Privileged Traders and Asset Market Efficiency: A Laboratory Study

Daniel Friedman*

Abstract

The 39 experiments reported here examine the impact on trading profits and on market performance of awarding special trading privileges to some traders and not others. In call market experiments, the last-mover and orderflow access privileges are both modestly profitable and neither impairs market performance. In continuous market experiments, quicker access to orderflow information is quite profitable and more detailed access is possibly profitable; both privileges seem to enhance market performance slightly. By contrast, privileged marketmaking is extremely profitable and greatly impairs market performance.

I. Introduction

Most major contemporary financial markets award special privileges to some traders. For example, securities traded in the New York Stock Exchange (NYSE) have only one trader, the specialist, who is allowed to make the market by publicly posting bid and ask prices. The Chicago Board of Trade and the Chicago Mercantile Exchange have a wider but still limited set of pit traders who are allowed to make bids and asks and to respond immediately to other traders' bids and asks. There has been a tendency over the years to expand the set of privileged traders and the rights of less privileged traders, but few observers expect to see privileges disappear entirely.

The main reason for granting trading privileges is to overcome a public goods problem. A financial marketplace generates valuable information that becomes freely available to all observers.[1] Market participants, therefore, have an incentive not to pay the costs of creating and maintaining the market. Consumption of market-generated information (and other market services) is excludable to some

*Economics Department, University of California, Santa Cruz, CA 95064. The author thanks the National Science Foundation for support under grant #RI-8812798, Wyn Hillier and Tim Kolar for expert programming, and Chanyong Park, Bret Carthew, Ron LaMont, and Debbie Carson for excellent research assistance. The paper has benefited from suggestions by Tom Copeland, Charles de Bartolome, Pete Kyle, **JFQA** Referee Corinne Bronfman, and an anonymous **JFQA** referee, and from the editorial advice of **JFQA** Managing Editor Jonathan Karpoff.

[1]There is a second (and perhaps even more important) sense in which financial markets provide a public good. Traders desire liquidity, and the more traders participating in a market, the greater the liquidity. This is concurrent consumption with a vengeance: consumption by one trader enhances consumption by another.

degree, so a natural solution to the public goods problem is to charge for selected participation rights. This solution is often realized in the form of transferable "seat" ownership tied to certain trading privileges, as in the NYSE, or in the form of several tiers of participation rights available at different prices, as in NASDAQ's Level 1 to Level 3 screens (see Schwartz (1988), pp. 54–55). Either way, some traders are privileged relative to others, and the privileged traders have the incentive to cover the costs of maintaining the marketplace.

Given this justification for trading privileges, the natural goal is to find an efficient tradeoff between the profitability to privileged traders and the overall market performance. Some privileges, such as dual trading in the Chicago exchanges, are profitable but create conflicts of interest that may impair market performance.[2] There may be other privileges, such as more rapid access to orderflow information, that provide the same profitability and better market performance.

The issue is timely because major financial markets are entering an unprecedented era of rapid change in rules and technology. Some, such as the Toronto Stock Exchange, have done quite well with new electronic technology and others, such as the London Stock Exchange, have done rather badly. A key determinant of success or failure appears to be whether the rule makers find an efficient tradeoff.

What guidance is available to an exchange looking for privileges that provide adequate profits at minimal loss of market performance? The exchange cannot rely much on its own historical experience because the trading environment and the technology are changing so rapidly. Neither can it rely on economic theory. Standard theory assumes either a fictitious Walrasian auctioneer or unorganized search. Theoretical literature that analyzes specific rules of organized exchange is still in its infancy and as yet offers few useful insights into the effects of trading privileges.

Empirical data clearly are essential. Unfortunately, standard econometric analysis of contemporary financial markets is not much help. Comparisons of trading rules are hamstrung by the infrequency of rule changes within a given market and by dissimilarities in the traded securities and the environments across different markets. The econometrician also faces the problem that available measures of trading profits and market efficiency are typically incomplete and/or indirect at best.

Two alternative empirical approaches are available. First, one can write computer programs to simulate trader behavior and compare market outcomes for alternative trading institutions, for example, as in Domowitz and Bollerslev (forthcoming). Such computer simulations are relatively cheap and are a reasonable first cut. The main problem is that the simulations are unable to properly account for changes in trader behavior induced by changes in the rules. Yet the market impact and the profitability of a rule change depend fundamentally on the behavioral response of (privileged and nonprivileged) traders.

This paper pursues a second alternative, laboratory experiments. One can create a small-scale electronic market in the laboratory and allow experienced and profit-motivated human traders to respond to various trading privileges and rule

[2]Dual trading may cause investors to lose confidence and reduce participation. External and internal regulatory bodies may try to counteract this by increasing the resources devoted to policing the marketplace. In either case, market performance is indirectly impaired.

changes. Laboratory methods for analyzing general market behavior have matured and are now a widely accepted empirical technique (Smith (1982), Plott (1982)). In the last decade, there have been dozens of asset market studies in the laboratory; see Sunder (1992) for a survey.

The experimental approach offers several important advantages for studying the market impact and profitability of trading privileges. To begin with, one works with humans who have the means (and motive and opportunity) to adapt their behavior to privileges and other market rules. The measurement problems virtually disappear because the experimenter can observe (or even control) traders' private information and, therefore, can directly measure market efficiency and traders' profits. Moreover, the experimenter can systematically vary the market rules and trading privileges and, therefore, make valid causal inferences about their impact.

An important caveat is in order. Research budgets limit the scale of laboratory markets. One's first impulse in designing a laboratory market environment often is to replicate a field setting (e.g., the NYSE) as closely as the budget permits. Experimentalists have learned to resist this impulse because it usually leads to effective loss of experimental control. They have learned patiently to introduce only a few novel features at a time into the laboratory environment, because only then can they confidently disentangle each feature's direct and interactive effects.

For that reason, the experiments presented below do not depart radically from previous laboratory designs. Traders are exogenously endowed with various privileges in otherwise familiar relevant laboratory asset markets. The immediate goal is to distinguish major effects from minor effects, not to estimate precise dollar values. A longer term goal is to provide a sound basis for larger-scale or more sophisticated laboratory experiments.

The next section briefly describes the trading institutions and privileges, the market environments, and the experimental design. The market environments include the arrival of private information during the course of trade, so we observe the process of price discovery in each trading period. The experiments examine two basic trading institutions: the continuous trading procedure known to academics as the Double Auction, and the periodic call market referred to below as the Clearinghouse. Several privileges are defined for each trading institution. Approximate field counterparts of each privilege are mentioned, but inevitably the field privileges are more complex and are bundled with other privileges or obligations. Therefore, there is no precise mapping between the field privileges and the elementary laboratory privileges tested here.

The third section presents results from 39 laboratory sessions, each consisting of a dozen or more trading periods.[3] The section begins with a qualitative overview and a summary of performance measures. A statistical analysis follows for the market impact and profitability of various trading privileges. The main findings for the Clearinghouse institution (or call market) are that the last-mover and orderflow access privileges both are modestly profitable and neither impairs

[3]A companion paper, Friedman (1993), examines the same data set. It compares asset market performance across the Clearinghouse and Double Auction trading institutions. The present paper focuses on the profitability and market impact of trading privileges within each institution. Portions of Section II and of Section III.A below are adapted from the companion paper and are included here to make the exposition more self-contained. The companion paper should be consulted for a literature survey.

market performance. The main findings for the Double Auction institution (or continuous market) are that quicker access to orderflow information is quite profitable and more detailed access is possibly profitable; both privileges seem to enhance market performance slightly. By contrast, privileged marketmaking is extremely profitable and greatly impairs market performance.

The final section offers a summary and a few interpretive remarks. Official instructions, including descriptions of the market institutions and reproductions of the trading screens, are collected in Appendix A.

