



American Finance Association

Partial Revelation of Information in Experimental Asset Markets

Author(s): Thomas E. Copeland and Daniel Friedman

Source: *The Journal of Finance*, Vol. 46, No. 1 (Mar., 1991), pp. 265-295

Published by: Blackwell Publishing for the American Finance Association

Stable URL: <http://www.jstor.org/stable/2328696>

Accessed: 25/11/2009 12:05

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=black>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Blackwell Publishing and American Finance Association are collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Finance*.

<http://www.jstor.org>

Partial Revelation of Information in Experimental Asset Markets

THOMAS E. COPELAND and DANIEL FRIEDMAN*

ABSTRACT

We develop a model of market efficiency assuming private information is partially revealed to uninformed traders via the behavior of those who are informed. This partial revelation of information (PRE) model is tested in fourteen computerized double auction laboratory markets. It explains the market value and allocation of purchased information, and asset allocations, better than either a fully revealing information model (FRE strong-form efficiency) or a nonrevealing expectations model; but it takes second place to FRE in explaining asset prices. We conjecture that refined versions of PRE may provide insight into “technical analysis” and minibubbles in securities markets.

FROM THE LARGE BODY of literature on informational efficiency beginning with Fama (1970), there now appears to be a general (but not universal) consensus that most important modern securities markets are at least semistrong form efficient but probably less than strong form efficient.¹ That is, all public information but probably not all private information is fully reflected in security prices. The question then becomes *when and to what extent* private information becomes incorporated, or, from the opposite perspective, *what is the value* of private information to the investor? Furthermore, as Latham (1985) and Rubinstein (1975) point out, theories of efficient markets should explain asset allocations as well as asset prices, and clearly allocations also depend on how the market incorporates private information. The empirical difficulty is that private information by definition is not contemporaneously observable in major financial markets, so further progress using existing market data is problematic.

Laboratory asset markets are a natural setting to study the issue of market efficiency because private information can be controlled and allocations can be directly observed. Early laboratory studies (e.g. Plott and Sunder (1982); Forsythe, Palfrey and Plott (1982); Friedman, Harrison, and Salmon (1984)

*Copeland is a partner at McKinsey & Co., Inc., New York. Friedman is at the Economics Department, University of California, Santa Cruz. We are indebted to the National Science Foundation (grant SES 8411382) and the UCSC Social Science Division for financial support, to Heikki Ketola, James Whitaker, and Christopher McRae for programming support, and to Gregory Allen, Alexander von Borries, and Peter Carr for research assistance. We are also indebted to Colin Camerer, Jack Michaelsen, René Stulz, Donald Wittman, and an anonymous referee for suggesting improvements in the exposition. The usual caveat applies.

¹Some observers regard anomalies such as the “January effect” as violations of semistrong efficiency. See Thaler (1987) for a recent discussion.

seem consistent with the general consensus cited above. More recently, Copeland and Friedman (1987) compare a strong form theory which we will call FRE (for fully revealed expectations, formerly called TRE for telepathic or true rational expectations; the idea is that somehow all private information immediately becomes public) to a semistrong form theory called NRE (for no revelation of expectations, formerly called ORE for ordinary rational expectations; the idea is that private information is never inferred by other traders, but in other respects all traders are fully rational). They find that FRE better explains the asset price data and (despite some anomalies regarding asset allocation and trading volume) offers a better overall explanation of the market data than NRE.

Although strong-form efficiency (FRE) explains prices better than alternative theories, it is lacking in some important dimensions. For example, if private information becomes public instantaneously, then prices will adjust to it without transactions. Hence FRE cannot explain the changes in asset allocations that we typically see when private information arrives. Furthermore, under FRE the value of privately held information is zero Sunder (1988). The NRE alternative remains unattractive because of its poorer empirical performance and also because the theory is based on naive behavior. The NRE model assumes that traders respond to private information using only knowledge concerning the possible final states of nature and that traders do not respond to trade-generated market signals that might reveal other traders' private information.

The purpose of the present paper is to introduce a model of partial revelation of information (PRE) that assumes traders can infer some privately held information from observable market signals. For example, a trader observing an increase in transaction prices, or even in the bid price, could infer that with positive probability some other traders received bullish private information. The PRE model is intended to bridge the gap between strong and semistrong theories of informational efficiency and to make specific, testable predictions regarding asset prices and allocations as well as the value and allocation of purchased information. By finding a specific form of PRE that explains all (or almost all) laboratory asset market data better than alternatives such as FRE and NRE, we hope to provide insight into the process by which private information becomes incorporated into asset prices in any asset market.

There is by now a substantial body of theoretical literature on information revelation. Grossman and Stiglitz (1976) sparked recent interest in the idea that equilibrium asset prices can in some circumstances fully reveal (i.e., be a sufficient statistic for) all private information. This insight was formalized in an abstract setting by Allen (1981) as a question of the invertibility of the equilibrium price correspondence. However, Jordan (1983) pointed out that more plausible dimensional assumptions make invertibility and full revelation very unlikely (i.e., nonfull revelation is generic).

Hellwig (1980) presented a more careful parametric model of the Grossman type to show that the equilibrium price typically is inefficient in aggregating information and that (as previously noted by other authors) the equilibrium

price is not fully revealing if there are more sources of uncertainty than prices. Consequently private information is valuable. Some of these points were amplified by Diamond and Verrecchia (1981) and Verrecchia (1982) under the rubric of “noisy rational expectations.”

