/o Outlier*	*Model 6-2 Successive*	*Model 6-2b Successive w/o Outlier*
Buy Series	(n = 48)	(n = 47)	(n = 74)	(n = 73)
Intercept	0.693[a]	0.804[a]	0.439[a]	0.535[a]
	(1.360)	(1.359)	(1.418)	(1.462)
ordervolume	−0.007	−0.012	−0.005	−0.012
	(−0.591)	(−0.673)	(−0.567)	(−0.736)
Sell Series	(n = 44)	(n = 43)	(n = 62)	(n = 61)
Intercept	0.311	0.282	0.275	0.254
	(0.741)	(0.602)	(0.834)	(0.722)
ordervolume	0.001	0.003	0.000	0.001
	(0.244)	(0.247)	(−0.002)	(0.011)

[a]Significant at the 0.10 level.

large orders into the interbank currency markets than small ones because standard-sized transactions are far larger in interbank trading. It is also consistent with a market maker whose variable costs (such as adverse-information costs, risk-bearing costs, and inventory-holding costs) are less substantial than its fixed costs [see Stohl (1989)]. Such a market maker would prefer to process larger orders, and might adjust prices to encourage such behavior. While this may characterize the situation for a back-office trading room making markets by arbitraging positions into the interbank currency markets, it is not how one would characterize a local whose variable costs associated with risk bearing and inventory holding far outweigh their fixed costs of doing business.

Finally, if one simply examines the scatter plots presented in Figures 8 to 11, it is fairly clear that once the lot size exceeds about 20 contracts, the price change rarely exceeds plus or minus one tick, and hardly ever exceeds plus or minus two. Thus, even were one to obtain substantially more data, one would not expect to find a significant upward slope for buy-series transactions nor a significant downward slope for sell-series transactions. In fact, one might conceivably obtain significance for the opposite relationship.

CONCLUSION AND SUMMARY COMMENTS

In conclusion, the authors hope that this article accomplishes two things. First, that it has a methodological impact. Traditionally, it has been hard to measure liquidity directly in markets. By capitalizing on specific institutional details of open-outcry trading and common trading practices, the article pioneers a mechanism for identifying series of buyer- and seller-initiated trades. With that identification in hand, it is possible to conduct far more powerful tests for liquidity than have been possible previously.

Second, the article questions the validity of using volume as a proxy for liquidity. Whenever arbitrage (even loosely defined risk arbitrage) is a significant possibility in a derivative market, it is very possible for the market to be liquid even without substantial volume.

Indeed, the analyses suggest that the USDX market in Dublin, a very low volume market, is far more liquid than its volume or other readily available market proxies would indicate. Effective spreads in place prior to active trading are significantly narrower than any calculation of average spreads or end-of-day spreads would indicate. Location calculations show surprisingly little difference in the prices paid by buyers and sellers, whether one looks at the location relative to the final spread or to the prior, wider spread before the narrowing process began. And most importantly, (1) the larger the trade, the narrower the spread, which suggests a market whose liquidity comes from arbitrage rather than the trading of locals and (2) there is no significant relationship between order size and price response over the volumes tested.

Indeed, the tests of the spread data show a sharp division between spreads active just before a trade and spreads at other times. Like at other low-volume markets, this result is due to the market dynamic in effect at FINEX. When there is little activity in the market, the market is relatively wide, reflecting the lack of interest. As interest increases, however, the spread tightens until a transaction is completed, which means that the effective transaction cost in this market is substantially less than would be indicated by a casual look at average daily spreads or closing spreads.

Moreover, the market exhibits no price response to relatively high volume orders. Of course, one cannot speculate about the price responsiveness to volume orders outside the range tested, and one certainly cannot estimate what the impact would be if one were to enter buy or sell orders that exceed the overall scope of the market. Nevertheless, the findings presented in this article are applicable over ranges that are significant given the size of this marketplace. A series of buy orders totaling 50 contracts represents one sixth of the volume during a normal trading day, yet for volumes of up to about 50 contracts (a notional value of 4 to 4 ½ million dollars), order size appears to have no appreciable impact on price.

In conclusion, therefore, while it is true that examining average daily spreads and daily trading volume and using them as proxies for liquidity might keep one from trading in truly illiquid markets, it also has the potential to cause a trader to overlook an attractive, liquid market. Given the competitive nature of securities trading, market participants might prefer to take a more comprehensive view of liquidity, for doing so might increase the number of markets one considers suitable for trading.

As long as risk arbitrageurs stand ready to trade in a derivative market and in deeper, underlying markets—and as long as there is enough competition among the arbitrageurs to prevent them from charging excess rents—then the derivative market may well be sufficiently liquid, even when volume is low. Only by identifying such markets and then examining both the behavior of bid–ask spreads and price response to order volume can one fully assess whether such a market is liquid.

BIBLIOGRAPHY

Baesel, J., Shows, G., and Thorp, E. (1983): "The Cost of Liquidity Services in Listed Options: A Note," *The Journal of Finance*, 38:989–995.

Branch, B., and Freed, W. (1977): "Bid–Asked Spreads on the AMEX and the Big Board," *The Journal of Finance*, 32:159–163.

Cheng, M., and Madhavan, A. (1994): "In Search of Liquidity: Block Trades in the Upstairs and Downstairs Markets," *NYSE Working Paper* 94-02.

Cohen, K., Maier, S., Schwartz, R., and Whitcomb, D. (1981): "Transaction Costs, Order Placement Strategy, and Existence of the Bid–Ask Spread," *Journal of Political Economics*, 89:287–305.

Eytan, T., Harpaz, G., and Krull, S. (1988): "The Pricing of Dollar Index Futures Contracts," *The Journal of Futures Markets*, 8:127–139.

Figlewski, S. (1987): "The Interaction Between Derivative Securities on Financial Instruments and the Underlying Cash Markets: An Overview," *Journal of Accounting, Auditing and Finance*, 2(3):299–318.

George T., and Longstaff, F. (1993): "Bid–Ask Spreads and Trading Activity in the S&P 100 Index Options Market," *Journal of Financial and Quantitative Analysis*, 28:381–397.

Jones, C., Kaul, G., and Lipson, M. (1994): "Transactions, Volume, and Volatility," *The Review of Financial Studies*, 7:631–651.

Laux, P. (1993): "Trade Sizes and Theories of the Bid–Ask Spread," *The Journal of Financial Research*, 14:237–249.

Massimb, M., and Phelps, B. (1994, January–February): "Electronic Trading, Market Structure and Liquidity," *Financial Analysts Journal*, 39–50.

O'Hara, M. (1995): *Market Microstructure Theory.* Oxford: Blackwell.

Redfield, C. (1986): "A Theoretical Analysis of the Volatility Premium in the Dollar Index Contract," *The Journal of Futures Markets*, 6:619–627.

Stoll, H. (1989): "Inferring the Components of the Bid–Ask Spread: Theory and Empirical Tests," *The Journal of Finance*, 44:115–134.
Vijh, A. (1990): "Liquidity of the CBOE Equity Options," *The Journal of Finance*, 45:1157–1179.