II. Experimental Design

A. Basic Procedures

Each experiment reported here consists of a series of 12 or more trading periods (sometimes called "Market Days"), each lasting at most five minutes. The market participants in each experiment, the "traders," are typically eight to 12 University of California undergraduates who buy and sell asset units (called "shares") for cash, using various computerized trading mechanisms described below. At the end of the experiment, the traders are paid the profit they earn, ranging from $10 to $30 in a typical experiment. The stakes seem sufficient to strongly motivate the traders to seek strategies that will increase profit. Due to the market complexities, traders generally appear to require experience in one or two experiments before they become comfortable with their strategies. The data reported here exclude experiments using inexperienced traders.[4]

Asset units are valuable because each share pays a trader-specific liquidating dividend (the "payout") at the end of a trading period. Differences in payout values provide traders with gains from trade[5] in a risky environment. More specifically, in each experiment, there are two or three different trader types with each type consisting of three or more individual traders. The trader type with the highest

[4]Training procedures were as follows. Traders were recruited from large sophomore and junior-level economics classes at the University of California, Santa Cruz. Those who agreed to participate were given copies of the instructions and invited to attend a training experiment using the basic DA institution. Each training experiment began with a 10–15 minute oral review of the instructions, a question and answer period, and a short written quiz. Then three or four practice trading periods (no cash payments) were conducted on the computer system with questions permitted. When all traders were ready, a computerized 8–14 period experiment was conducted. A few individuals with unusually low profits and quiz scores were eliminated and the remaining (80–95 percent) participants were entered into the pool of trained traders, which typically numbered about 40 individuals. Except for a few last-minute replacements, the traders in reported experiments were all drawn from this pool of experienced traders. The data from training experiments have been saved but are not analyzed here because these experiments contain relatively few trading days, are usually dominated by beginner errors, and often contain computer bugs, since beta testing for new versions of the program often was conducted with inexperienced subjects.

[5]Differing payouts are intended as counterparts of trading incentives for participants in contemporary asset markets such as differing tax brackets, differing nonmarketable assets held in portfolios, and differing risk preferences. Traders begin each trading period with a new endowment, typically three shares and $20.00 cash. They earn trading profit by purchasing shares at prices below their own payout and by selling shares for prices above own payout. Hence, both traders earn trading profits when a trader with lower payout sells to a trader with higher payout at an intermediate transactions price. Traders accumulate profits from one trading period to the next, and take home the total earned for all periods in the experiment (or, in some cases, a preannounced fraction of accumulated trading profit, e.g., 50 cents on the dollar).

payout varies across trading periods according to random events. For example, in some experiments, there are four traders of type 1 (payout either $1.80 or $0.40 per share) and four traders of type 2 (payout either $1.40 or $0.80). Each type receives the higher (G = good) or lower (B = bad) payout with 50-percent probability independently each period, so there are four equally likely payout-relevant states, e.g., GB = [G for type 1 and B for type 2], in the HET (for heterogeneous) treatment of uncertainty. The simpler alternative HOM (for homogeneous) treatment of uncertainty, in which the only permitted states are GG and BB, is seldom employed in the experiments reported here. Under either treatment, the traded asset is a claim whose value depends on the agent's type as well as on the "state of nature."

Another information variable deals with the timing of news concerning own payout. The simplest treatment is IM (immediate news): traders receive news of their own actual payout just before the beginning of the trading period. Even with this simple news treatment, traders face uncertainty regarding the market value of the asset because they do not know other traders' realized payouts. In more complex treatments, traders begin each trading period uncertain of the payout, but are privately notified by the computer of their own realized payout for that period before it ends. The usual treatment is SEQ (sequential news): traders of different types receive payout news at different times, the sequence being random. For example, the bad news that their payout is $0.80 might appear privately on type 2 traders' screens at 90 seconds into a trading period, and the good news that type 1 traders' payout is $1.80 might appear privately on their screens at 180 seconds. The final treatment is SIM (simultaneous news): traders of all types receive news simultaneously, for example at 150 seconds into a trading period.

These treatments allow for a considerable range of environmental complexity for asset trading, ranging from rather transparent (IM/HOM news with two trader types) to quite opaque (SEQ/HET news with three trader types). Copeland and Friedman (1991) and the papers it cites show that these treatments do systematically affect market performance. As explained in Section II.C below, one can compute a priori equilibrium predictions for trading activity and market efficiency for each environment independent of the trading institution, and the actual market outcomes can be compared to the equilibrium forecast.

B. Market Institutions and Privileges

All market institutions and privileges examined here are implemented by computer programs that collect traders' orders at the keyboard and compute and report outcomes to traders' screens. Currently, the programs are run under UNIX on a Sun Sparc1+ workstation and networked terminals or PCs. Complete rules and examples of screen displays can be found in Appendix A.

In the continuous Double Auction (DA) trading institution, each trader at any moment can enter a bid (an offer to buy an asset unit for a specified amount of cash) or an ask (a similar offer to sell) from her interactive terminal, can use the terminal to accept the current best (highest) bid or best (lowest) ask offered by her fellow traders, and can cancel her outstanding bid or ask. The computer serves as the only communications link between traders. It also serves as the record-keeping device and enforces the rules. For instance, transaction requests that would result in a negative cash or asset position are not executed, but rather generate descriptive error messages. Most major field asset markets with continuous trading use some sort of noncomputerized Double Auction institution to organize trade. The new Globex system of the Chicago futures exchanges is a computerized version.

Several variants of the DA have been tested. The crudest variant involves a treatment variable called *post*, which can disable the facility for entering ("posting") bids and asks for some traders. (Of course, the market could not function if *all* traders were so disabled.) In effect, the nondisabled (*post* = 1) traders retain dealership or marketmaking privileges, while the remaining (*post* = 0) traders are price-taking customers.[6,7]

A more subtle variant involves a delay in orderflow information. Normally, when some trader improves the best bid or ask, the information appears on all traders' screens within a tenth of a second. Under the *delay* treatment, some or all traders receive the information with a lag. Typically *delay* = 5 seconds. Delays to some traders occur naturally in noncomputerized trading systems, and some partially computerized systems (e.g., quote screens for U.S. Treasury securities) have employed them in the past.

The time lag in *delay* sometimes leads to the presence of crossing offers, e.g., a delayed trader enters a bid of $1.25 not knowing that some other trader has already entered a $1.20 ask. The default in this case is to execute the transaction at the older price ($1.20 in the example), but sometimes nondelayed traders have an additional privilege called *arb*, for arbitrage. All traders with this privilege are immediately notified when offers cross, and the first to hit a single key ("y") automatically buys at the lower ask price and simultaneously sells at the higher bid price, thus obtaining an arbitrage profit.

A final DA variant called *book* provides enhanced orderflow information. In the default treatment *book* = 1, only the best bid and ask are displayed. In the *book* = 2 treatment, some or all traders have a modified screen display that shows all bids and asks, arrayed from best to worst. NASDAQ's distinction between Level 1 and Level 2 screens is similar, but anonymity is preserved in the laboratory markets in that a trader does not see trader identification for orders other than her own.

The other main market institution is the Clearinghouse market (CH), also known as the call market. In the CH, traders enter bids and asks at their terminals as in the DA, but these offers are interpreted as limit orders and are not executed immediately. Rather, at the end of the clearing period (typically lasting 60 seconds) or when all traders indicate they are ready, the bids and asks are aggregated respectively into market demand and supply curves, and the market is cleared in the usual fashion. That is, the price (or the midpoint of the range of prices) is found

---

[6]Some theoretical microstructure literature, beginning with Demsetz (1968), models the NYSE specialist as the sole trader with posting privileges. The discussion in Section IV emphasizes that specialist privileges are not so simple under current NYSE rules.