Unfortunately this theoretical literature is restricted to a very special parametric example (joint Normal distribution for 2-dimensional random shocks and price, constant absolute risk aversion by investors, etc.) of an asset market run by a Walrasian auctioneer.² By contrast, actual asset markets including those created in the laboratory have no auctioneer to call out equilibrium prices before trade takes place and may have much different information structures. Consequently the existing theoretical literature provides no “off the shelf” model for analyzing experimental data, although it certainly provides guidance in conceptualizing and clarifying the issues.

We begin in Section I with a brief summary of the design employed in the fourteen experiments we report here. Section II introduces a theoretical framework that allows fairly general sorts of private information, and signals that are based on *observable* market behavior such as transaction prices. We derive a proposition that shows that even in our most complex laboratory environment there is a rational, logistically feasible process that in the absence of endogenous noise can fully reveal all private information. Thus, FRE is actually a special case of PRE. (So is NRE, but that will become obvious from the definitions.)

Section III derives a specific version of PRE which is distinct from both NRE and FRE, and shows how experimental data can be used to evaluate the competing theories. In Section IV we analyze the data, beginning with a qualitative overview. We present a battery of statistical tests that suggest that the market does provide reliable signals and that the specific version of PRE produces good price and allocation forecasts. Indeed, despite superior performance by FRE for asset price forecasts, PRE forecasts are at least as accurate as all known rivals in three important dimensions: (1) asset allocations, (2) the market price of purchased information, and (3) the allocation of purchased information. Furthermore, PRE predicts asset prices significantly better than NRE. We conclude in Section V that a refined version of PRE could well dominate in all dimensions and close with the conjecture that our findings may shed light on “technical analysis” and minibubbles in asset markets.

I. Experimental Design

A. Basic Features

Our fourteen experiments all employ a computerized continuous-time double auction (DA) program. The experiments consist of 12 to 20 trading periods, each period lasting 3 to 5 minutes. At each moment in a trading

²Perhaps the most relevant exception is Kyle (1986) who finds partial revelation of private information in the context of a call market. However, his results do not apply directly to most laboratory (or nonlaboratory) asset markets which are organized as double auctions.




Figure 1. Illustration of the trading screen in the computerized Double Auction Market. Features on the left half of the screen include state probabilities (0.50) and payout (\$1.50 or 0.50), the current market ("Best") and own ("Your Current") bid and ask prices, a ticker tape ("Last Transacted Prices") a space for entering prices ($\Rightarrow$), and reminders for special function keys (*q*, *e*, and *r*) for selling shares. The right side of the screen shows time remaining in the trading period (1 minute 28 seconds), the period number (5) and trader number (3), cash and share inventory, and reminder for special function keys (*u*, *i*, and *p*) for buying shares.

period, each participant can enter a bid (a statement of willingness to buy an asset unit for a specified amount of cash) or an ask (a similar statement of willingness to sell) from his interactive terminal, can use the terminal to accept the current best (highest) bid or best (lowest) ask tendered by his fellow traders, or can cancel his own outstanding bid or ask. Terminals are visually isolated from each other so that trading is anonymous. The computer serves primarily as a communications and record-keeping device and also enforces the rules. For example, transaction requests that would result in a negative cash or asset position are not executed but rather generate descriptive error messages. Trader confirmation is not required to execute a transaction, and trading can be brisk—we have observed as many as 42 transactions in a 5-minute trading period along with several times that number of unaccepted bids and asks. Bids or asks by different traders are queued as received, with the number tied for the best bid and best ask being publicly displayed. In addition to the best bid and ask, the trader's screen displays his inventory of cash and securities, an updated list of transacted prices (analogous to an exchange ticker tape), the amount of time left before trading stops, his own current bid and ask, and his possible end-of-period per share payouts. See Figure 1 for an illustration of the trader's screen display.

We recruited and trained two pools of about 25 subjects each, one pool consisting of MBA students at site M and the other of undergraduates at site U. Reported experiments all involve nine experienced subjects, each endowed

Table I
State-Contingent Payout Schedules for Traders of Each Type, in Cents^a

Experiment	Type 1		Type 2		Type 3	
	G	B	G	B	G	B
Exp2	200	30	170	80	120	100
Exp4	195	95	165	105	135	115
Exp5	205	105	175	115	145	125
Exp6	195	95	165	105	135	115
M3	205	105	175	115	145	125
M6	200	30	170	80	120	100
M7	200	30	170	80	120	100
M8	200	30	170	80	120	100
M9	275	105	245	155	195	175
M10	185	15	155	65	105	85
Info1	200	30	170	80	120	100
Info2	185	15	155	65	105	85
Info3	200	30	170	80	120	100
Info4	200	30	170	80	120	100

^aThe experiment names (Exp2, . . . , Info4) are arbitrary labels. In each experiment there are three traders of each clone type (Type 1, 2, and 3), and the two states G and B are equally likely.

initially with \$20 cash and three units (shares) of the asset per trading period.³ We randomly assign three traders (“clones”) to each of three payout schedules as shown in Table I. At the beginning of each of the trading periods (referred to below as repetitions or “reps”) in an experiment, each subject was informed that there is a 50–50 probability for each of two possible payouts for that rep. In Table I these are referred to as the good (*G*) and the bad (*B*) payout. For example, in experiment Exp2 the two possible payouts are \$2.00 or \$0.30 per share for Type 1 clones. (Before the end of the rep they are told which payout actually applies.) At the end of each rep the \$20 cash was reclaimed and subject profits were determined as the sum of (1) status quo profit = payout per share $\times$ endowed number of shares, plus (2) gross trading profit = (payout per share – purchase price) summed over share purchases + (sale price – payout per share) summed over shares sold. Subjects received these profits (or a stated fraction such as 50% of profits) in cash at the end of the experiment. Thus, we induced (information- contingent) asset values in the sense of Smith (1976) and potential gains from trade.

