Combining Quote-Driven and Order-Driven Trading Systems in Next-Generation Stock Markets: An Experimental Investigation

Robert A. Schwartz and Bruce W. Weber

Department of Information, Operations and Management Sciences

Leonard N. Stern School of Business, New York University

44 West 4th Street, New York, NY 10012

bweber@stern.nyu.edu

IS-96-13

COMBINING QUOTE-DRIVEN AND ORDER-DRIVEN TRADING SYSTEMS IN NEXT-GENERATION STOCK MARKETS: AN EXPERIMENTAL INVESTIGATION

Abstract

We use computer-based simulations of stock market as a background environment for experimental tests of the integration of an order-driven trading system into a dealer/quote-driven market. Experimental subjects traded using a traditional dealer quote screen (such as Nasdaq in the U.S. or the London Stock Exchange's SEAQ), to which was added a public limit order facility. Data captured on subjects' trading decisions under different market structures revealed that: (1) When available, the limit order facility was used by the subjects, attracting some orders that would have otherwise gone to dealers, and reducing investor trading costs. (2) The relative use of market orders and limit orders was related to the bid-ask spread; wider spreads (higher cost of immediate trading) led subjects to enter fewer market orders. (3) Limit order use was reduced when the dealers were provided with an "informational advantage." (4) While the introduction of a limit order facility did <u>not</u> have a substantial effect on a dealer profit *margins*, dealers' activities as a percentage of total market volume declined. Overall, we find the simulation environment is a workable device for analyzing the effect of market design changes on trader behavior and market quality. It can provide solid guidance on market structure issues, such as how best to incorporate a limit order facility in a competing dealer market.

Combining Quote-Driven and Order-Driven Trading Systems in Next-Generation Stock Markets: An Experimental Investigation

Robert A. Schwartz
Professor of Finance and Economics
and Yamaichi Faculty Fellow

Bruce W. Weber
Assistant Professor of
Information Systems

Leonard N. Stern School of Business
New York University

July 1996

Abstract

We use computer-based simulations of a stock market as a background environment for experimental tests of the integration of an order-driven trading system into a dealer/quote-driven market. Experimental subjects traded using a traditional dealer quote screen (such as Nasdaq in the U.S. or the London Stock Exchange's SEAQ), to which was added a public limit order facility. Data captured on subjects' trading decisions under different market structures revealed that: (1) When available, the limit order facility was used by the subjects, attracting some orders that would have otherwise gone to dealers, and reducing investor trading costs. (2) The relative use of market orders and limit orders was related to the bid-ask spread; wider spreads (higher cost of immediate trading) led subjects to enter fewer market orders. (3) Limit order use was reduced when the dealers were provided with an "informational advantage." (4) While the introduction of a limit order facility did <u>not</u> have a substantial effect on dealer profit *margins*, dealers' activities as a percentage of total market volume declined. Overall, we find the simulation environment is a workable device for analyzing the effect of market design changes on trader behavior and market quality. It can provide solid guidance on market structure issues, such as how best to incorporate a limit order facility in a competing dealer market.

1. Introduction

Market structure research makes a strong distinction between financial markets that are *quote-driven* markets and those that are *order-driven* markets. Quote-driven markets rely on dealers to compete to provide the most attractive quotes to buy and to sell shares. Examples are the U.S. government bond market, the foreign exchange market, and the Nasdaq stock market. In an order-driven market, prices are established by limit orders to buy and sell that are submitted to a limit order book by traders and investors [9]. Examples are the Paris Bourse and the Toronto Stock Exchange. Both structures have their advantages, and neither can claim to be closer to the economic ideal of a perfect market. More advanced technology and growing

competition between markets, however, has given rise to opportunities for stock exchanges to combine the two trading mechanisms.

Several of the major stock exchanges that are dealer markets — including the London Stock Exchange (LSE) and Nasdaq in the U.S. — are in the midst of designing new trading systems and evaluating changes to their market structures. In the U.S., the Securities and Exchange Commission (SEC) has recommended that Nasdaq modify its systems to:

> "... display immediately customer limit orders ... This would improve competition among market participants by providing investors enhanced access to the market and, consistent with the statutory directive of achieving a national market system, would provide greater opportunities for investors' orders to interact with one another." **("Report Pursuant to Section 21(a) of the Securities Exchange Act of 1934 Regarding the NASD and the Nasdaq Market", SEC, August, 1996, p. 56)**

In London,

> "[Because] the current [LSE] structure does not easily allow trades to be executed on a 'patient' basis, without paying for the cost of risk capital ... the Exchange intends to provide a vibrant and attractive order book providing liquidity and immediacy for a significant proportion of trades on the Exchange." **(London Stock Exchange, "New Electronic Trading Services: Proposal for the Introduction of a Public Limit Order Book", May 1996)**

However, introducing a limit order mechanism into dealer markets is controversial and has many detractors. In the U.S., the NASD has

> "... considered proposing changes that would allow market orders to interact with limit orders between the inside spread, thereby increasing the number of trades executed inside the spread ... The NASD staff anticipated that many market makers would oppose this change ... " **(Report Pursuant to Section 21(a) of the Securities Exchange Act of 1934 Regarding the NASD and the Nasdaq Market, SEC, Appendix, p. 31)**

And, as reported in the London *Financial Times* ...

> "... there could be a combined quote- and order-driven system [for the LSE], a possibility known as the 'hybrid' approach. Marketmakers oppose a hybrid, arguing it would not provide high enough returns ... [However] large investors have made it clear they want some form of computerized order-matching system. This would keep costs down by cutting out the middleman." **(Gapper, J. "Options for order-driven system published, *Financial Times*, 13 January 1996, p. 4)**

"... many of the UK's marketmaking firms are resisting order-driven trading. ... Ministers of Parliament questioned whether a cartel of large market-making companies engineered Mr. Michael Lawrence's ousting [as LSE Chief Executive] in order to block the introduction of order-driven trading in place of their traditional quote-driven system." **("Backtracking Angers Brokers",** *Financial Times,* **24 April 1996)**

One contention is that a limit order facility will benefit the market by reducing costs and providing greater flexibility to investors. A counter-argument is that an order-driven facility will damage dealer firms, and eliminate their profits and willingness to provide liquidity to the market. It is also possible that a limit order facility in competition with dealers will simply not attract enough orders to be relevant and to have an effect on market quality.

To test these conjectures, we developed an experimental trading environment designed to facilitate the assessment of these market design modifications. This paper:[*]

- Shows how computer simulation can be used to evaluate the effects of market structure changes by capturing data on subjects' order placement behavior and on market quality.

- Details an experimental design in which live participants, playing the role of customer order entry firms, interact with the simulated market under three different market structure conditions.

  - The first scenario is a base case with a pure dealer market in which users can only submit market orders.
  - The second and third scenarios integrate a limit order book with the dealer quotes, and enable users to enter both market orders and limit orders.

- Presents summary data from the tests and analyzes the findings.

- Contains our conclusions concerning the effectiveness of simulation as an analytic tool for stock markets that are considering structural changes.

The paper is organized into six sections. Section 2 will discuss the simulation and the functions the computer performs in providing a controlled trading environment. Section 3 provides background on our experimental design, and states the hypotheses concerning

---

[*] Hugo Levecq, a Ph.D. student in the Information Systems Department, Stern School, NYU, tested the simulation software and conducted several of the trading experiments.

participant order placement behavior that we test. Section 4 explains how we conducted the experiments. Sixteen students and eight practitioners participated in one of five experiments. Each experiment consisted of three different scenarios. The first was a "pure" dealer market; all buying and selling was accomplished through dealer-intermediaries. The second scenario added a limit order book to the dealer market. In the third scenario, the dealers are given an informational advantage and trade more aggressively. Section 5 presents the results of the experiments, and the final section, Section 6, provides a conclusion and an outline of our future work. An appendix contains the structural details of the simulation model, and describes the layout of the computer screens, the generation of quotes and orders, dealer behavior, and the benchmark statistic used to measure participant performance.