[7]In standard nonasset (perishables or commodities) versions of DA, beginning with Smith (1962), bids are disabled for some traders (the sellers) and asks are disabled for other traders (the buyers). In the Posted Offer (PO) trading institution, one side of a perishables market is completely disabled, e.g., buyers cannot bid (or ask) and can only accept sellers' asks or remain silent. The basic results are that the DA institution is very efficient and the PO institution is fairly efficient (and the disabled side earns lower profits)—see Davis and Holt (1992) for a recent survey. In the present DA experiments, each trader can enter both bids and asks unless he is disabled (*post* = 0), in which case he can enter neither bids nor asks.

at which the supply revealed by the asks equals the demand revealed by the bids, and all higher bids and lower asks are filled at this clearing price. Thus, the CH can be described as a "batch" (discrete time) mechanism that provides a uniform price to all transacters in each clearing. Versions of this institution are used to open trade in major stock exchanges. The Arizona Stock Exchange (formerly the Wunsch Auction System) uses a fully computerized version of this institution. See Friedman and Rust (1993) for recent research and surveys on both the DA and CH institutions.

There can be more than one clearing in a trading period. Most experiments with two trader types have three clearings per period: one prior to news arrival, a second after one trader type has been notified of actual payout, and the last after the second type has been notified.

Variants of the CH institution provide additional orderflow information. In the basic version, denoted *book* = 0, traders submit bid and ask orders "blind" in that they have no direct knowledge of what orders other traders are submitting. In the variant *book* = 1, traders' screens display a continuously updated "indicated price," the price at which the market would clear if no further orders were received. The Toronto Stock Exchange offers this sort of indicated price to all traders. The other orderflow variant, *book* = 2, also gives traders a continuously updated summary description of the order book. Near-marginal orders (five orders on either side of the indicated price) are displayed and allow the trader to see the ceteris paribus price impact of any new orders she might contemplate. The Arizona Stock Exchange originally offered this sort of complete orderflow information, but recently has allowed traders to temporarily "hide" large extramarginal orders.

A second variant on the CH institution, called *extratime,* allows some traders the privilege of submitting (or perhaps removing or offsetting) new orders after time has expired for all other traders and their orders are committed, as in the opening procedures at the Toronto Stock Exchange. In my experiments, the privileged traders are usually allowed 15 or 20 extra seconds for this purpose.

In the basic CH, traders have the right to delete (or "pull") bids and asks during the clearing period, but are not allowed to enter offsetting orders (e.g., a bid price lower than one's own ask price). These rules are summarized in the default treatment *pull* = 1. In the variant *pull* = 0, orders must be confirmed before they are recognized by the computer and they are "committed," i.e., they cannot be offset or pulled. In the variant *pull* = 2, orders are offsetable as well as pullable. The final variant, *pull* = 3, allows traders at a single keystroke to improve their offers (increase bids and decrease asks) to meet the market. The Arizona Stock Exchange (Wunsch auction) again features an improvement rule of this sort.

Table 1 lists the main privileges used in the experiments. Recall that most DA experiments feature several news events that divide each trading period into *subperiods*—time intervals in which price discovery recurs. Likewise, there typically is a clearing after each news event in CH experiments. Hence, the number of trading periods listed in Table 1 understates the number of available observations. A three-page table presenting the design of each experimental session is available on request.

TABLE 1

Privileges Tested

Privilege	No. of Periods	Present in Experiment (number of periods)
Double Auction		
Marketmaking	114	7(4), 10(6^m), 13(24^m), 35(40^m) , 36(40^m)
Rapid Access	92	6(16^a), 8(14), 9(18^a), 12(12), 19(16), 37(16)
Detailed Access	69	6(16), 12(12), 10(9), 17(8), 18(8), 37(16)
Clearinghouse		
Detailed Access	44	5(16), 14(12), 24(8), 39(8)
Last Mover	98	1(8), 2(6), 3(16), 4(16), 14(12), 26(8), 27(12), 28(20)

The main privileges in the Double Auction market institution are marketmaking (*post** = 1), more detailed access (*book** = 2), and more rapid access (*delay** = 0) to orderflow information. The main privileges in the Clearinghouse market institution are last-mover (*extratime**) and more detailed access to orderflow information (*book** = 2). (Other treatments not emphasized in the present paper include temporary trading bonuses or commissions (*comm*) and alternative Clearinghouse rules on order cancellation or improvement (*pull*).)

The experiments are numbered 1–39, as listed in Table 2.

Unless otherwise noted, privileges are granted to one trader (out of three or more traders) of each type on a rotating basis. An m superscript indicates that multiply privileged traders are present in some periods. An a superscript indicates that the abritrage (*arb**) privilege is also included.

C. Equilibrium Forecasts

Traders' direct knowledge of their laboratory environment is quite limited. Each trader knows own payout and endowment parameters and knows (ex post) news arrival time, but does not know the parameters or the news (or even the news arrival times) of other traders. To analyze the situation faced by traders as a game of incomplete information is a daunting task, particularly in the case of DA markets (and CH markets with *book* > 0) since continuous-time strategies then must be chosen. Fortunately, much simpler complete information approaches seem successful at organizing the data in market experiments with several trading days (see Smith (1990)). Presumably, traders learn to behave as if they acquire the relevant information from market outcomes.

The simplest complete information model assumes risk neutrality and treats all private information as if it were public. The literature (e.g., Sunder (1992)) refers to this model as RE, TRE, or FRE for "true" (or "telepathic") or "fully revealing" rational expectations, and notes that it characterizes a fully efficient market and represents the fundamental value of an asset. FRE also generally outperforms alternative models in predicting asset prices and is relatively simple to compute. For these reasons, I use it as the benchmark for prices. One computes FRE equilibrium prices as follows. First, for each payout relevant state z (e.g., z = GB), set the final equilibrium price $p(\text{FE}, z)$ equal to the highest payout in that state.[8] For example, if there are two trader types with respective payouts [G:$2.00, B:0.30] and [G:1.70, B:0.80], then p(FE, GB) = max{$2.00, $0.80} = $2.00. Next, for each time of interest, look at all news messages received so far in the trading period and update the state probabilities $\pi(z)$. To continue the

[8]The logic here basically is that the three or four traders with the highest payout will bid the price up to their payout level because their demand is very large at that price (given the $20 per capita cash endowment), while asset supply is fixed at three shares per capita. It follows that the traders with highest payout will hold all shares at the end of the trading period.

example, each of the four possible states has initial probability 0.25 but, after B news to type 2 traders, the probabilities become $\pi(\text{GB}) = \pi(\text{BB}) = 0.50$ and $\pi(\text{BG}) = \pi(\text{GG}) = 0$. Finally, set the FRE equilibrium price p^* equal to the expected FE price, using the updated state probabilities. In the example, we get $p^* = (0.5)p(\text{FE}, \text{GB}) + (0.5)p(\text{FE}, \text{BB}) = (0.5)(\$2.00) + (0.5)(\$0.80) = \$1.40$ as the equilibrium price when the news 2B arrives. Thus, one obtains a price forecast for every subperiod of a DA market and for every clearing in a CH market.

D. Market Performance Measures

I employ three measures of market efficiency. Informational efficiency is measured in each subperiod (or clearing) in which transactions occur as the root mean squared deviation (*RMSE*) of transaction prices (or clearing price) from the fully efficient FRE price forecast. Allocational efficiency is defined by the statistic *AIE*, the unrealized trading profit as a percentage of potential total trading profit in a given trading period. An example of this calculation appears at the end of the first paragraph in Section III.A below. The third efficiency concept is market depth or *Spread*, measured as the time-weighted average difference between the ask and bid price in a DA subperiod, and measured in a CH clearing as the difference between the best rejected (extramarginal) bid and ask prices.

Another market performance measure is trading volume, measured in each subperiod or clearing as the number of shares sold (or bought). Volume has no direct efficiency implications other than that a minimum volume is required to move from the initial allocation to an efficient final allocation, but it is of interest to theorists who seek fuller characterizations of the trading process and to practitioners whose income depends on the volume of orders.

The other relevant market outcome to be measured is privilege profitability. For this purpose I use *excess profit*, defined as the mean trading profit earned by privileged traders minus the mean trading profit earned by nonprivileged traders in a given trading period.