B. Treatment Variables

Several important environmental factors differ across our experiments. Subjects always know that they will costlessly receive a “news” message revealing whether their higher (*G*) or lower (*B*) payout will apply in that rep but are never told when the message will arrive. In our Seq (sequential)

³By “experienced”, we refer to a trader who has previously participated in a paid experiment involving at least 60 minutes of computerized asset market trading.

experiments, we send the messages to clone types sequentially in random order. The three clones of Type 2 might get their news at $t = 60$ seconds, those of Type 3 at $t = 120$, and those of Type 1 at $t = 180$ in a 240 second trading period. In a few of the experiments, we employ the alternative Sim treatment in which all traders receive news simultaneously.⁴

The content of the news message (G or B , i.e., the lower or the higher payout) for any given rep is the same for all traders in the Hom (homogeneous) experiments. The Het experiments feature a more complex 8-state environment that is heterogeneous in the sense that the payout (G or B with equal probability) is determined separately for each of the three trader types. Hence the set of states is $\{GGG, GGB, GBG, \dots, BBG, BBB\}$.

Eight of the experiments reported here feature a (noncomputerized) uniform price sealed bid auction before each rep for advance information on own payout. Before an asset trading period began, the three highest bidders anonymously received their news message at a price (deducted from profits) equal to the fourth highest bid; other traders received no news, only notification of the price of information. The trading period (including news messages) then proceeded as usual.

Figure 2 lays out the timing of events in the most complicated experiments which feature a market for information and sequential news messages. Table 2 summarizes the design of each experiment.

It may be worth noting that although the basic trading and profit calculation procedures were of course well known to our subjects, the parametric structures were never announced. That is, subjects were never told that some subjects have the same but others have different payouts and message arrival times, etc. In particular, subjects never had direct knowledge of whether or when their trading partners possessed superior or inferior information. We presume in our theoretical analysis that subjects came to (behave as if they) understand these matters after sufficient experience.⁵

II. Theoretical Models

A full theoretical analysis of rational behavior and strategic equilibrium in our markets (which feature uncertain payoffs and a news arrival process that creates information asymmetries) is far beyond the scope of this paper. Indeed, simple Double Auction markets (with neither of these features) still await definitive theoretical analysis.⁶ Given this gap in theory, experimen-

⁴For reasons explained in Section III A below, the Sim treatment is employed only in some experiments featuring an information market.

⁵Smith (1989) and other experimentalists have observed that data from repetitive experiments conducted with no public announcement of key parameters are often best explained by theories which counterfactually assume common knowledge. The results we report here are consistent with that observation. We note that such data offer much more convincing evidence for the "real world" relevance of common knowledge theories than would data from experiments in which the parameters are publicly announced in order to implement literally the theoretical assumptions.

⁶See Easley and Ledyard (1986) and Friedman (1984) for partial analyses and Wilson (1987) for a first attempt to construct a sequential equilibrium in bid/ask/acceptance strategies.




Figure 2. Time line for a typical sequential information experiment with an information market. This is the most complex design used. Simultaneous information experiments have only a single news event in each trading period (rep), and many experiments have no sealed-bid auction for advance information.

talists have explicitly or implicitly assumed simple reservation price strategies in deriving theoretical predictions in virtually all previous asset market studies. For the sake of tractability we generally follow that tradition here, although at times we touch on alternative strategies.

A. A General Framework

Let Z denote the set of all possible payout-relevant *states* of nature, with typical element z . Similarly, let $m \in M$ denote a private *message* and $s \in S$ denote a publicly observable *signal*. For simplicity (and implementability) we assume that Z , M , and S are finite sets. Let $[0, T]$ denote the time interval for a trading period and let $0 < t_1 < \dots < t_k < \dots < t_K < T$ denote the

Table II
Design Features of All Reported Experiments

Subjects at site U were undergraduates, and were MBA students at site M. Information arrival is either sequential (Seq) or simultaneous (Sim). Traders of all clone types receive either homogeneous (Hom) messages, i.e., all good or all bad, or they receive heterogeneous (Het) messages, e.g., good for Type 2 traders but bad for Types 1 and 3. Information market refers to whether (Y) or not (N) a sealed bid auction for information was conducted prior to the start of trading. Rep (short for repetition number) refers to the trading period, e.g., rep 9-16 means the ninth through the sixteenth trading periods.

Site	Experiment	Length (seconds)	Information arrival	Information content	Information market	Rep
U	Exp2	300	Seq	Het	N	1-8
U	Exp2	300	Seq	Hom	N	9-16
U	Exp4	300	Seq	Hom	N	1-8
U	Exp4	300	Seq	Het	N	9-16
U	Exp5	240	Seq	Hom	N	1-20
U	Exp6	240	Seq	Het	N	1-20
M	M3	300	Seq	Het	N	1-8
M	M3	300	Seq	Hom	N	9-16
M	M6	300	Seq	Hom	N	1-8
M	M6	300	Seq	Het	N	9-16
M	M7	300	Sim	Hom	Y	1-12
M	M8	300	Sim	Het	Y	1-12
M	M9	300	Seq	Hom	Y	1-12
M	M10	300	Seq	Het	Y	1-12
U	Info1	240	Seq	Het	Y	1-20
U	Info2	240	Seq	Hom	Y	1-20
U	Info3	150	Sim	Hom	Y	1-16
U	Info4	150	Sim	Het	Y	1-20

message arrival times.⁷ Finally, let $p(z)$ denote the final equilibrium asset price in state z . We assume in our exposition of theory that the sets Z , M , and S , the times t_k , and the prices $p(z)$ are all common knowledge for all traders, who are indexed $i \in I$.