2.   Using Simulation to Create a Realistic Market Environment

The objective of our computer simulation is to provide a backdrop for assessing the market decisions of live participants. Specifically, we have created a controlled background of public order flow and market maker quote setting that can interact with a user. This has enabled us to analyze a dealer market, and a combined dealer and order book market. The success of real-world markets in combining dealer and order book systems depends on adequate acceptance and use by traders, and the results from our experiment shed light on how such combined markets will be used.

Standard tests in experimental economics have not thus far used computers to create background order flow into which participants individually enter orders [11]. Our test environment does. This methodological difference has important implications for the development of experimental economics results for market design questions [10]. The computer-generated market provides the controlled background in which the live participants operate. The decisions made by participants in our experiments show that hybrid quote- and order-driven trading systems can succeed and provide market quality benefits.

Simulation is an alternative to analytical modeling and empirical tests of actual market data. To be used effectively, simulation models must reflect real world dynamics without being burdened by unnecessary real world detail [7]. A simulation model also requires a strong theoretical foundation. The advantage of simulation over theoretical modeling is that "theorizing" requires abstracting away from the very details of market structure that exchange officials and regulators wish to study. Consequently, theoretical modeling can give only limited insight into the effects of market design changes on the behavior of market participants. The

advantage of simulation vis-a-vis empirically testing of new market structures is that the simulated experiments can be run at much lower cost, and across a broader range of alternatives. — For our experiments, discrete event computer simulation accomplishes the following. It:

- Establishes the individual bid and ask quotes of five dealers. The quotes are dynamically updated as orders arrive, as trades are made, and as the dealers' inventories reach position limits (both long and short).

- Generates a public order flow that can (i) execute against the dealer quotes, and in two of the three test scenarios, can (ii) be placed on a public limit order book for later execution, or (iii) execute against the public limit order book.

- Gives the live participants instructions to buy a certain quantity of stock. Depending on the test scenario, the live participants can (i) execute these orders against dealer quotes, (ii) place them on a public limit order book, or (iii) execute them against the public limit order book.

- Maintains the screen which displays (i) the quotes of the five dealers, (ii) orders on the public limit order book, (iii) a time stamped record of all transaction sizes and prices for each trading session, and (iv) the value of the market index. (see Figure 1)

- Captures information concerning (i) the live participants' decisions and (ii) market quality measures such as bid-ask spreads. This information is assessed to determine the effect of market structure on the behavior of market participants.

In our person-machine interactive environment, the computer can generate orders from an unlimited number of "machine-resident" traders and investors. This enables us easily to satisfy the conditions for an active, competitive market. This facilitates assessing the relative performance of all participants in an experiment, and economizes on our use of live subjects.

Our use of simulation to observe participants operating under alternative market structures is analogous to the use of flight simulators to train airline pilots. Flight simulation reduces the cost of training new pilots. Securities market simulation can greatly reduce the cost of designing a new market structure. We caution, however, that a flight simulator can replicate atmospheric conditions far more realistically than our simulation can replicate the real intricacies of securities markets (Mother Nature is often easier to model than human nature).

Overall, our objective is not to describe subtleties in markets, but to gain insights that will be useful in analyzing a combined dealer and order book market structure, and for making

modifications. We report on tests run with students and industry executives as live users of the trading simulation.

Assumptions in the Simulated Market. We have kept the simulation as simple as possible. Key processes in the market are represented, but specific institutional arrangements need not be included. For example, interdealer trading is essential and is included in the simulation, but a separate interdealer trading facility (e.g., SelectNet, Instinet, or London's IDBs) need not be explicitly incorporated into the model. The model allows for trading just one asset; thus, arbitrage and "pairs" trading are not possible. The subjects' computers are <u>not</u> connected in a network, and each live participant interacts with the model and their screen individually. Thus, negotiation between two or more live participants is not possible. Further complexity along these and other lines could be added in the future, and the live participants could be given further training, if needed, before participating in more sophisticated experimentation.

The simulation reflects the runs and reversals in price changes that characterize the dynamic process of price adjustments in real world markets. These patterns give the live participants a rich environment to operate within, and in which to make decisions. The runs and reversal are introduced as follows. We specify an equilibrium (balance) price, P^*, that reflects the broad market's assessment of the value of the stock being traded. The live participants do not see P^*. Given P^* and the current level of the best bid and offer on the market, orders are generated by three types of machine traders: liquidity traders, informed traders, and momentum traders. (See Figure 2)

- Liquidity orders are obtained by drawing orders from a distribution located in relation to the market bid and ask quotes.

- Informed orders are obtained by identifying an equilibrium price, P^*, and by drawing orders from a tight distribution around P^* whenever P^* lies <u>outside</u> the bid-ask spread.

- Momentum orders are obtained by increasing the probability that the next order will be a buy (or a sell) whenever three or more buy orders (or sell orders), in part by chance, arrive in a row.

P^* follows a random walk jump process (see Graph 1). After the equilibrium value jumps randomly from one level to another (which it does from time-to-time), the orders of the informed traders pull the quotes and transaction prices to — and cause them to trend towards — the new level. Occasionally, market prices can also trend away from P^* because of the

orders of momentum traders. However, movements away from P* are unsustainable; eventually, price reverses and heads back towards P*. The appendix provides additional details on the model.

3.   Experimental Design and Hypotheses

In the tests, subjects played the simulation game in groups. Each laboratory session lasted two hours. Roughly one hour was devoted to explaining the simulation and to allowing the subjects to practice entering orders in the market, and one hour was devoted to formal testing. The experiment had three different simulation environments, each of which took about twenty minutes to complete.

Subjects' Decisions. In the simulation, live participants played the role of public order entry firms, and at different points during the experiment received from the computer "instructions" to buy a quantity of shares. Giving only buy orders was a way of simplifying the subjects' decisions and controlling complexity. The quantity a user was given was large enough that it could <u>not</u> be fully executed in one trade; this required subjects to exercise discretion in when to submit market and limit orders. Each test run covered three simulated days of trading each of which lasted a simulated 6 1/2 hours (from 9:30 am to 4:00 pm). In a typical simulated day, a subject received two buy "instructions." By entering his or her orders, and from trading resulting from the computer-generated orders, about 100 trades occurred a day. The subjects were participants in about ten percent of the trading activity. At the end of a simulation run, any unfilled buy instructions were automatically executed by the computer. This treatment of otherwise unfilled orders encouraged participants to obtain fills before the end of a trading session, but did not penalize them unduly, for this could have undermined their willingness to trade patiently.

Incentives. The participants in a laboratory session were ranked by a performance measure, and each was given a payment in dollars at the end of the session that depended on his or her relative ranking for the session.[1] The buy instructions given would create "market impact" and drive up prices if submitted as market orders in large pieces. The performance measure for each participant was constructed as follows. Each time the participant brings all or part of an order to the market, the computer assesses the average share price that would be obtained if the order

---

[1] Cash payments ranged from $20 to $30 for two hours. Payments were not made in the industry practitioner experimental session. Instead, results were tabulated and displayed.

were brought to the market in its entirety and executed at one or more prices. We refer to this price as the "benchmark price." The participant's "profit" for any trade is the benchmark price minus the actual purchase price, times the number of shares bought. The participant's performance measure is the cumulative profits over all trades. After each experimental session, participants in the session were ranked by their performance measure and paid accordingly. Knowing that their payouts depended on their performance induced the subjects to perform well.

Careful order handling enabled subjects to increase their performance measure by buying at lower prices. Orders can be "worked" by breaking them up into smaller pieces, and by placing limit orders [4][5]. However, if a limit order executes because the balance price (P^*) has decreased, the benchmark price will be marked down accordingly and the subject receives a relatively poor profit score for the trade. In effect, the profit score on a trade relects "ex-post regret" of trading. If a trader buys, and subsequently realizes he could have bought for less later, the benchmark will fall and profit will be small, or negative. If the trader buys and prices rise, the benchmark and profits will be greater. This provides realistic incentives to trade strategically and cost-effectively. In total, overly eager order handling is likely to result in low scores because of market impact, and excessively patient trading commonly resulted in poor scores because of undesirable limit order executions following informational change.