III. Results

A. Overview of the Data

To provide the reader with a background against which the statistical analysis can be viewed, I begin with a description of events in an early DA trading period and an early CH trading period. Figure 1 shows the first trading period of the first DA experiment. The three news events (notification of realized payout to each of the three trader types) divide the four-minute trading period into four one-minute subperiods, as indicated by the vertical lines in Figure 1. The market bid (indicated by the lower step-wise line) opened at 20 cents about 15 seconds into the trading period and rose to \$1.00 a few seconds later. Shortly thereafter the best ask opened at about \$1.10 and, after three quick transactions (indicated by stars), bounced up repeatedly to \$1.50 as four more transactions (all accepted asks) occurred in the first subperiod. At the end of the subperiod, 11 shares were held by type 1 traders, 11 by type 2s, and five by type 3s. The transaction prices are considerably below the (unconditional) FRE price forecast of $1.65, resulting in a RMSE of 52.3 in this initial subperiod. Type 3 traders received B (low payout) news to begin the second subperiod, lowering the FRE price forecast to $1.625 for this subperiod. The news appeared to have little effect on the market since the eight transaction prices were generally a bit higher and share allocation changed little. In the third subperiod, transactions prices again generally rose slightly, and type 2 traders (who received G news) were net purchasers of five shares from type 3s on a volume of nine shares. In the final subperiod, type 1 traders received B news and sold most of their shares to the type 2s at prices mostly 10–20 cents below the equilibrium price of $1.65 (RMSE = 12.5 cents). The final allocation deviated from the equilibrium forecast (recall that in final equilibrium, all shares are held by the high payout type, here type 2 traders) by the 2 shares still held by the type 1 traders. The final allocation failed to realize 2 × ($1.65 − 0.25) = $2.80 of the $18.90 in potential gains from trade, for an allocational inefficiency AIE) of 14.8 percent.

FIGURE 1

A DA Price-Time Graph

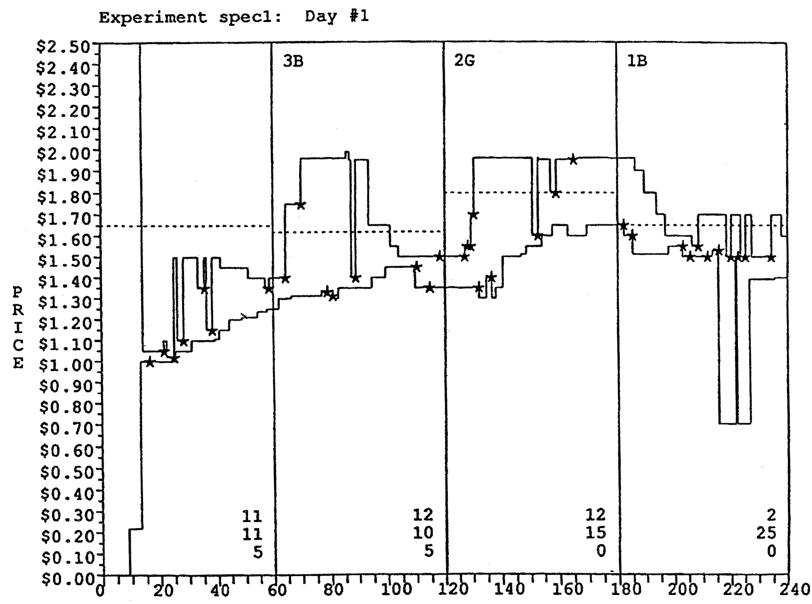

The upper step function is the best ask price, the lower step function is the best bid price, and stars indicate transaction prices. The horizontal dashed lines indicate the FRE price, the fundamental value of the asset. Vertical lines indicate news events The trader type and content (e.g., 3B for the trader type 3 receiving notification of the lower payout) are noted at the top of the line, and the asset allocation (e.g., 11, 11, and 5 shares held, respectively, by traders of types 1, 2, and 3 at the time of the news event) is noted at the bottom of the line.

Figure 2 shows that in the first clearing of the first period of the CH experiment Chm1, type 1 traders sold three shares to type 2 traders at the price of $1.45. The second clearing occurred after type 2 traders knew they would receive the higher payout (indicated by the "2G" in the upper right corner of the panel); no trades occurred but the best rejected bid and ask were within one or two cents of each other, near $1.55. The third clearing was preceded by B (low payout) news to type 1 traders, who sold 11 of the 15 shares they held at a price of $1.50. The final clearing price then turned out to be 20 cents below the equilibrium value (indicated by a dotted line), and four shares were misallocated.

FIGURE 2

A CH Supply-Demand Graph

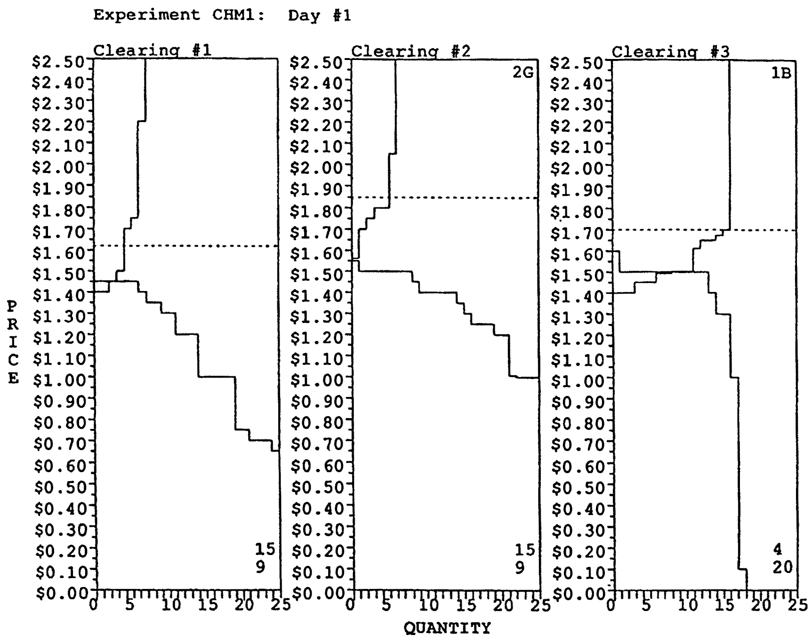

The increasing step function is the supply revealed by all submitted asks in the given clearing, and the decreasing step function is the demand revealed by the submitted bids The intersection of supply and demand determines the clearing price and transaction quantity (use the midpoint if the intersection is a vertical line segment, and use the right endpoint if the intersection is a horizontal line segment.) The horizontal dashed lines indicate the FRE price, the fundamental value of the asset The news event is indicated in the upper right corner, e.g., 2G for trader type 2 receiving notification of the higher payout. The asset allocation after market clearing is indicated in the bottom right corner, e.g , 15 and 9 shares held, respectively, by type 1 and type 2 traders after the second clearing.

Figure 3 summarizes trading period #15 of Chm4 and illustrates the *extratime* and *book* = 2 privileges. At the end of regular time in the third clearing, the indicated price was $2.00, well below the equilibrium price of $2.15. Trader 2, the type 1 privileged trader, then entered two additional bids at $2.05, which (as he knew from orderbook display) did not affect the price. A third bid at the same price moved the price up two cents, for a positive marginal revenue of 13¢ − 2 × 2¢ = 9¢. Additional bids at this price would have had negative marginal revenue and were not entered. The result of activity was a 39-cent profit increment, and a slight *increase* in market efficiency (misallocations unchanged at 1 share, deviation of price from equilibrium reduced from 15 cents to 13 cents.) The other privileged trader was of type 2 and evidently had no profitable *extratime* opportunity.