Let m_k and s_k denote a message and a signal respectively, associated with the message time t_k . We assume that the joint probabilities $\pi(z, m_1, \dots, m_k, s_1, \dots, s_k)$ are common knowledge and, in particular, that traders know the conditional probabilities $\pi(z | m_1, \dots, m_k, s_1, \dots, s_k)$ and unconditional (or prior) probabilities $\pi(z)$. Thus traders can calculate conditional expected final equilibrium prices $E(p | m_1, \dots, m_k, s_1, \dots, s_k) = \sum_{z \in Z} p(z) \pi(z | m_1, \dots, m_k, s_1, \dots, s_k)$ for any $k \leq K$ as well as the unconditional expectation $E(p) = \sum_{z \in Z} p(z) \pi(z)$.

Following (implicit) tradition we proceed as follows. Each alternative model provides a specific reservation price for each trader in each subperiod (t_k, t_{k+1}) ; traders with sufficient cash will purchase shares and/or raise the bid at prices below their reservation level, and similarly traders holding

⁷For example, in our Seq experiments $K = 3$, so there are four subintervals. In the first, no traders have information, and all do in the last. In the middle two subperiods, traders are differentially informed.

shares will sell and/or lower the ask at prices above their reservation level. The result of such behavior is a price equal to the second highest reservation price among individual traders (which, given our use of clones, is the highest reservation price among trader types) with all shares held by the traders with the highest reservation price.⁸ Thus, each alternative model yields a price and an allocation forecast for each subperiod. For example, all models we consider forecast that in the final subperiod of our experiments, when every traders knows his own realized payout, the final equilibrium price $p(z)$ will be the highest realized payout and that all shares will be held by traders with this payout.

We begin by defining the NRE and FRE models.⁹ Let $\max_2$ denote the second largest element in a set, let $(m_{i1}, \dots, m_{ij})$ be the messages received by trader i up to time t , and let $(m_1, \dots, m_k)$ be the sequence of *all* messages received by one or more traders up to time t . Then $p_t(\text{NRE}) = \max_2\{E(p | m_{i1}, \dots, m_{ij}): i \in I\}$ is the *no revelation equilibrium* price and $p_t(\text{FRE}) = E(p | m_1, \dots, m_k)$ is the *full revelation equilibrium* price. The reservation price in each model is the expected final equilibrium price conditioned on news messages received personally (in NRE) or conditioned on all messages sent (in FRE).¹⁰

For example, with the NRE model a Type 1 trader in a Het environment can rule out four of the eight possible states of nature when she receives her news message, so her reservation price then is the expected final equilibrium price over the remaining four states. In the absence of a personal news message, however, she is assumed to ignore market signals and to stick to her unconditional expected final equilibrium price. By contrast, in the FRE model it is as if private messages somehow immediately become public.

Neither FRE nor NRE refers to market signals observed during trade. To investigate the possibility that signals are important but imperfect we make the following general definition. For given signal set S and the associated conditional probabilities we define the *partially revealed expectations* price as $p_t(\text{PRE}) = \max_2\{E(p | m_{i1}, \dots, m_{ik}, s_1, \dots, s_k): i \in I\}$, where $(s_1, \dots, s_k)$ are the signals publicly observed up to time t . Thus, in PRE traders again use the expected final equilibrium price as the reservation price but now conditioned on realized market signals as well as on private news messages.

If no informative signals can be found by traders—i.e., if the conditional probabilities are independent of $(s_1, \dots, s_k)$ —then PRE clearly reduces to NRE. Likewise, if perfectly revealing signals can be found—e.g., if $(s_1, \dots, s_k)$ is a sufficient statistic for $(m_1, \dots, m_k)$ —then PRE reduces to the other extreme case FRE. The next subsection explores this last possibility.

⁸An underlying reason is that induced demand is very elastic and induced supply is very inelastic in asset market experiments. See Copeland and Friedman (1987) and Section II B below for further elaboration.

⁹Previously introduced in Copeland and Friedman (1987) as ORE and TRE.

¹⁰Several authors have used expected own payout (rather than expected final equilibrium price) to define reservation prices in the so-called PI equilibrium. Given the generally poor predictive value of the PI model and its severe myopia in our context (implicitly it ignores important resale or repurchase opportunities), we will not discuss it further.

B. On the Rationality of PRE

An important theoretical issue is whether a specific PRE model yields a rational expectations equilibrium. The standard literature takes the signal set S to be the set of market-clearing prices and the associated conditional probabilities to be those obtained from Bayes Theorem using the objective likelihood function (arising in part from an exogenously specified noise process) and the objective prior probabilities. For markets organized as a double auction (or for any other realistic market mechanism) a broader view of the signal set must be allowed, but the question remains whether the conditional probabilities are consistent with Bayesian rationality.