FIGURE 3: Market Order and Limit Order Illustration

Assume that in the combined quote- and order-driven market structure, the expressions of trading interest reflected below exist. The stock currently has a bid quote of 24 1/4, which reflects the limit order to buy 15. Three dealers have a bid quote of 24 1/8, which is inferior to the limit order buyer willing to pay 24 1/4. The best offer to sell is at 24 1/2 by Dealers #1 and #3. The lowest limit order to sell is at 24 5/8, which is inferior to the dealers' ask quotes. A dealer's quote is good for up to 25 units. However, after a trade occurs, the dealers have the option to move their quotes. This implies that it may not be possible to buy 50 at the current offer price of 24 1/2.

COMBINED DEALER QUOTES AND LIMIT ORDER BOOK MARKET STRUCTURE

Limit Orders to Buy (No. of Shares)	Dealers Bid	Price	Dealers Offer	Limit Orders to Sell (No. of Shares)
		$24 3/4	--	12
		24 5/8	#2,#4,#5	10
		24 1/2	#1,#3	
		24 3/8		
15		24 1/4		
6	#2,#4,#5	24 1/8		
22	#1,#3	24		

LIMIT ORDER STRATEGY

To execute an order to buy, an investor could place a limit buy order (a bid) at 24 3/8 or less. If entered at 24 3/8, it does not immediately execute, but supersedes the bid of 24 1/4. The hope is that a market sell order arrives and executes against the limit buy order at 24 3/8, or that a dealer that has become long the stock sells to the limit order.

MARKET ORDER STRATEGY

Under the same circumstances, an investor using a market order can buy immediately at the offer price of 24 1/2.

The limit order trader can avoid the cost of immediate execution ($1/4 in this example). Limit orders however face the risk of not executing, and the risk that the price will move adversely while they are on the book. For instance, the offer price could move higher without the limit order at 24 3/8 executing. Better bid prices could be entered, and a buyer that initially attempted unsuccessfully to use a limit order may end up facing regret and paying more than 24 1/2 for the stock.

Empirical data on limit order and market order use are limited. The New York Stock Exchange (NYSE) market is based on a limit order book and a single dealer, called a specialist, in each stock. Although not directly comparable to the competing dealer and order book market considered here, NYSE data indicates that about half of all orders are limit orders, and 73

percent of the *shares* in orders submitted are limit orders. However, just 44 percent of the shares submitted as limit orders eventually execute. [8]

A combined trading system will be viable only if its component trading mechanisms are adequately utilized. Of fundamental interest in our experiments is the extent to which the live participants actually use limit orders. The use of limit orders can be measured in two ways: the *level* (the actual frequency with which limit orders are placed) and the *change of the level* (how the frequency of use responds to a change in the simulation structure). The level is easier to measure, but may be overly dependent on simulation parameters that are somewhat arbitrary. The change in the level is difficult to capture but is more meaningful. We place primary emphasis on the *change of the level*. We hypothesize that limit order use and its changes will depend on conditions in the market.

The experiments enable us to test three hypotheses:

H1) Live participants will use more limit orders when market bid-ask spreads are wider.

The bid-ask spread is the price of immediacy for investors, and at higher values it will encourage limit order use. The greater the price of immediacy, the more apt a public participant should be to enter limit orders.

H2) Participants will use limit orders <u>less</u> frequently when dealers have an informational advantage regarding the location of P*.

In actual markets, some public customers are at times at an informational disadvantage vis-a-vis dealers who are professionals, and whose contact with a broad array of customers (including institutional clients) is frequent. Dealers may capitalize on any informational advantage via the placement of their quotes and by "picking off" mispriced limit orders. We hypothesize that, if the dealers do so, the live participants will be discouraged from using limit orders. The dealers were given more of an informational advantage in Scenario 3 than in Scenario 2.

H3) The limit order book will reduce dealer profits by "disintermediating" some of the trading volume, but the order-driven facility will give dealers a new way to lay off positions and risk.

A limit order book can provide liquidity to dealers as well as other investors. Since market makers avoid large long or short positions, the presence of limit orders can help dealers to control their risk and reduce the size of their losses from adverse price movements. The reduction in dealer profits from a limit order book could be offset by the decrease in their risks.

To test the hypotheses above, we vary market conditions by defining three alternative scenarios, and by assessing the participants' differential use of limit orders in these three

environments.   The first environment includes a dealer screen only (that is, there is no limit order facility).   The second and third scenarios, which include a limit order facility, are differentiated by (i) the average price of dealer-provided immediacy, and (ii) the dealers' ability to have superior knowledge of the equilibrium (balance) price, P^*.

Subjects were put into three scenarios, detailed below:

Scenario #1	Uninformed dealers; Dealers have no special information concerning the location of P^* and its changes. Dealer quotes are relatively wide and hence the price of dealer provided immediacy is relatively high. In Scenario 1, all trading takes place via market orders trading at dealers' quoted bid and ask prices.
Scenario #2	Uninformed dealers with quote-setting as in #1. The limit order book is integrated into the market, and price priority is enforced so that a limit order at a better price than the best dealer quote will trade first. If the best dealer quote and the best limit order at tied at the same price, then the quote or order placed earlier has priority.
Scenario #3	Dealers are informed about half of P^* changes giving them better information concerning the location of P^* than have the live participants. Dealer quotes can be narrowed and the price of dealer-provided immediacy made relatively low without losses. Dealers will reduce the bid-ask spread from \$3/4 in Scenarios 1 and 2, to \$1/2.

The informational advantage of dealers in Scenario 3 was controlled as followed.   As noted in Section 2, the simulation includes an underlying equilibrium (or balance) price, P^*. Displayed quotes and transaction prices <u>can</u> differ from the P^* level after a change in P^* or when temporarily pushed away from P^* by momentum traders or chance one-sided liquidity order flow.   Within this structure, we can give the dealers a controllable informational advantage vis-a-vis the live participant by conditioning dealer trading and quote-adjustment behavior, in part, on their ability to know P^*.   Operationally this works as follows.   In Scenario 1 and 2, we gave the dealers no special information about P^*, but in Scenario 3 we inform the dealers of P^*'s new value <u>half</u> of the time that P^* changes.   When the dealers know the new value of P^* and it is <u>below</u> public buy limit orders (or <u>above</u> public sell limit orders), the dealers sell to the overpriced public bids with their own sell orders (or buy from the underpriced public offers with their own buy orders); in the process, they adjust their bid and offer quotes to straddle P^*.   We have established that the profits of the machine dealers are controlled via the strength of the information signal given to them (see Section 5 and Table 5 in particular).

4.   Conduct of the Experiments

Our tests comprised four laboratory sessions conducted with four students each (16 subjects in total), and one with 8 industry participants (an additional 8 subjects). Eight of the subjects were Ph.D. students at NYU's Stern School of Business, and eight were graduate students at the Wharton School. Each experimental session lasted two hours. Roughly one hour was spent introducing the students to the experiment and the simulation, and to letting them practice entering orders into the simulated market. One hour was then devoted to running the formal tests.

Each formal test period was divided into three 20-minute windows. The subjects completed one scenario in each 20-minute interval. The first 20-minute period included the dealer only environment, while a limit order facility was included in the second and third periods. In the second 20-minute period we used the "uninformed dealer" environment, and in the third 20-minute period we used the "informed dealer" environment.

Fundamental parametric values were held constant across all simulation runs, and the tests were structured so that valid, ceteris paribus contrasts could be made. However, it was necessary to start each new simulation runs for each of the 20-minute periods with a different random number seed to keep the subjects ignorant about what the new price changes would be. Changing the random number seed, however, results in a different sample of returns, and some random number seeds produce more volatile returns distributions than others. To mitigate any bias that this might cause, the random number seeds were rotated between the three 20-minute windows for the four different experimental groups. Our ordering of Scenario 2 and 3 was also reversed for about half of the subjects to offset any ordering or learning effects.

For each participant in each experimental group, the performance measure was tracked over the three different environments to compute the final score for the participant for the simulation session. The dollar payments for the participants in each session were then determined with the subjects achieving the best score receiving the largest payments.