FIGURE 3

Supply-Demand Graph for Chm4, Day 5

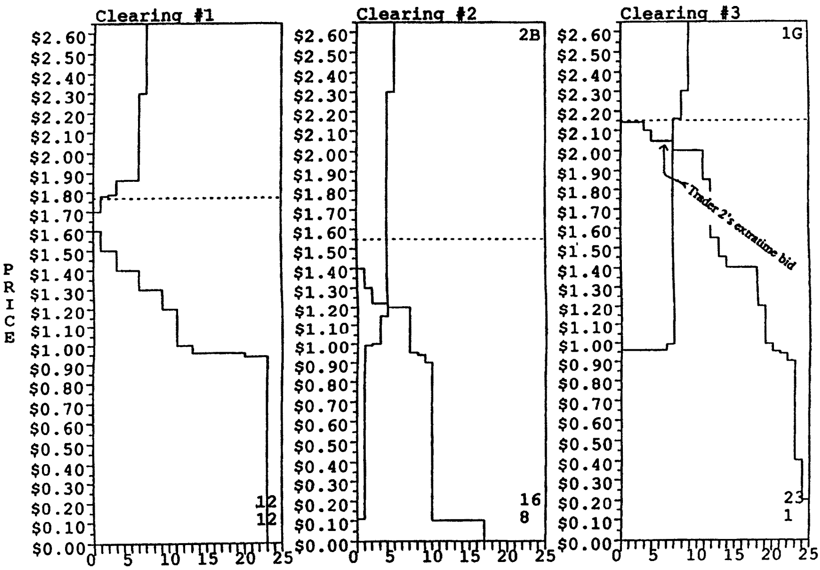

This illustration suggests that measuring privilege profitability is not entirely straightforward. To begin with, large period to period and trader to trader fluctuations in profits disguise the effects of privileges. In Chm5, for example, nonprivileged traders earned profits in a single period ranging from \$3.97 to −\$1.91, while privileged traders' profits ranged from \$3.43 to −\$1.67. Inspection of individual bid/ask data suggests that opportunities for either of the two specialists to profitably use *extratime* opportunity only exist about every second day, and simultaneous opportunities for both specialists are rare. Even when an opportunity exists, the extra time often is insufficient for a trader to fully exploit it, and occasionally an overeager trader will bungle the extra opportunities and actually reduce profit. The excess profit measure averages across traders to reduce the noise, but since there are typically two or three times as many nonprivileged as privileged traders over which to average, the data remain rather ragged and certainly not normally distributed.

Table 2 provides an overall summary of market performance in each experiment, reporting the mean performance measures and excess profits (and, in parentheses, the standard deviation and number of observations). Considerable variation is seen across experiments in all measures. These variations should be thought of as arising from differences in the institutions and privileges, differences in the environments (e.g., 3×3 SEQ versus 2×4 IM), and differences in the participating traders.

The main task in the rest of this section is to disentangle the effects of the trading privileges from the other influences. Statistical hypothesis tests will be the primary tool, including both standard Student *t*-tests and also (because much of the data are far from normally distributed) two popular nonparametric tests. The

TABLE 2

Performance Summary
Mean (Standard Deviation; NOBs)

Experiment Code	RMSE in Cents	AIE in Percent	Spread in Cents	Volume in Shares	Excess Profit in Cents
1. Ce	11.6 (16.4, 53)	6 9 (10 4, 18)	12.2 (17 5; 53)	4.9 (3 6, 54)	11 8 (75.8, 8)
2 Ce	15.4 (16 7, 41)	10 6 (20 7, 14)	28.2 (31 3, 42)	3 3 (3 1, 42)	8.3 (118 7, 6)
3. Ce s	20 5 (20 8; 60)	5 7 (6 6; 20)	20 1 (26 8; 59)	5 2 (3 4; 60)	32 8 (131 0, 8)
4 Ce	17 1 (16 8; 72)	2 6 (5.4; 24)	23 1 (30 5, 72)	5 1 (3 7, 72)	−4.5 (90 7, 16)
5 Ce	10 3 (11 6, 60)	2 7 (4 4, 20)	11 8 (16 4, 60)	4 4 (3 9, 60)	8 5 (47 6, 16)
6. Dr	29 8 (25.8; 44)	3 5 (6.7; 20)	41 1 (22.4, 60)	5 8 (6 4, 60)	68 9 (182 8, 16)
7 Dm c	34.9 (18.8; 45)	6.6 (9 9, 12)	47 0 (23.3, 48)	7 8 (4 8; 48)	74 2 (145 0, 4)
8 Dr c	20.7 (14 8, 56)	2 2 (4 1, 18)	30 3 (17.3; 72)	5 5 (5 1, 72)	13 4 (41.5; 14)
9 Dr c	33.9 (25 6; 72)	1 4 (2.3, 20)	39 6 (20 5, 80)	6.7 (5 8; 80)	45 7 (70.7, 18)
10. Dm.c	31 1 (22.9, 73)	3.7 (9 3, 20)	46 3 (30 4, 80)	5 7 (5 2; 80)	108 7 (176 5, 15)
11 Dc	26.0 (16 0, 62)	3 5 (5 5, 16)	30 1 (20.3, 64)	7 6 (5 6, 64)	NA
12 Dr.s	13 7 (14.4, 30)	1 7 (5 7, 30)	32 4 (14 6, 30)	14 5 (2 8, 30)	16 6 (40 3; 24)
13 Dm s	78.7 (68 6; 30)	1.8 (4 0, 32)	66 3 (37 3, 32)	12 6 (2 4, 32)	86 0 (125 3; 24)
14. Ce.s	4 4 (3 5, 17)	1 3 (2 4, 17)	5 8 (4 5, 17)	11.1 (1 7, 17)	3 7 (13 6, 12)
15 Cb.s	1 7 (1.5, 23)	1 1 (2.8, 23)	5 2 (6.7, 23)	11 8 (1 1, 23)	NA
16. Cb.s	4 2 (13.2, 57)	0.2 (0.7, 19)	6.0 (11.8, 57)	4.2 (3.2, 57)	NA

continued on next page

excess profit data are based on paired observations, so the signs test is appropriate. The null hypothesis is that excess profits are as likely to be negative as positive, and the alternative hypothesis is that positive excess profits are more likely. One counts the number r of instances of positive excess profits and the number w of negative. The signs test statistic then is $z = (r - w)/\sqrt{r + w}$. Its asymptotic distribution is unit normal under the null hypothesis. The market performance comparisons involve a set of trading periods in which privileged traders are present and a control set of otherwise similar trading periods with no privileged traders. Absent paired observations, the nonparametric Wilcoxon statistic seems appropriate. It measures the tendency of one group of observations (e.g., RMSE with privileged traders present) to be larger than another control group of observations (RMSE with privileged traders absent) by rank-ordering the combined observations and summing the ranks of observations from the first group. The asymptotic distribution of the standardized test statistic again is unit normal under the null hypothesis that both groups have the same distribution.