To begin to answer this question we first consider reservation price strategies and some alternatives more closely. Say that $R_i(t)$ is the *single reservation price* for trader i if she accepts (refuses) opportunities to sell shares at time t at prices above (below) $R_i(t)$ and refuses (accepts) similar opportunities to buy shares. Refer to such a trader as *risk-neutral* if $R_i(t)$ is the expected final equilibrium price conditioned on all messages and signals available to her at time t . Finally, refer to such a trader as *aggressive* if she undercuts (raises) the existing best ask (bid) if held by another trader whenever she can do so at a price exceeding (below) her risk-neutral reservation price $R_i(t)$. It is intuitively clear that an aggressive trader maximizes her expectation of final wealth as long as she (1) believes that no other trader has superior information and (2) believes that she can not systematically affect her future terms of trade by her current actions.¹¹

At the opposite extreme, a trader will not wish to trade with another who is believed to possess superior information, due to the well-known adverse selection problem discussed in related contexts by Copeland and Galai (1983) and Glosten and Milgrom (1985). Such a trader then may wish to pursue a *withdrawal strategy*, defined as the refusal to accept opportunities to buy or sell, and the setting of her own bid (ask) below the minimum (above the maximum) possible final equilibrium price given her current information.

We are particularly interested in signals that are created by traders responding to messages. An aggressive trader, upon receiving a message which alters her conditional expectation, will immediately take observable actions such as raising the bid. If several aggressive traders receive messages at the same time, it may take a while for the best bid to stabilize. We assume that there is some length of time ε which suffices for aggressive traders to fully interact after one or more of them revise their reservation prices. Since we do not allow the interaction process to cause further revisions, it is reasonable to assume ε is small relative to the interval between message times. Thus we assume $0 \leq \varepsilon < \min \Delta t_k$ and refer to this assumption by saying that the *calibration interval is short*.

¹¹See Friedman (1984) and (1987) for some justification of this claim in a related context. Two qualifications apply. Nonnegativity constraints on share and cash holdings should be taken into account in the obvious way. And if the common knowledge assumptions, (e.g., for $\pi(z)$ and $p(z)$) are relaxed, then in general a trader does better with separate reservation prices for buying and selling.

PROPOSITION 1: Assume:

- (1) $E(p | m_1, \dots, m_k)$ is a 1:1 function of $m_k \in M$ for each $k = 1, \dots, K$;
- (2) It is common knowledge that the message m_k is (privately) received by at least three traders at each message time t_k , $k = 1, \dots, K$;
- (3) Each trader pursues an aggressive strategy except when some other trader has superior information, in which case she reverts to a withdrawal strategy; and
- (4) The calibration interval is short.

Then there is a set S of observable signals and associated conditional probabilities so that PRE prices and allocations coincide with FRE, except perhaps during the calibration intervals $[t_k, t_k + \varepsilon]$.

Proof: Let $S_k = \{E(p | m_1, \dots, m_k) : m_j \in M, j = 1, \dots, k\}$ and let $S = \bigcup_{k=1}^K S_k$, with the convention that a signal s_k is observed at time $t_k + \varepsilon$ as the average of the best bid and ask. By assumption (1), for each $s \in S_k$ there is a unique $m_k = \phi_k(s) \in M$. Define the conditional probabilities on S from the (given) probabilities on M by $\pi(z | m_{i1}, \dots, m_{ik}, s_1, \dots, s_k) = \pi(z | m_{i1}, \dots, m_{ik}, \phi_1(s_1), \dots, \phi_k(s_k))$.

The resulting PRE works as follows: at the first message time t_1 , the traders receiving the message $m_1 \in M$ will reset their reservation prices to $E(p | m_1)$. Pursuing an aggressive strategy (while other traders pursue withdrawal strategies) they will drive the best bid and ask to $E(p | m_1)$ by the time $t_1 + \varepsilon$ by proposition assumptions (3), (2), and (4).¹² Other traders then observe the resulting signal $s_1 = \phi_1(m_1)$ and by assumption (1) obtain the same conditional probabilities (hence expectation and reservation price) as the informed traders, who thus no longer possess superior information. Hence, $p_t(\text{PRE}) = p_t(\text{FRE})$ and similarly for allocations, for $t \in [t_1 + \varepsilon, t_2]$. The same argument gives the same conclusion for $[t_2 + \varepsilon, t_3], \dots, [t_K + \varepsilon, T]$. Of course, the equilibrium concepts coincide for $t \in [0, t_1]$. Q.E.D.

The idea simply is that traders learn to anticipate the information arrival times t_k and withdraw briefly from the market if they do not receive a message. Absent noise, the bid/ask behavior of three informed traders will fully reveal their price expectation and thus their private information during the “calibration” interval of length ε . At that point it is as if the messages were public.

The assumptions driving this result are strong but not ridiculous. The 1:1 function between conditional expectations and states of nature, which plays a role similar to the invertibility assumption in the fully revealing rational expectations theory, may seem very special because it might easily be the case that m_k could be replaced by some $\hat{m}_k \neq m_k \in M$ without affecting $E(p | m_1, \dots, m_k)$. If, for each fixed k such a substitution also has no effect on $E(p | m_1, \dots, m_k, \dots, m_{k+h})$ for $h \leq k - K$, then the proof becomes

¹²Perhaps the role of the withdrawal strategy in revealing “bad” news should be underlined. If uninformed trades do not withdraw (as implicitly assumed in NRE), then they will not detect “bad” news. Such traders would still detect “good” news which leads informed traders to raise the best bid and ask.