A simulation experiment was conducted with eight industry executives. After a training period, all eight participated in the dealer-only (Scenario 1) environment. In the next round, four participants used the limit order scenario with uninformed dealers (Scenario 2) and four used the limit order scenario with informed dealers (Scenario 3).

5.   Test Results

Tables 1a and 1b present a breakdown of the subjects' orders by type of order and Scenario: Scenario 1 is dealers only, Scenario 2 is uninformed dealer and limit order facility, and Scenario 3 is informed dealers and limit order facility.   The interesting observation concerning Tables 1a and 1b is that the live subjects used limit orders about as much as they used market orders in Scenarios 2 and 3 (the only scenarios where they are allowed).   Limit orders ranged from 46% to 57% of the total orders submitted.   Subjects found that both limit orders and market orders had value to them in executing their buy instructions.   Because not all limit orders execute, market orders were used for over half of the shares actually bought.   The data from the industry subjects are only reported separately in Table 1a, but because meaningful differences did not appear, we report all subsequent data in aggregate.

TABLES 1A AND 1B
1A) BREAKDOWN OF SUBJECTS' ORDERS BY TYPE AND SCENARIO

No. of Shares	Students (n=16)			Industry Practitioners (n=8)		
	Scen. 1	Scen. 2	Scen. 3	Scen. 1	Scen. 2	Scen. 3
Market Orders	4,917	3,148	3,003	2,618	1,530	1,748
Limit Orders Placed	-0-	4,123	2,523	-0-	1,736	1,748
Limit Orders Executed	-0-	1,763	1,345	-0-	1,032	870
Limit Orders as % of Total Placed	NA	56.7%	45.7%	NA	53.2%	50.0%
Limit Order as % of Total Executed	NA	35.9%	30.9%	NA	40.3%	33.2%

TABLE 1B) AGGREGATED DATA

No. of Shares	All Subjects (n=24)		
	Scen. 1	Scen. 2	Scen. 3
Market Orders	7,535	4,678	4,751
Limit Orders Placed	-0-	5,859	4,271
Limit Orders Executed	-0-	2,795	2,215
Limit Orders as % of Total Placed	NA	56.1%	46.7%
Limit Orders as % of Total Executed	NA	36.8%	31.5%
Limit Order Execution Rate	NA	45.7%	52.4%

Table 2 shows the frequency with which the various spread values occurred, and the frequency with which *market orders* are entered at different values of the spread. Examining columns (3), (6), and (9) in Table 2, we see that the use of market orders diminishes when the spread is at its widest level (e.g., 5/8 and 3/4). E.g., in Scenario 1 the bid-ask spread was 5/8 or 3/4 36.9% of the time, but only 23.4% of the total market orders were placed when the spread was that wide. The dealers' spread was reduced from 3/4 to 1/2 in Scenario 3, and as a result the maximum market spread is 1/2, and no data cells are needed in the lower right of the table.

TABLE 2

SUBJECTS' MARKET ORDERS PLACED AT DIFFERENT BID-ASK SPREADS

Spread (in $)	Scenario 1			Scenario 2			Scenario 3		
	(1) Spread Freq.	(2) Market Orders Entered	(3) (2)÷(1)	(4) Spread Freq.	(5) Market Orders Entered	(6) (5)÷(4)	(7) Spread Freq.	(8) Market Orders Entered	(9) (8)÷(7)
1/8 & 1/4	24.3%	26.7%	1.10	39.2%	44.9%	1.15	71.7%	72.7%	1.01
3/8 & 1/2	38.8%	49.9%	1.28	41.9%	38.5%	0.92	28.3%	27.3%	0.96
5/8 & 3/4	36.9%	23.4%	0.63	18.9%	16.6%	0.88			

Table 3 shows the frequency with which *limit orders* are entered at different values of the spread. Examining columns (3) and (6) in Table 3, we see that limit orders are used at all values of the spread, but they are used least when the spread is narrowest (1/8 and 1/4).

TABLE 3
SUBJECTS' LIMIT ORDERS PLACED AT DIFFERENT SPREADS

Bid-Ask Spread	Scenario 2			Scenario 3		
	(1) Spread Frequency	(2) % Limit Orders Entered at this Spread	(3) (2) ÷ (1)	(4) Spread Frequency	(5) % Limit Orders Entered at this Spread	(6) (5) ÷ (4)
1/8 & 1/4	39.1%	31.1%	0.79	71.7%	69.5%	0.97
3/8 & 1/2	41.9%	51.8%	1.24	28.3%	30.5%	1.08
5/8 & 3/4	18.9%	17.1%	0.90			

Table 4 focuses on the use of limit orders versus market orders. A meaningful pattern appears to exist between the relative importance of these two types of orders and spread size. When spread size is greater, subjects used proportionately more limit orders, which is sensible, and supports our first hypothesis. Overall, limit orders are used less than market orders in Scenario 3 than in Scenario 2 for a given size of spread. The diminished attractiveness of limit order use supports the second hypothesis: a narrower bid-ask spread and the fact that the (partially) informed dealers will "pick off" limit orders in Scenario 3 (following a change in the balance price, P*) reduces the incentive to use limit orders. This holds even when controlling for the narrower spreads in Scenario 3. Comparing Scenario 2 and 3, we find that the order book will be better utilized the smaller the information advantage enjoyed by dealers relative to public investors.

TABLE 4
SUBJECTS' LIMIT ORDERS PLACED RELATIVE TO SUBJECTS' MARKET ORDERS

Spread	Scenario 2 Limit ÷ Market	Scenario 3 Limit ÷ Market
All Values	1.28	0.90
1/8 & 1/4	0.88	0.86
3/8 & 1/2	1.72	1.00
5/8 & 3/4	1.31	

Based on the results shown in Tables 1−4, we conclude that the subjects are responding rationally to the existence of the limit order book. They use limit orders when allowed to do so. They cut back on their use of market orders at very wide spreads. Use of limit orders is positively related to the costs of immediacy as they vary during a simulation run in the experimental environment. Use of limit orders increases as the spread widens from 1/8 and 1/4 to 3/8 and 1/2. And, the subjects cut back on their use of limit orders when the limit orders are at risk of being picked off by better informed dealers (Scenario 3 compared with Scenario 2).

Table 5 presents various summary statistics for the three scenarios. Comparing Scenario 1 to the others, the introduction of the limit order facility narrows the bid-ask spread, reducing trading costs for investors. The dealers' share of trading volume falls from 100% in Scenario 1 to about 64% in Scenario 2. This indicates that investor limit orders are able to provide liquidity to arriving market orders over a third of the time. In the hybrid market, the need for dealer intermediaries is reduced, but not eliminated. In Scenario 3, when dealers are informed and tighten the bid-ask spreads that they each quote to 1/2 from 3/4, they regain market share, and are counterparties to over three-quarters of trading volume.

In Scenarios 1 and 2, dealer profits were negative on average. Adjusting the simulation parameters, such as lowering the P* volatility, would be likely to raise dealer profits. However, the more important result is that dealer profits in basis points are nearly unchanged from Scenario 1 to Scenario 2. Only part of the third hypothesis is supported; contrary to our expectations, dealer margins were not reduced. Moreover, the dealers' average position size is reduced about 25%, and hence dealer risk is *lower* in the presence of an active order book. In Scenario 3, when dealers are informed about P* changes, profits are positive, but this is

certainly due to the ability to quote the "right" price and avoid losses from trading with informed investors.

The implications of steady dealer margins but a reduced share of trading occurring with the dealers, is that <u>absolute</u> profit levels of dealers will fall in a combined market structure <u>unless</u> overall order flow and trading volume increases. Another effect of the order-driven facility is a reduction in inter-dealer trading volumes.