A final statistical remark is in order. Trading periods (and subperiods) are not really independent observations. Traders' expectations from previous periods and subperiods create some interperiod dependence. In the present case, this will tend to bias test statistics toward zero, since interperiod dependence will tend to blur the effects of changing the rules or privileges. (As usual, there is also the opposite effect that the effective sample sizes may be smaller than the number of periods.) These sampling problems preclude the use of precise small sample distributions. I believe that, on balance, conventional asymptotic test statistics will be conservative in the sense of underreporting the significance of measured performance differences. Most natural alternative hypotheses in my analysis are

TABLE 2 (continued)

Performance Summary
Mean (Standard Deviation; NOBs)

Experiment Code	RMSE in Cents	AIE in Percent	Spread in Cents	Volume in Shares	Excess Profit in Cents
17 Db	24 2 (17 3; 56)	4.7 (11.3; 23)	57.6 (30.3; 69)	4.4 (4.2; 69)	−51.0 (136 1; 6)
18 Db	27 1 (16.4; 69)	4 9 (10.2; 23)	43 3 (20 0; 69)	9 2 (7 0; 69)	−6 8 (147.5, 6)
19 Dr	36 8 (18 8; 72)	6.2 (7.7; 24)	42.3 (17.7; 71)	7.5 (3.3; 72)	NA
20 Db	21 7 (15 0, 36)	2 3 (4 5, 12)	34.8 (18.3; 36)	6.4 (4.2; 36)	NA
(31)	25 7 (22.6; 26)	8 3 (12.9, 11)	63 4 (36.2; 33)	5.2 (5 4; 33)	NA
21 Db	22 1 (8.4, 16)	2.5 (4 3, 6)	32 9 (15.3; 18)	8 1 (5 4; 18)	NA
(32)	24 2 (16 7, 28)	0.6 (1.5; 12)	42 0 (20 5, 36)	5 4 (5 5, 36)	NA
22 Db s	17 7 (12 5, 20)	1 3 (2 1, 20)	26 1 (7.5, 20)	15 8 (12 5; 20)	NA
(33)	23 8 (12 1, 20)	2 4 (2.4, 20)	35 6 (8 3, 20)	15 6 (2 6, 20)	NA
23 Db.c	15 7 (11 9, 34)	1.2 (1.9; 34)	26 4 (14 3; 40)	5 8 (11.9, 40)	NA
24. Cb	14 3 (10 9; 69)	6.3 (7.3, 23)	10 3 (12.5; 69)	4 8 (3.0, 69)	21.1 (58.2; 8)
25 Cb	16.3 (14.1; 66)	9 1 (10.3; 22)	9.6 (6.7, 66)	3 7 (2 5; 66)	NA
26. Ce	21.5 (19.0, 42)	11 4 (7.9; 14)	14 1 (20.0, 42)	6.5 (2.9; 42)	65.4 (255.9; 8)
27. Ce	18 0 (15 4; 60)	2.5 (4 5; 20)	15.0 (15.8; 55)	5.1 (4 2; 60)	−8.5 (89 6; 12)
28. Ce	24.0 (33.2, 60)	3.5 (7 0; 20)	19 1 (15.9, 60)	4 0 (3.6; 60)	29.5 (93.4, 12)
29. Cb	11.3 (12.5, 72)	3 8 (6 5; 24)	15 2 (16 8, 72)	4.4 (3.4; 72)	NA
30. Cb c	25 1 (18.5, 64)	3.1 (4.9; 16)	14 8 (13 1, 64)	5.1 (4.1; 64)	NA
31 Cb	20 4 (16 9; 33)	2.6 (3.4, 11)	18.6 (12.1, 33)	6.4 (4.3, 33)	NA
(20)	27 2 (25 9; 36)	4.2 (3 9; 12)	31.5 (30.7, 36)	4 1 (4.0; 36)	NA
32. Cb	20 9 (17.3, 36)	9.3 (13.2, 12)	15.5 (23 9; 36)	4.8 (3.9, 36)	NA
(21)	12.4 (12.4; 36)	2 5 (4 2; 12)	26.2 (25.6; 36)	4.8 (4.8, 36)	NA
33 Cb.s	22.5 (25.0; 20)	2 7 (3.1; 20)	12.9 (10.4; 20)	14.1 (3.1, 20)	NA
(22)	26 0 (17 2; 20)	4.2 (7 3; 20)	22.7 (19.6; 20)	13 1 (4 9; 20)	NA
34. Cb c	20 0 (19.8, 40)	2.8 (2.5, 10)	19.7 (27 7; 40)	5 2 (3.7, 40)	NA
(23)	22 9 (14 6; 40)	7.4 (10 1, 10)	28.7 (21.7; 39)	4 5 (2 9; 40)	NA
35. Dm.s	40 5 (24.8; 52)	5 6 (6 4; 52)	76.6 (39.9; 52)	22.4 (7 3; 52)	398.2 (352.9; 40)
36. Dm.s	27.2 (19.0, 52)	4 2 (5.6; 52)	49.1 (28 5, 52)	23.1 (5.7; 52)	306.9 (379.9; 40)
37. Dr	18.3 (14.9; 53)	0 7 (1.4; 24)	34.1 (16.4, 72)	4.4 (4.4; 72)	57.1 (114.3; 16)
38 Cb.s	22 7 (26 9; 43)	6 1 (7 6; 43)	20.7 (32.4, 37)	9 9 (3 7; 43)	NA
39. Cb	14 2 (21.0; 69)	3.8 (7.7; 23)	14.1 (14.8, 69)	5.1 (3.6; 69)	13 (76.6; 7)

The experiments are listed in approximate chronological order.

The codes identify some of the more important treatments used in each experiment as follows: C indicates the Clearinghouse or call market institution; e indicates the last-mover (*extratime**) privilege and b indicates detailed access to orderflow information (*book* = 2) D indicates the continuous Double Auction market institution; m indicates privileged marketmaking (*post** = 1), r indicates rapid access to orderflow information (*delay** = 0), and c indicates temporary trading commissions (*comm*) The suffixes " s" and " c" indicate, respectively, informationally simple (e g , two trader types and immediate notification of payouts) and informationally complex (e.g., three trader types and sequential notification of heterogeneous payouts) environments. A three-page table laying out the design of each experiment is available on request.

RMSE measures informational inefficiency as the root mean squared deviation of transaction price from fundamental value in each subperiod. AIE measures allocational inefficiency as unrealized gains from trade as a percentage of maximal gains from trade in each period. Spread measures market depth as the average difference between ask and bid price (or best rejected bid and ask in the clearinghouse) in each subperiod. Volume is trading volume by subperiod. Excess profit is the difference between privileged and unprivileged traders' average profits in each period.

Some sessions are split between Double Auction (D) and Clearinghouse (C) The double-line entries indicate companion split sessions, with the companion experiment number indicated in parentheses, e.g., (31) for experiment 20.

one-sided. Therefore, I will refer to a positive statistic in excess of 1.20 as possibly significant and one in excess of 2.00 as significant, since these correspond roughly to the 10-percent and 2-percent confidence levels, respectively, for the asymptotic distribution and the one-sided alternative hypothesis.

First-Round Results

The first 16 experiments listed in Table 2 were conducted and analyzed before the remaining experiments. The analysis suggested several conclusions. Among the Double Auction privileges, the posting (marketmaking) privilege clearly had the most dramatic impact. Traders with the privilege, on average, earned more than twice the profit of unprivileged traders. Market efficiency was impaired: RMSE and Spread also roughly doubled, and AIE increased and volume decreased, but not as significantly.

The other DA privileges had some detectable effects. Delaying orderflow information to unprivileged traders increased profit of privileged traders by about 25 percent, but had little impact on market performance. Universal access to full orderflow information surprisingly appeared to *increase* spread and RMSE, but privileged access was not tested in these experiments.

Clearinghouse privileges also produced some surprising results in the first-round experiments. Typically, *extratime* and *book* = 2 privileges were combined, which increased market efficiency as well as profitability. The apparent impacts of committed (*pull* = 0) and offsetable (*pull* = 2) offers by all traders both were minor.

The first set of experiments was designed to cover a lot of territory superficially (a 1989 working paper, available on request, provides a detailed discussion of design and results). More reliable conclusions regarding the profitability and market impact of trading privileges required a new set of experiments designed for the purpose. These follow-up experiments are numbered 17 and above in Table 2.

C. Follow-Up Results

Based primarily on follow-up experiments, Tables 3 and 4 report, respectively, the excess profit and the market performance impact for the most important trading privileges. The first pair of rows in Table 3 consider privileged access to Clearinghouse orderbook information (*book** = 2). The top row shows that of the 44 trading periods in which the privilege was present, excess profit (average privileged trader profit minus average nonprivileged trader profit) was positive 26 times and negative 18 times (Nobs); that it averaged seven cents per period overall with a \$1.32 standard deviation, for an insignificant *t*-statistic of 0.37 and a possibly significant signs statistic of 1.20. The second row ("narrow def.") restricts the comparison to the follow-up experiments 24 and 39, the subset of the data in which no privileges were present other than *book** = 2. The sample size here is small and, in two periods, the privileged traders earned very negative profits. Examination of the raw data reveals trader confusion as the cause; if these two aberrant data points were removed, both the narrow and broad definitions would yield signs statistics in excess of 1.20.