messier but still goes through. The latter version of the assumption holds for all our laboratory environments. The lack of information monopolists in the second assumption is a feature of our basic laboratory asset markets. It also helps reconcile the first assumption, which implies that traders' actions reveal their private information, with the last assumption, which is based on the trader's belief that his own actions do not change others' behavior. If it is common knowledge that each realized message m_k is privately received by at least three traders, then each of them can correctly assume that the actions of the other two will suffice to reveal it.¹³

A formal proof that FRE is a rational expectations equilibrium is beyond the scope of this paper, but PROPOSITION 1 (together with our assertion that the postulated aggressive and withdrawal strategies are optimal) are key ingredients. We make no such claim for NRE. In implementing NRE we use the objective state probabilities conditioned on news messages and the true state-contingent final equilibrium prices, so traders are rational in the sense of Muth. However, as Copeland and Friedman (1988, page 23) implicitly recognize, a type of "winner's curse" arises from aggressive strategies in NRE—an uninformed trader holding the market bid will tend to lose money when better informed traders are present, but NRE assumes that uninformed traders never learn to avoid such behavior. Consequently NRE is not a rational expectations equilibrium. Unless signals are perfectly revealing, the same problem will persist (albeit in attenuated form) in our PRE implementations. Evidently separate reservation prices for buying and selling or even more complex strategies would be required to overcome this problem.

A final theoretical issue concerns endogenous noise. A rational trader with superior information might find it advantageous to deliberately confuse the price signal (e.g., by bidding above his expected final equilibrium price), but if others (clones) share his information, such tactics seem unprofitable. It is hard to see any rationale for less well-informed traders to deliberately introduce noise. Yet we have observed considerable noise (i.e., bids, asks, and transactions that appear to sacrifice profit opportunities) in every double auction experiment we have examined. Whether such noise is due to human error or represents learning or subtle strategic behavior, it certainly must be taken into account in interpreting market signals. The derivation of FRE presented above assumes this problem away (in that traders flawlessly pursue aggressive or withdrawal strategies), and its logic suggests that full revelation is hard to avoid in the absence of noise.¹⁴ One could treat noise as an exogenous process to be estimated from the data, but at best there is a

¹³See Kyle (1986, Theorem 7.5) for a related result based on similar intuition in the context of a call market. One of the present authors conjectures that, given the practice in our experiments of sending simultaneous messages to 3 clones of a given type at each information time, the common knowledge assumption for the information arrival times is not necessary in PROPOSITION 1.

¹⁴Experiments can be designed to introduce noise or its strategic equivalent. For example, if the news arrival times or presence of insiders is uncertain, then simple nonfully revealing versions of PRE are possible. See von Borries and Friedman (1988) for an example.

delicate interplay between a noise process and the corresponding version of PRE in rational expectations equilibrium.

These theoretical difficulties, together with the practical necessity of avoiding free parameters, dictate that versions of PRE to be tested in our experimental environment as alternatives to NRE and FRE must be less than fully rational. Our goal in the next section will be to construct a simple but robust implementation whose predictive power will be representative of potentially more refined versions of PRE.

III. Parametric Models and Forecasts

We first specialize the general framework of Section II to our Het:Seq experimental environment. There are 8 states in $Z = \{z_1 z_2 z_3: z_i \in \{B, G\}\} = \{BBB, BBG, \dots, GGG\}$ and 6 messages in $M = \{iy: y \in \{B, G\}, i = 1, 2, 3\} = \{1B, 1G, \dots, 3B\}$, where i now indexes clone *types* rather than individual traders. The prior probabilities $p(z)$ are all $1/8$. There are 3 message arrival times (i.e., $K = 3$), at $t_1 = 60$ sec, $t_2 = 120$, and $t_3 = 180$ for our experiments with 4 minute ($T = 240$ sec) reps. We arbitrarily (and a priori) assume a calibration interval of length $\varepsilon = 20$ seconds, although one might argue for a figure anywhere between 5 and 30 seconds for our computerized double auction. The empirical significance of the calibration interval is twofold: specific PRE forecasts will depend on events taking place during the interval and forecast comparisons will be restricted to times outside the interval.

Since signals are endogenous, one must consider all publicly observable events over the interval $[t_k, t_k + \varepsilon]$ that might be triggered by private messages as possible candidates for market signals. We noticed even in our earliest training experiments that traders usually fix their gaze on the best bid/best ask display on their screens (see Figure 1 and Section I A above). The simplest signal set we could think of uses the display to distinguish “upticks” from “downticks” as follows. An accepted best ask indicates buyer initiative; count each such event as $+1$. Count an accepted best bid as -1 for the analogous reason. Sum over the interval $(t_k, t_k + \varepsilon)$. If the result is positive, call it an uptick; if negative call it a downtick; and if zero call it a nulltick. These tick signals may be of some interest in their own right, but we use them mainly as an ingredient of our PRE implementation.

The signal set $S = \{0, \dots, n\}$ for our parametric PRE model partitions the price line into intervals, say $I_0, \dots, I_n$, corresponding to possible news messages. The signal s_k for message arrival time t_k then is $j \in S$ if the observed price at $t_k + \varepsilon$ lies in interval I_j . The “observed price” is the best ask (best bid) if most transactions have been at that price since the message arrived, i.e., given an uptick (downtick). Given a nulltick, the observed price is the average of the bid and ask prices.