TABLE 5
MARKET AGGREGATES

	Scenario 1	Scenario 2	Scenario 3
Average Bid-Ask Spread	50.4¢	38.0¢	26.0¢
% of Price	2.0%	1.6%	1.1%
Percentage of Total Trading Volume with a Dealer counterparty	100.0%	63.7%	76.8%
Average Dealer Profits in $	−229.9	−137.0	45.6
in Basis Points	−2.07	−2.17	0.63
Average Absolute Dealer Inventory Position	19.9	14.3	19.3
Inter-Dealer Trading Volume as % of Total Volume	33.2%	16.3%	16.7%

6.   Conclusions and Future Research

The market simulation experiment provided a trading environment that we use to test the complex characteristics of a hybrid dealer screen and a public limit order facility. Critical to its success, the simulation was kept simple enough to be workable for the live subjects, but also rich enough to generate useful results. Twenty-four live subjects had no difficulty learning how

to interact with the simulation. The subjects appeared to find the exercise interesting and enjoyable.

The simulation runs generated a considerable amount of data which we captured. The important results can be highlighted:

- The subjects used the limit order facility actively, and dealers saw some of the trading activity bypass them as investor limit orders met investor market orders in the order book.

- The subjects use of the order book reduced the market's bid-ask spread.

- The relative use of limit orders and market orders was correlated with spread width in a reasonable fashion; wider spreads led to greater limit order submission.

- The introduction of a limit order facility did <u>not</u> reduce dealer profits in basis points, but lowered their percentage of trading volume. This implies that in a hybrid market, dealer profits will decrease in total dollar amount unless overall trading activity increases to offset the order book's volumes.

Our tests indicate that the market simulation environment is a workable device for experimental analysis of the effect of market design changes on trader behavior and market performance. We anticipate that, when further developed, it will give strong guidance on market structure issues such as how best to incorporate a limit order facility into a competing dealer environment, or how to add a periodic call market. We expect to make the following enhancements in the conduct of the experiments and in the structure of the simulation.

1. <u>Conduct of the experiments:</u> The current version of the simulation was kept as simple as possible so as to run valid experiments with inexperienced subjects after just one hour of training. We suggest that far more could be accomplished if the subjects are given two to three hours of training.

2. <u>The inclusion of buy and sell orders in the order entry firm environment:</u> In the current version of the simulation, the live subjects in the order entry firm environment were given buy orders only so as to keep their role as order entry firms as simple as possible. With more training, the subjects could be given both buy and sell orders. This would make their task more realistic and more challenging. It would also enable a simpler performance

measure to be used. We would balance the buy and sell orders given to each subject so that, for each, the sum of all shares bought equaled the sum of all shares sold over the course of a simulation run. We would then assess each participant's performance by a simple profits calculation: the sum of all positive cash flows from sales, less the sum of all negative cash flows from purchases.

3. **Live Subjects Should Play the Role of Dealers:** In the current tests, the live subjects played the role of order entry firms only. Tests should also be run with the subjects playing the role of dealers.

To summarize, we use computer-based simulations of a stock market as a background for examining the effects of a combining order-driven trading system into a dealer/quote-driven market. The results indicate that a limit order facility will be used and will reduce investor trading costs, and an active limit order facility, in and of itself, will not eliminate the profitability of dealers' activities. We found that interactive simulation is a workable device for analyzing the effect of market design changes on trader behavior and market quality. It can provide helpful direction on market structure issues to securities exchange such as Nasdaq and the LSE that are considering combining order-driven and quote-driven trading systems.

7. Bibliography

[1]   Clemons, E. and Weber, B. "London's Big Bang: A Case Study of Information Technology, Competitive Impact, and Organizational Change", *Journal of Management Information Systems*, Vol. 6, No. 4, Spring 1990, pp. 41-60.

[2]   Clemons, E. and Weber, B. "Alternative Securities Trading Systems: Tests and Regulatory Implications of the Adoption of Technology", *Information Systems Research*, Vol. 7, No. 2, June 1996, pp. 163-188.

[3]   Cohen, K., Conroy, R., and Maier, S. "Order Flow and the Quality of the Market", in Amihud, Y, Ho, T., and Schwartz, R. (eds.) *Market Making and the Changing Structure of the Securities Industry*, Lexington Books, 1985, pp. 93-109.

[4]   Conroy, R. and Winkler, R. "Informational Differences Between Limit and Market Orders for a Market Maker", *Journal of Financial and Quantitative Analysis*, December 1981, pp. 703-724.

[5]   Handa, P. and Schwartz, R. "Limit Order Trading", *Journal of Finance*, forthcoming, 1997.

[6]   Ho, T. and Macris R. "Dealer Market Structure and Performance", in Amihud, Y, Ho, T., and Schwartz, R. *Market Making and the Changing Structure of the Securities Industry*, Lexington Books, 1985.

[7]   Law, A. and Kelton, W.D., *Simulation Modeling and Analysis*, McGraw-Hill, New York, Second Edition, 1991.

[8]   New York Stock Exchange, *Fact Books*, and correspondence with Research Department, 1989-1995.

[9]   Schwartz, R. , *"Reshaping the Equities Markets: A Guide for the 1990s"*, Business One Irwin, 1993.

[10]   Smith, V. "Microeconomic Systems as an Experimental Science", *American Economic Review Proceedings*, (1982), pp. 923-955.

[11]   Smith, V. and Williams, A. "Experimental Market Economics", *Scientific American*, December 1992, pp. 116-121.

[12]   Mendelson, H. "Consolidation, Fragmentation, and Market Performance", *Journal of Financial and Quantitative Analysis*, Vol. 22, June 1987, pp. 189-207.

[13]   Schwartz, R. *Reshaping the Equity Markets: A Guide for the 1990s*, Harper Business, New York, 1991.

[14]   Weber, B. "Bypass Trading and Market Quality in Electronic Securities Exchanges", *Journal of Organizational Computing*, Vol. 5, No. 3, pp. 327-353 (1995).

[15]   Weber, B. "Assessing Alternative Market Structures Using Simulation Modeling", *Global Equity Markets: Technological, Competitive, and Regulatory Challenges*, R. Schwartz (ed.), Business One Irwin, 1994.

Appendix: Details of the Simulation Model

In this section, we describe the major features of the simulation model and the user instructions. The simulation involves one live person interacting with a computer-driven "market background." There is one risky security, and the live user plays the role of a public order entry firm.

The computer screen has two parts: (1) A dealer montage that shows the quotes of five dealers. This screen is similar to, but not a replication of the Nasdaq screen. The display is simplified to economize on information load and to facilitate comprehension by the live participants. (2) A limit order file that resembles an open order book-type facility. The screens

are shown in Figure 1.

Quotes and orders are generated by the computer. The computer driven orders arrive stochastically, according to a poisson order arrival process. Order interarrival time is a parameter in the model. We set the order arrival rate to 30 orders per hour. Dealer quote setting is based on heuristic rules built into the model that result in market makers adjusting their bids and offers according to their inventory and position limits. (See Graph 1)

In the dealer-only simulation (Scenario 1), all public orders are market orders. Market orders trade immediately at the best bid quote (if it is a market sell order) or at the best offer quote (if it is a market buy order). In the dealer and limit order file model, a limit order that arrives in the market indicates the maximum price at which a purchase is to be made, or the lowest price at which a sale is to be made. The procedures used for generating machine orders are discussed in further detail below.

All orders, whether generated by the computer or by the live participant, are integer values of from one to twenty-five units of the security. The initial value of the equilibrium price is $24.00, and moves in $1/8 (12.5 cents) increments.[2]

Order size. On the basis of its fit with empirical data, the Beta distribution was selected for order sizes in the simulation (see Graph 2). The Beta distribution is useful for modeling activity completion times and quantity demands that are bounded on a finite interval. In the simulation, order sizes are distributed as a discrete linear transformation of a Beta(a,b) random variable, resulting in integers between 1 and 25.