TABLE 3

Excess Profit

Privilege	Nobs	Mean	Std Dev.	Signs	t
	Summary			Tests	
Clearinghouse					
1. Detailed Access	26+18	$0.07	$1.32	1.20	0.37
(narrow def)	10+6	−0 03	2 00	1.00	−0.05
2. Last Mover	45+41	$0.22	$2.14	0.43	0 94
(narrow def)	29+29	0 26	2 56	0	0.77
Double Auction					
3 Detailed Access	37+24	$0.42	$2.40	1.67	1 36
(narrow def)	17+16	−0.00	1 95	0 17	−0 00
4 Rapid Access	47+17	$0 96	$2 29	3 75	3.36
(narrow def)	29+10	$0 70	$1 54	3 04	2.82
5 Arbitrage	26+10	$1.00	$2.64	2.67	2.28
6 Marketmaking	101+13	$3.32	$3.79	8.24	9.35

The data sources underlying the broad samples are listed in Table 1 The "narrow def." samples consist of the subset of data from sessions using only the given privilege and no other privileges.

A Nobs entry of $r + w$ means that excess profit for privileged traders was positive r times and negative w times in the sample. The signs test is computed directly from Nobs, and the t-test is computed in the usual manner from the mean and standard deviation (Std. Dev) of excess profit.

The privileges are detailed access to orderflow information (*book** = 2) and last mover (*extratime** > 0) in the Clearinghouse or call market institution; and detailed access (*book** = 2) and rapid access (*delay* * = 0) to orderflow information and the simultaneous buy low/sell high arbitrage (*arb**) and marketmaking (*post** = 1) in the continuous Double Auction market institution.

The last-mover privilege *extratime** > 0, on average, increases trading profit by about 25 cents per period (or about 10 percent) in both the broadly and narrowly defined samples. Although the point estimates suggest that the last move is the more profitable CH privilege, the intense noise in the excess profit data renders the result statistically insignificant.

Of the DA privileges, access to the complete orderflow (*book** = 2) seems the least profitable, indeed worthless in the narrow definition. (Again, two aberrant periods with very negative privileged trader profits account for the insignificance.) More timely access to orderflow information (*delay** = 0) was more profitable, averaging 70 cents or more per period, a very significant difference with all signs and t-statistics above 2.80. The additional privilege of arbitrage (*arb**) increases the average excess profit to $1.00 per period. Finally, privileged marketmaking (*post** = 1) is confirmed to be exceptionally profitable: excess profits are positive in 101 of 114 relevant trading periods and average $3.32, about three times the average of unprivileged trader profits.

Table 4 summarizes the impact of privileged traders on market performance. Beginning again with privileged access to CH orderflow information (*book** = 2) and the narrowly defined comparison (experiments 24 and 39 only), the table indicates slightly greater efficiency (lower RMSE, AIE, and Spread) and slightly lower trading volume in the 48 clearings in which privileged traders were present than in the 90 clearings with no privileged traders. The differences, however, are not statistically significant. Similarly, the last-mover privilege *extratime** has no

TABLE 4
Market Impact

Privilege [Data Sources]	Statistic	Performance Measures: RMSE	AIE	Spread	Volume
Clearinghouse					
1 Detailed Access	Nobs	48, 90	16, 30	48, 90	48, 90
[Exp 24, 39]	Means	12 7, 15 0	4.9, 5 2	10 8, 12 8	4 5, 5.2
	t	−0 76	−0 12	−0 80	−1 19
	Wilcoxon	−0 17	0 59	−0 38	−0.85
2. Last Mover	Nobs	96, 66	32, 22	96, 66	96, 66
[Exp 26–28]	Means	17.3, 26 7	5 8, 4.3	17 1, 18 0	5 2, 4 8
	t	−2.43	0 73	−0 28	0 57
	Wilcoxon	−1 10	0 89	−1.12	0 68
3. Improve Rule	Nobs	39, 53	21, 23	39, 52	39, 53
[Exp. 29, 38]	Means	15.8, 13 0	6 3, 4 8	13 8, 12 8	6 2, 4.9
	t	0 70	0 63	0 26	1.60
	Wilcoxon	0.10	0 70	0 20	1.27
Double Auction					
4 Detailed Access	Nobs	77, 75	33, 30	84, 85	84, 85
[Exp 17–18]	Means	21.5, 24 4	1 3, 3.6	37 7, 48.6	7 4, 6 4
	t	−1 10	−1 95	−2 83	1.15
	Wilcoxon	−1 54	−2 28	−3.04	1.59
5 Rapid Access	Nobs	111, 125	39, 42	140, 140	140, 140
[Exp 19, 37]	Means	22 1, 25 6	1 0, 3 3	32.7, 37.5	5 3, 6 0
	t	−1.58	−2 33	−2 20	−1 10
	Wilcoxon	−2 13	−2 86	−2 36	−1 28
6 Marketmaking	Nobs	139, 70	114, 42	144, 72	144, 72
[Exp 35–36]	Means	45 1, 34.8	5 0, 3.4	70 5, 41 2	16 0, 14.1
	t	1 96	1 36	6.45	1 45
	Wilcoxon	2 25	1 26	6 19	1 50

The data sources and privileges are presented in Tables 1 and 2, the main data sources are indicated here in brackets, e.g , [Exp. 24, 39]

RMSE, AIE, and Spread are measures of informational, allocational, and transaction cost inefficiency, respectively Volume is trading volume

Nobs *x*, *y* indicates that privileged traders are present in *x* of the observations in the relevant sample and absent in *y* observations The Means of the *x* and *y* observations are listed underneath the Nobs The standard (2 group) *t*-test and Wilcoxon test statistics are discussed in Section II.A of the text; positive entries indicate larger performance measures (greater inefficiency or trading volume) when privileges are present.

The privileges are detailed access to orderflow information (*book** = 2) and last-mover (*extratime** > 0) in the Clearinghouse or call market instituion; and detailed access (*book** = 2) and rapid access (*delay** = 0) to orderflow information and marketmaking (*post** = 1) in the continuous Double Auction market institutions The Clearinghouse improve rule allows all traders to increase the bid price (or decrease the ask price) automatically

effect (or perhaps a beneficial effect[9]) on market efficiency and has no effect on trading volume.

The third part of Table 4 examines the impact of allowing all traders (not just a privileged few) to improve orders at a single keystroke. The default rule *pull* = 1 was employed in 39 clearings and the improvement rule *pull* = 3 in 53 clearings of the relevant experiments 29 and 38. The data show insignificant enhancements in the efficiency measures and a possibly significant *decrease* in trading volume under the order improvement treatment.

[9]The null hypothesis is that market performance is unaffected by the presence of privileged traders, and the natural alternative hypothesis is that performance is impaired. Negative signs indicate the null hypothesis is rejected in the wrong direction, so the one-tailed significance levels do not apply.

The rest of the table deals with DA privileges. Market efficiency improves a bit in the presence of the orderflow access privilege *book** = 2; the gains in allocational efficiency and depth are significant even according to two-tailed tests. Privileged timely orderflow information (*delay** = 0) also appears to improve market efficiency, but some caution is in order.[10] The appropriate conclusion is that modestly delayed orderflow information to nonprivileged traders does not appear to impair market efficiency, even though it provoked many complaints from unprivileged traders. Preventing nonprivileged traders from posting bids and asks (*post** = 1) not only is unpopular with the nonprivileged but also is harmful to market efficiency. Informational efficiency and market depth are very significantly impaired (test statistics around 2.0 and 6.0, respectively) and volume and allocational inefficiencies also seem to increase somewhat.