The next step in constructing the parametric PRE model is to specify the conditional probabilities, and here we face a dilemma. If we allow calibration of these probabilities in the light of actual experience, we may obtain accuracy and realism but must deal with several deep theoretical issues (e.g.,

concerning Bayesian learning in a disequilibrium environment) and many free parameters in the data analysis. For present purposes we prefer an *a priori* specification using the same basic concepts as the NRE and FRE models as defined in Copeland and Friedman (1987). Hence, we assume that information can always be summarized by a subset A of Z that identifies the states that remain possible, states $z \notin A$ are ruled out by the information. That is, the conditional probabilities can always be expressed in the form $\pi(z | A) = \pi(z) / \pi(A)$ if $z \in A$ and $= 0$ if $z \notin A$, where $\pi(A) = \sum_{z \in A} \pi(z)$.

We now summarize the information content of message $m = iy$ (recall that i is the trader type and y refers to the realized payout G or B) by the set $A(m) = \{z = z_1 z_2 z_3 \in Z : iy = z_i\}$, i.e., states inconsistent with the message are ruled out. Then $\pi(z | m)$ is defined as $\pi(z | A(m))$. This definition works well for message sequences because $A(m_1, \dots, m_k) = \bigcap_{j=1}^k A(m_j)$ enforces the natural requirement that a state that remains possible if and only if it is consistent with all messages received.

It is not difficult to extend this approach by means of a set-valued function, or correspondence $B: S \rightarrow M$ that assigns to each signal s the set $B(s)$ of possible messages that might have triggered that signal. The information content of s is then summarized by $\hat{A}(s) = \bigcup_{m \in B(s)} A(m)$. Thus the conditional probabilities in our parametric PRE model are

$$\pi(z | m_1, \dots, m_k, s_1, \dots, s_k) = \pi\left(z | \bigcap_{j=1}^k A(m_j) \cap \hat{A}(s_j)\right).$$

Table III spells out the parametric PRE model for Het:Seq experiments, with numerical values computed for Exp2 parameters from Table I. The conditional expectations $E(p | m)$ adjusted by a 3¢ tolerance (chosen *a priori* to reflect a minimal level of background noise) define the endpoints of the intervals I_j . The messages $B(s)$ corresponding to the signal $s_k = j$ are simply those that could have generated a price at least as extreme as the observed priced P . For example, the second most bullish signal $s = 2$ arises when P exceeds the reservation price for Type 3 traders with good news by at least the 3¢ tolerance, but does not exceed that for Type 2 traders. The interpretation is that Type 1 or 2 traders must have received good news.¹⁵

¹⁵We made three further conventions in empirical work: (a) since signals are endogenous, traders who receive a private message at t_k ignore the corresponding signal; (b) elements $m = iy \in B(s)$ are ignored by traders of type i when inconsistent with their own experience (as happens occasionally), and (c) in the rare case that the intersection is empty (all states ruled out), then the oldest signals are ignored.

For comparative purposes, we also implemented the underlying “tick rule” whose signal set $\{u, d, 0\}$ distinguishes only upticks, downticks and nullticks. Its signal correspondence is simply $B(u) = \{1G, 2G, 3G\}$, $B(d) = \{1B, 2B, 3B\}$ and $B(0) = M$. That is, an uptick (downtick) signals that some clone type got good (bad) news, and a nulltick (0) eliminates no states and therefore provides no information. Roughly speaking, the tick rule is a crude signal based on buying or selling “pressure” and is intended to distinguish nonspecifically between “good” and “bad” news.

Table III
The Price Rule Signal Correspondence

Variable P refers to the observed price, and p refers to the final equilibrium price. The observed price is compared to ranges of expected final equilibrium prices to establish the signal s . The ranges use an arbitrary tolerance parameter of 3¢ which was chosen a priori. $B(s)$ is the set of possible messages that might have triggered signal s . M means no signal. 1G, for example, means that the signal implies that clone Type 1 received good news. Trader types are indexed so that news is more informative for types with lower indices, i.e., $E(p|iG) \geq E(p|jG)$ when $i < j$. The numerical example [in brackets] is taken from the payouts in the first row of Table 1.

s	Signal	Definition
1	$E(p 2G) + .03 < P < \infty$ [188 < P < ∞]	{1G}
2	$E(p 3G) + .03 < P \leq E(p 2G) + .03$ [175.5 < $P \leq 188$]	{1G, 2G}
3	$E(p) + .03 < P \leq E(p 3G) + .03$ [173 < $P \leq 175.5$]	{1G, 2G, 3G}
4	$E(p) - .03 \leq P \leq E(p) + .03$ [167 ≤ $P \leq 173$]	M
5	$E(p 3B) - .03 \leq P < E(p) - .03$ [164.5 ≤ $P < 167$]	{1B, 2B, 3B}
6	$E(p 2B) - .03 \leq P < E(p 3B) - .03$ [152 ≤ $P < 164.5$]	{1B, 2B}
7	$0 < P < E(p 2B) - .03$ [0 < $P < 152$]	{1B}

To summarize, our parametric PRE model employs a partition of bid/ask prices that is cruder than but analogous to that employed in the fully revealing version of PRE in PROPOSITION 1. The rationale is that in practice signals may be too noisy for full revelation.

A. Forecasts

It is now a straightforward exercise to generate the PRE reservation prices $E(p|m_{i1}, \dots, m_{ik}, s_1, \dots, s_k)$ for the realized message sequence and price history of any experiment. Using the definitions presented in Section II. A, one then obtains asset price and allocation forecasts for NRE, FRE, and PRE, all of which can be compared to the actual prices and allocations. Since all the forecasts agree in the first and last subperiods $[0, t_1]$ and $[t_3, T]$, comparisons are meaningful only for the two intermediate subperiods $[t_1 + \varepsilon, t_2]$ and $[t_2 + \varepsilon, t_3]$ of our Seq experiments, when traders are heterogeneously informed. In our experiments with no information market, the Sim environment allows no distinction between NRE, FRE, and PRE. Since FRE recognizes no information asymmetries, it makes no asset allocation forecasts except for the final subperiod, as noted in Copeland and Friedman (1987).