Order preferencing. Public market orders that execute against dealer quotes are directed to specific dealers by the following procedure. To simplify the discussion, consider the arrival of a market sell order (the handling of a market buy order is symmetrical). If _all_ of the dealers are making the best bid price, an incoming sell order is allocated to a dealer on the basis of market share. For instance, a dealer with a 20 percent market share has a one-fifth chance of getting the next incoming order. If only some dealers are making the best price, the order may be preferenced to a dealer who is not setting the best bid or offer. In this case, the order will go to a dealer based on the following market share adjustment: the market shares of all dealers _making the best quote_ are multiplied by 1.3 and added to the market shares of the other dealers. This inflated sum is used to scale back the individual market shares of all dealers so that the

---

[2] If one unit is viewed as representing 300 shares, the size and price values would reflect a convenient normalization which is reasonably realistic for the Nasdaq market (trade sizes of $4,000 to $100,000 account for about 70% percent of total value of Nasdaq trades). Beyond 25 units, we assume the trade would be handled as a negotiated block trade, or would arrive in the market in a sequence of smaller pieces.

final sum is unity.[3] This procedure leads to a smaller probability of receiving a customer order when a dealer is not making the inside quote. Orders that are preferenced to a dealer not making the best price are nevertheless executed at the inside market quote (i.e., we incorporate quote matching).

Order flow. Three sources of machine driven orders are included in the simulation model: the orders of liquidity traders, informational traders, and momentum traders. The liquidity traders provide a base level of trading volume. Orders from liquidity traders arrive stochastically, according to a poisson order arrival process. An individual liquidity order is either a buy order or a sell order with equal probability.

Each liquidity order is drawn from a double triangular distribution that has one probability peak at the market bid and a second probability peak at the market ask (see Graph 3). This probability structure assures the preservation of a non-trivial bid-ask spread in the market. By varying parameters of the double triangular distribution, we are able to control, stochastically, the size of the bid-ask spread. Both market and limit orders are drawn from the distribution. For a buy order, a draw from the distribution at or above the market ask is a market buy order; a draw from the distribution below the ask is a limit order, and it is placed on the limit order file. Sell orders are handled symmetrically.

As the market bid and ask quotes change with the course of the simulation, the location of the double triangular distribution shifts with them.[4] If the double triangular distribution were the only source of order flow, transaction prices generated by the simulation would follow a random walk with a bid-ask bounce. Over time, with a random walk, prices will drift from one level to another; the price drift would be consistent with random informational change and instantaneous, perfect adjustment to that change in the market place. The environment would be what financial economists refer to as "informationally efficient."

Informed traders generate one-sided (buy or sell) volume, above the base level, whenever the market makers' quotes do not straddle an underlying equilibrium price, P*, as discussed above. P* is not seen by the live traders. Changes in P* create profit opportunities for the informed traders. The trades of the informed participants cause the market bid and ask quotes to move toward, and eventually to straddle, P*. One-sided orders from informed traders are generated according to the Poisson order arrival process. The time between information change

---

[3] For instance, assume four dealers, that two of the four have 20 percent market shares and are making the best quote, and that the other two dealers have a 30 percent market shares are *not* making the best quote. The probability that one of the dealers on the quote will get the trade is 1.3*20/(2*1.3*20+2*30) = 23.2% > 20.0%. The probability of one of the other dealers receiving the order is 30/112 = 26.8% < 30.0%. Note that 2(23.2) + 2(26.8) = 100.

[4] Specifically, two double triangular distributions are used: one for the generation of buy orders and another, its mirror image, for the generation of sell orders.

is exponentially distributed with a mean of 4 hours.[5]

The expected arrival rate of informed orders depends on the relationship between market quotes and P*. If the bid and offer quotes straddle P* — e.g., P*=$24 5/8, and the bid is $24 1/2 and the offer is $24 7/8 — there is no informed order flow. When P* is above (below) the bid-offer range, buy (sell) orders are generated by a Poisson arrival process that reflects the rate at which "informed" machine traders become aware of the relatively high (low) value of P*. We control the arrival rate of informed orders by increasing the base order arrival rate for buy orders by 75% (when P* is above the ask) and for sell orders by 75% (when P* is below the bid). The increased order flow from informed traders provides the link that keeps the equilibrium value (P*) loosely aligned with market maker quotes. Adjustment of the quotes to change in P* is not immediate, and volume tends to increase when the equilibrium price and the quotes diverge.

Momentum traders are a class of investors who trade whenever they sense that trending in the market has been caused by a change of P*. Momentum traders enter orders over and above the base volume, whenever there is a sequence of three or more trades or quote improvements on one side of the market. If they follow sufficiently closely on the heels of informed traders, this can be a profitable short term strategy. The strategy is unlikely to be profitable if the trend was merely a chance arrival of three liquidity orders on the same side of the market.

Market makers. Each of the five dealers is assigned a market share and a position limit. In the course of the simulation, when a dealer's inventory position increases beyond a position limit, he or she takes a number of steps to reduce the size of the position. The first step taken is to raise or lower his or her quotes. When the position becomes twice the position limit, the dealer initiates trades with other dealers, hitting their bids or lifting their offers. A dealer also initiates trades with other dealer(s) whenever raising or lowering a quote would lead to a "locked" market (the bid and offer quotes are equal).

Instructions to the user are generated by the computer. Once an instruction to buy is given, orders can also be entered by the live participant at his/her discretion.

Market rules/Trade priorities. When the limit order book is available, the limit orders are

---

[5] P*, when it changes, follows a random walk without return drift. To assure nonnegative prices, P* is determined using a log-normal distribution. That is,

$$\ln P_t^* = \ln P_{t-T}^* + e_t \quad \text{where, } e_t \sim N(0,TV)$$

where t is an index on time, T is the elapsed time since the last price change, and V is a variance parameter.

"protected" from trade-throughs. For instance, a limit order to buy at 23 1/4 must be filled before any trades can occur at 23 1/8. Public limit orders are generally executed according to strict time and price priority vis-a-vis other public limit orders and dealer quotes. If a large order executes against two or more limit orders place at different prices, the entire block is executed at the lowest bid that it executes against (for a sell) or at the highest ask that it executes against (for a buy). From the user's perspective, there are no negotiations; all transactions are made at posted quotes.

Interdealer trading is an important component in dealer markets, as it is the means by which public buying/selling pressure on one dealer is transmitted to other dealers. Real world facilities for inter-dealer trading include Instinet (in the U.S.) and the IDBs (in the U.K.). It is not necessary for our purposes to include a specific facility in the simulation, and doing so would only complicate the screen for the live participants. We simply trigger direct dealer-to-dealer trading whenever the inventory adjustment rules call for one dealer to hit or lift the quote of another dealer. All *initiated* inter-dealer trades are of size 10 units. All inter-dealer orders *received by another dealer* are of size 1 to 10 units, depending on the size of public limit orders on the book that have priority. Inter-dealer trades are not preferenced, meaning that only dealers quoting the best price will receive the next order.

Order book/Dealer interactions. Our simulation allows for interaction between the dealer screen and the limit order screen. Dealers can hit public buy orders and lift public offers to sell on the limit order book. The dealers also take public limit order prices into account when setting their own quotes. At all times, strict time and price priority is given to orders on the public limit order file vis-a-vis each other. Time priority also determines the sequence of order execution when a public limit order is tied in price with a dealer quote. A large incoming market order can execute partially against the limit order file and partially against a dealer quote if the incoming market order is larger than the sum of the public limit orders that have priority over the dealer quotes. When this occurs, the entire block executes at a single price (the highest price reached for a buy, or the lowest price reached for a sell). Outcomes and market quality may be highly sensitive to rule such as these. For instance, the NYSE imposes public order priority, and all limit orders at a price will execute before a specialist is allowed to buy or sell at that price as a dealer.