The exceptional profitability and adverse market impact of the *post** = 1 privilege deserve closer examination. Experiments 35 and 36 employed 6 traders of each of two payoff types, with the number of privileged traders of each type ranging from 1 to 6. The data (available on request) show that excess profits, spread, and RMSE decline approximately linearly as the number of traders with posting privileges increases. Trading volume is depressed when privileged traders are few, but reaches a plateau when three or more traders of each type have posting privileges. Control observations from experiment 17 (three traders of each type) and experiment 21 (four traders of each type) give generally consistent results, suggesting that the key variable is the number of traders with posting privileges, not the number of unprivileged traders.

IV. Discussion

The main findings can be summarized briefly. The two main Clearinghouse (CH) privileges—last mover (*extratime** > 0) and detailed orderflow access (*book** = 2)—each increase profit a modest 5–10 percent. Excess profit for the last-mover privilege is larger, on average, but is less statistically significant because it has greater variance. Neither privilege appears to impair market efficiency; indeed, there is some evidence that each slightly *improves* informational efficiency. There is also weak evidence that the Wunsch improvement treatment (*pull* = 3) may increase market efficiency.

Of the three main Double Auction (DA) privileges—detailed orderflow access (*book** = 2), rapid access (*delay** = 0), and marketmaking (*post** = 1)—the second two are quite profitable. Quicker access provides a very significant average excess profit of at least 70 cents; when combined with automated price arbitrage (*arb**) the figure reaches $1.00, so traders with the combined privilege earn about 60 percent more than unprivileged traders. The marketmaking privilege was exceptionally profitable, allowing privileged traders to earn about three times more than the

[10]Only one follow-up experiment (numbered 37 in Table 2) was designed to compare market performance with and without the privilege. Consequently, the set of subperiods with the privilege was supplemented by the relevant first-round data (from experiments 8 and 9) and the comparison set was balanced using the relevant data from experiments 20 and 23. The choice was made a priori before running the test statistics and is conservative in the sense that the comparison set slightly oversamples the noisiest (3 traders $\times$ 3 clones SEQuential) environment. Hence, positive results (i.e., reduced market efficiency) would be quite significant, but the observed negative results are somewhat ambiguous.

unprivileged. Orderflow access was profitable overall, but not significantly so in the crucial follow-up experiments. Surprisingly, detailed orderflow access and rapid access did not impair market performance, and indeed appear to increase market efficiency somewhat. Restricting marketmaking to a few privileged traders clearly impaired market performance, especially informational efficiency and depth.

This paper is purely empirical, but a few interpretative remarks are in order. The lack of adverse market impact for most privileges is probably the biggest surprise in the results. Perhaps something here is akin to the revenue equivalence results of static one-sided auctions. As Vickrey (1961) first showed, the expected price in a simple first-price sealed-bid auction is the same as in the second-price auction because traders' optimal bidding strategies adjust to the institutional change. The result has been generalized considerably since then; see Bulow and Roberts (1989) for a recent exposition. In the complex continuous-time two-sided auctions studied here, traders' strategic adjustments may fully compensate for apparently adverse changes in the rules. That is, there may be no net effect on market efficiency even though there are distributional effects on trader profits.

The most dramatic finding reported here is the exceptional profitability and exceptionally adverse market impact of restricted marketmaking privileges; in this case, strategic readjustment evidently does not suffice to restore market efficiency. It is interesting to note that since it first awarded sole posting privileges to "specialists" some 120 years ago, the NYSE has undertaken a series of reforms and rule changes that reduce the adverse impact. Current rules include the imposition of an "affirmative obligation" on specialists to smooth price movements, awarding priority to public limit orders (an indirect form of posting privileges for nonspecialist traders), awarding nonspecialists access to the limit order book, and various "tick rule" limitations on specialists' transactions (NYSE Guide (1988)). An interpretation is that the rule changes were necessary because strategic adjustments were insufficient.

The design of current experiments was constrained by the need to build on previous laboratory studies. The current experiments relax the constraint on future experiments. Now it is possible to examine trading regulations in more detail, e.g., the effects of pairing privileges to obligations. It also is now possible to take a broader perspective on trading privileges, e.g., to compare a straight commissions regime to an efficient privileges regime, or to find the equilibrium price of freely purchased privileges. Another idea for future experiments is to vary market thickness, perhaps in a multimarket setting. The early history of the NYSE and some recent circumstantial evidence (*Economist* (July 6, 1991)) suggests that sole posting privileges may enhance efficiency in very thin markets. Laboratory research along these lines has the potential to improve our understanding of intermediation and marketmaking.

The results reported here are not in themselves an adequate basis for policy decisions, but they are practically important for two reasons. First, a conclusion established in laboratory markets (especially when consistent from first-round to follow-up experiments) should alter the presumption about the behavior of field markets (Grether and Plott (1984)). For example, the presumption (at least for reasonably thick markets) now should be that restricted marketmaking privileges are more likely to impair market performance (and to enhance privileged traders' profits) than modestly delaying orderflow information to less privileged traders. Second, the results presented here and in the companion paper (Friedman (1993)) provide a basis for policy-oriented studies. Exchange officials and regulators now can conduct a less costly and more reliable laboratory of field experiments to assess specific reform proposals.

References

Bulow, J., and J. D. Roberts. "The Simple Economics of Optimal Auctions." *Journal of Political Economy,* 97 (Oct. 1989), 1060–1090.

Columnist. "What am I Bid?" *Economist* (July 6, 1991), 83.

Copeland, T. E., and D. Friedman. "Partial Revelation of Information in Experimental Asset Markets." *Journal of Finance,* 46 (March 1991), 265–295.

Davis, D. D., and C. A. Holt. *Experimental Economics.* Princeton, NJ: Princeton Univ. Press (1992).

Demsetz, H. "The Cost of Transacting." *Quarterly Journal of Economics,* 82 (Feb. 1968), 33–53.

Domowitz, I., and T. Bollerslev. "Price Volatility, Spread Variability, and the Role of Alternative Market Mechanisms." *Review of Futures Markets,* (forthcoming).

Friedman, D. "How Trading Institutions Affect Market Performance: Some Laboratory Evidence." *Economic Inquiry,* 31 (July 1993), 410–435.

Friedman, D., and J. Rust, editors. *The Double Auction Market: Institutions, Theories and Evidence,* Sante Fe Institute Studies in the Sciences of Complexity XIV. Reading, MA: Addison-Wesley (1993).

Grether, D., and C. Plott. "The Effects of Market Practices in Oligopolistic Markets." *Economic Inquiry,* 22 (Oct. 1984), 479–507.

New York Stock Exchange, Inc. *Constitution and Rules.* Chicago, IL: Commerce Clearing House, Inc. (1988).

Plott, C. "Industrial Organization and Experimental Economics." *Journal of Economic Literature,* 20 (Dec. 1982), 1485–1527.

Plott, C., and S. Sunder. "Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models." *Journal of Political Economy,* 90 (Aug. 1982), 663-669.

Schwartz, R. A. *Equity Markets: Structure, Trading, and Performance.* New York, NY: Harper and Row (1988).

Smith, V. "An Experimental Study of Competitive Market Behavior." *Journal of Political Economy,* 70 (April 1962), 111–137.

________________. "Microeconomic Systems as Experimental Science." *American Economic Review,* 72 (Dec. 1982), 923–955.

________________. "Experimental Economics: Behavioral Lessons for Microeconomic Theory and Policy." Nancy L. Schwartz Memorial Lecture, Kellogg Graduate School of Management, Northwestern Univ. (1990).

Sunder, S. "Experimental Asset Markets: A Survey," GSIA Manuscript, Carnegie Mellon Univ. (March 1992), forthcoming in the *Handbook of Experimental Economics,* Kagel and Roth, eds. Princeton, NJ: Princeton Univ. Press.

Vickrey, W. "Counterspeculation, Auctions, and Sealed Tenders." *Journal of Finance,* 16 (1961), 8–37.