FIGURE 1
USER TRADING SCREENS FROM THE SIMULATION EXPERIMENT

Initial menu
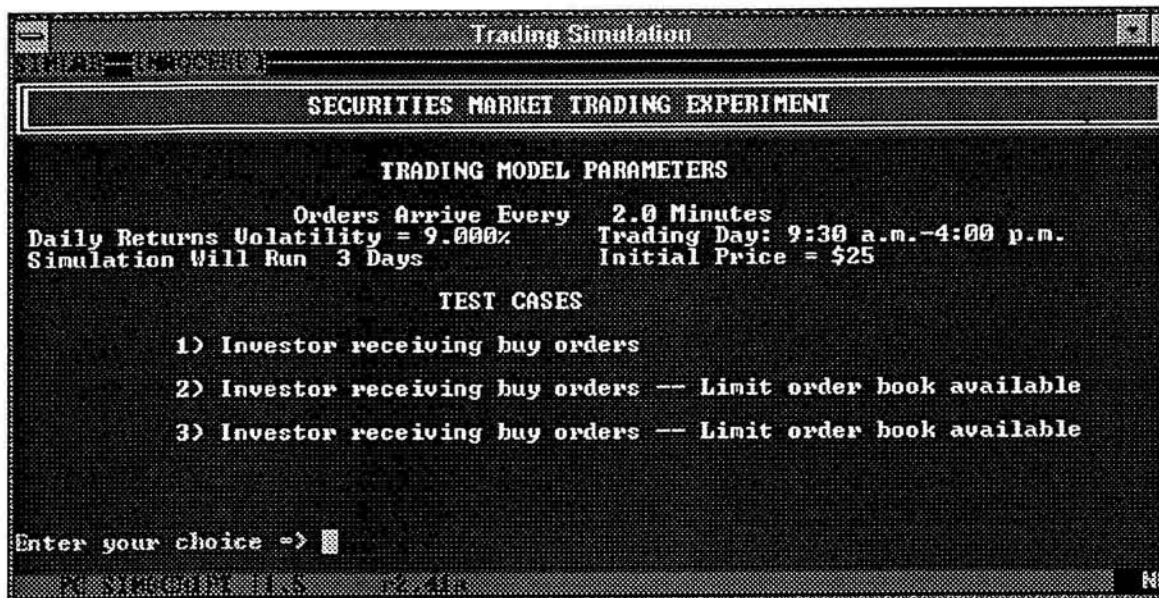
Beginning of Scenario 1 (dealer-only) experimental run. Current best bid is $25 and best offer is $25 4/8. Ticker at top of screen indicates two trades have occurred at $25 User has just received instructions to buy 56 units. Benchmark price that user seeks to buy as far below as possible is $25 7/8.
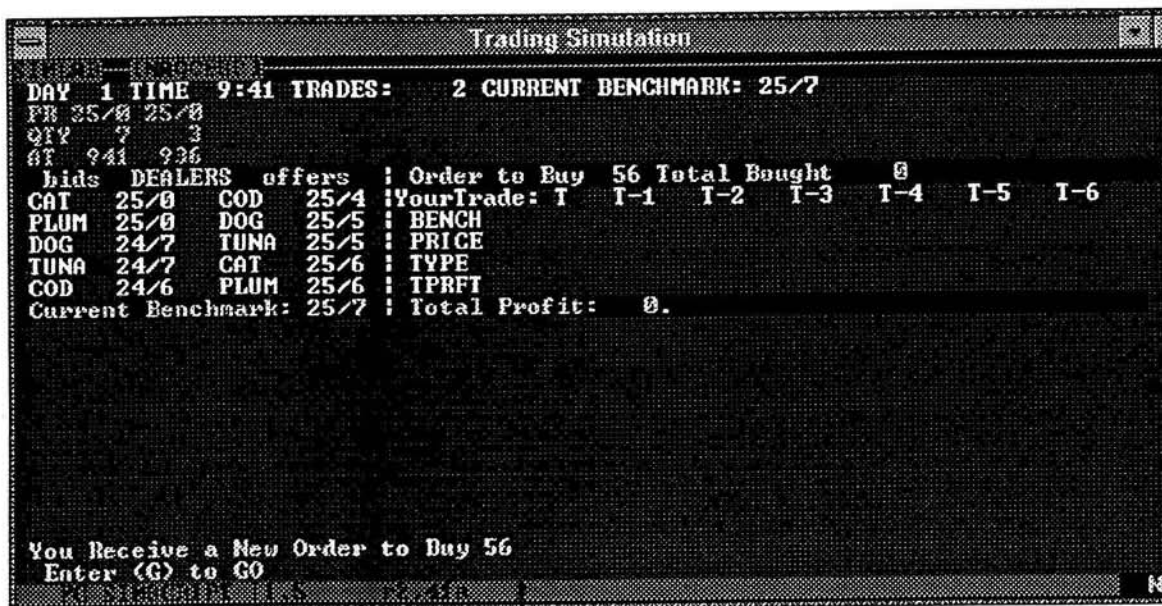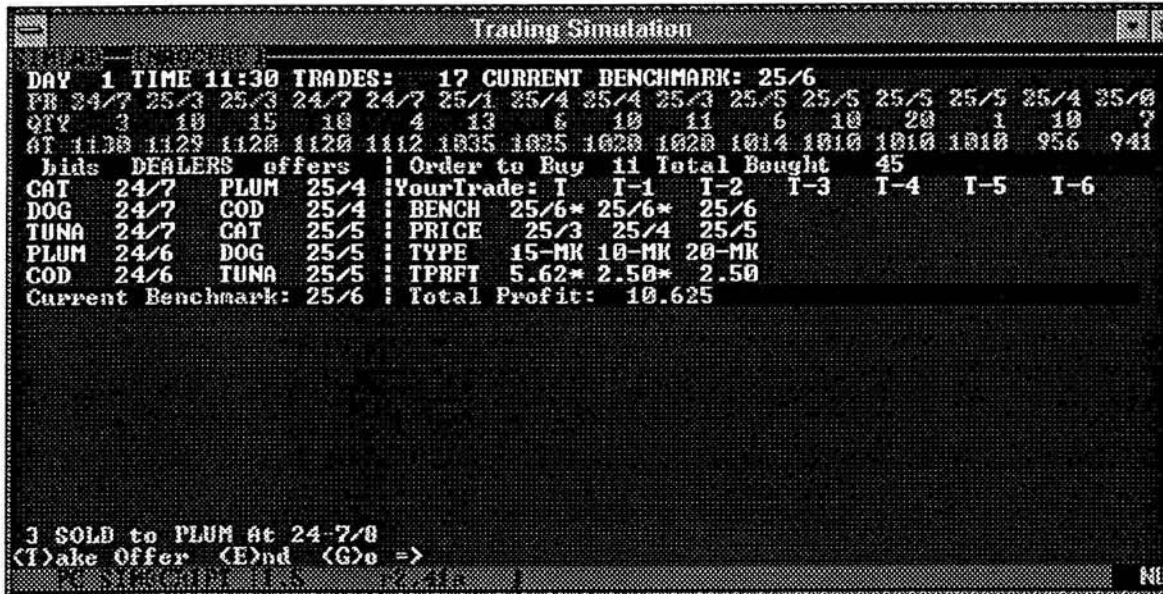
FIGURE 1 (CONT.)
USER TRADING SCREENS FROM THE SIMULATION EXPERIMENT

Two hours into the Scenario 1 simulated trading day (11:30). User has executed three trades, buying 20, then 10, and then 15 from dealers. User has 11 units remaining to buy. The most recent trade is the most profitable, and the user's cumulative profit is 10.625.

About one hour into Scenario 2 experimental run. The limit order book appears below the dealer montage, and show that the best bid and offer of $25 2/8 and $25 3/8, are both currently being made by limit orders. The user has one limit order to buy 20 at $25 1/8 on the limit order book. The user has executed one order so far, buying 15 with a limit order at $25 3/8.
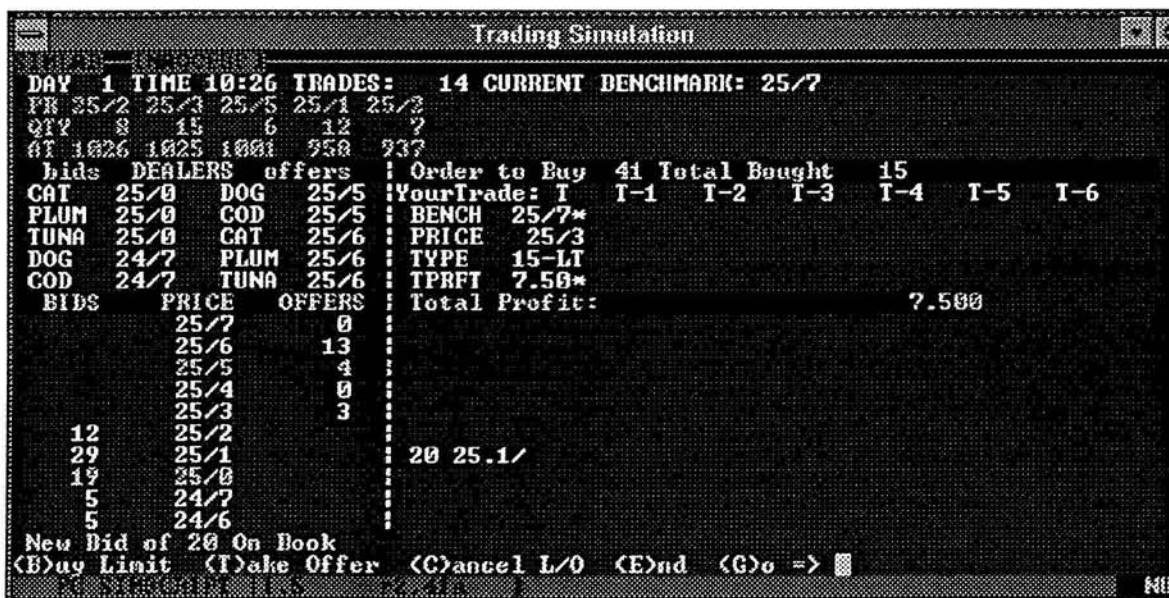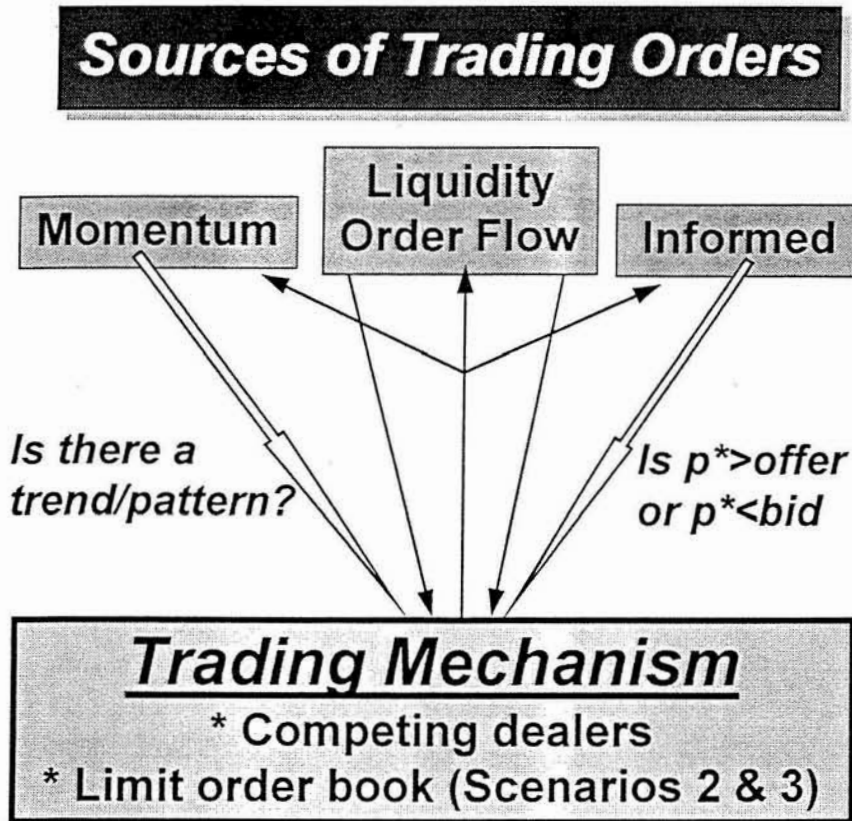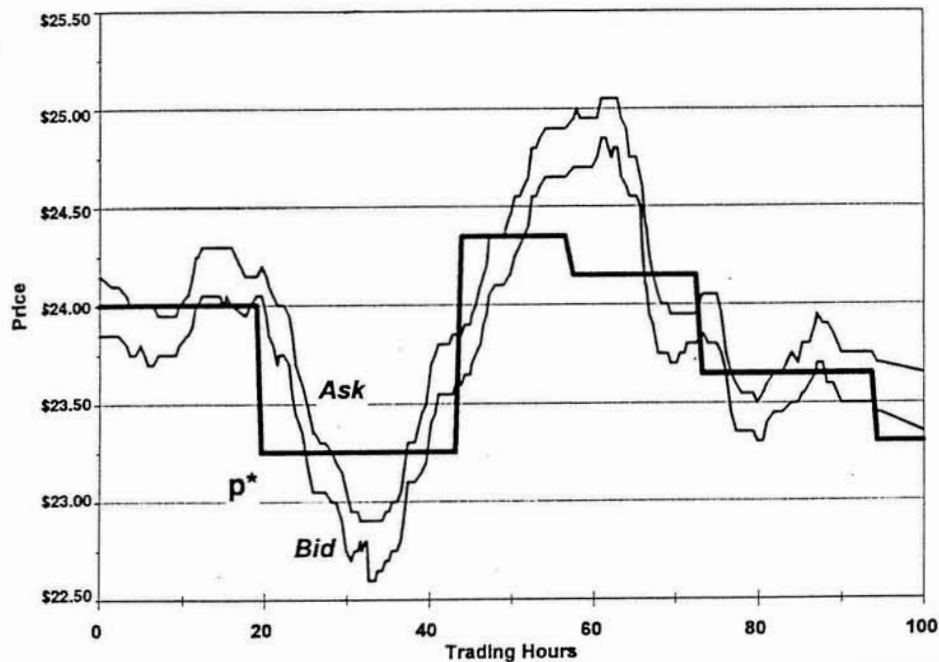
FIGURE 2
SCHEMATIC OF TRADING MODEL: ORDER FLOW, ORDER SOURCES, AND TRADING MECHANISM

Sources of Trading Orders

Momentum   **Liquidity Order Flow**   **Informed**

Is there a trend/pattern?   *Is p*>offer or p*<bid*

Trading Mechanism
* Competing dealers
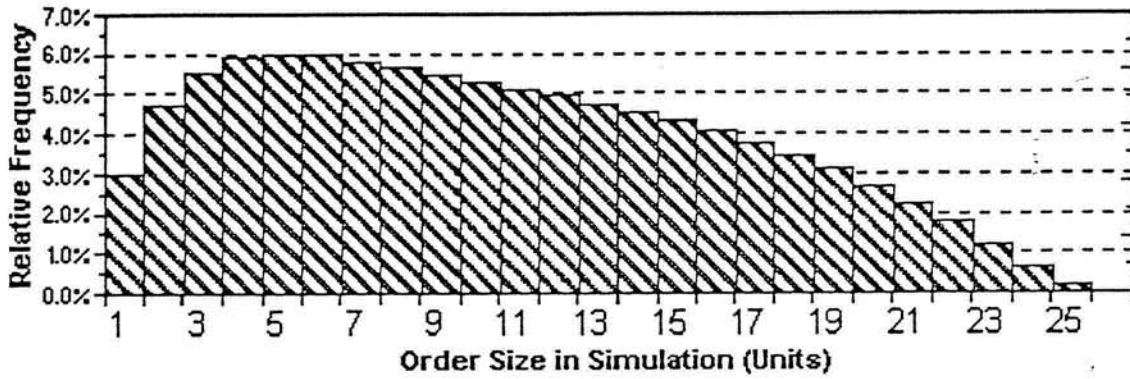* Limit order book (Scenarios 2 & 3)

GRAPH 1
TRACE OF MARKET SIMULATION SHOWS P* AND EVOLUTION OF INSIDE DEALERS QUOTES

GRAPH 2

ORDER SIZE DISTRIBUTION USED IN SIMULATION (BETA DISTRIBUTION)

GRAPH 3

DISTRIBUTION OF SELL LIMIT ORDER PRICES AROUND MARKET QUOTES
DOUBLE TRIANGULAR DISTRIBUTION (YAWL DISTRIBUTION)