Our criterion for comparing predicted with actual asset prices is root mean squared error (RMSE) over the subperiod; a lower RMSE means a better

forecast and is favorable evidence for that model. (We also note that observed prices below forecast might be due to traders' risk aversion or subjective transactions costs and therefore are less damaging to a model than observed transaction prices above forecast.) For asset allocations, we look at the number of shares actually held at times t_2 and t_3 by traders of the types forecast *not* to hold them under that model as a percentage of total shares. Other things equal, a lower misallocation percentage is favorable to the corresponding model.¹⁶

Another important set of forecasts concerns the market value of information. Recall from Section I.B that a sealed-bid auction for information was conducted before each asset trading period in eight of our experiments. The value of purchased information to a trader depends on the realized state of nature as well as on the types of other informed traders. For each possible realization one can compute the NRE asset price (assumed to be the second highest risk-neutral reservation price among traders), the optimal strategy for an informed trader, and the expected gains relative to an optimal uninformed strategy. Then one computes the fourth highest expected gain in Nash equilibrium, and this is the NRE information price forecast.¹⁷ FRE simply predicts that the value of information is zero.

The logic of PRE calculations is the same as for NRE, but of course there are fewer (and smaller) profit opportunities for information purchasers. Indeed, if market signals perfectly reveal purchased information (as in FRE), then the information has no market value. For concreteness in the PRE forecasts reported in Table IV we adopt the convention that information purchasers anticipate that the market will produce the least precise PRE signals ($s = 3, 4$, or 5).¹⁸

For example, refer to the parameters for experiment Info4 in Table I and consider the crucial case when two Type 1 traders purchase information which reveals that they will receive the lower payout. The presumed signal is $s = 5$ which allows all other traders to eliminate state *GGG*. An informed

¹⁶In the interest of brevity we omit tests of some less definitive asset market outcomes. From allocation forecasts one can derive trading volume forecasts as in Copeland and Friedman (1987). However, these forecasts are trivial for FRE and usually are in close agreement for PRE and NRE. The *timing* of trade could also be examined; our PRE model suggests that volume will be concentrated in the moments following the receipt of market signals, i.e., immediately after $t_k + \varepsilon$. Histograms of transaction times could be used to check the plausibility of our a priori choice of $\varepsilon = 20$ seconds. Finally, one often analyzes profit data with summary measures such as total profits earned across all traders in a rep as a proportion of the maximum possible. This measure is called efficiency and actually is an alternative view of the final allocation forecast. Since all three models predict the same final allocation, we will omit this measure here.

¹⁷This approach is adapted from Copeland and Friedman (1988) which should be consulted for details.

¹⁸If signals were more precise one would obtain PRE forecasts between those given for FRE and PRE in Table IV. If one assumed that some trade occurs in the calibration intervals, then one would obtain slightly higher forecasts. Consideration of purchase and resale strategies extending across subperiods also can raise the NRE and PRE forecasts slightly in the Seq environments. The "+" marks in Table IV are intended to suggest this last possibility. However, we think that the given PRE forecasts are representative.

Table IV

Forecast Price (and Allocation) of Information Purchases

Prices, denominated in dollars, are Nash equilibrium (NE) forecasts of the fourth highest bid in a sealed bid auction to purchase information. The plus (+) refers to the theoretical possibility of slightly higher values for experiments in which information arrives sequentially (Seq) than for experiments in which information arrives simultaneously (Sim). Allocations are forecasts of the trader types that purchase the information in NE. For example (1, 1, 1, or 2) means that the purchasers in NE are 3 Type 1 traders or 2 Type 1 traders and a Type 2 trader; (all) means that any combination of purchasers is possible in NE.

Experiment	NRE	FRE	PRE
M7	\$0.75 (all)	\$0 (all)	\$0 (all)
M8	0.45 (1, 1, 1 or 2)	0 (all)	0.39 (1, 1, 2)
M9	0.75 + (all)	0 (all)	0 (all)
M10	0.45 + (1, 1, 1, or 2)	0 (all)	0.39 + (1, 1, 2)
Info1	0.45 + (1, 1, 1 or 2)	0 (all)	0.39 + (1, 1, 2)
Info2	0.75 + (all)	0 (all)	0 (all)
Info3	0.75 (all)	0 (all)	0 (all)
Info4	0.45 (1, 1, 1 or 2)	0 (all)	0.39 (1, 1, 2)

Type 1 trader can also rule out states *GGB*, *GBG*, and *GBB*, while an informed Type 2 trader who also receives $m = B$ can eliminate states *GGB*, *BGG*, and *BGB*. Hence, the latter will optimally sell his three shares at the resulting PRE price of \$1.657 in the first subperiod rather than at the expected final equilibrium price of \$1.10. His net gain in this realization (which has probability $0.5^2 = 0.25$) will be approximately \$0.557/share and will be 0 in the other realizations, so his expected gain from information purchase is about $(\$0.557/\text{share}) \times (3 \text{ shares}) \times (0.25) \approx 42\text{¢}$. The corresponding expected gain for a third Type 1 trader is $(0.257/\text{share}) \times (3 \text{ shares}) \times (0.50) \approx 39\text{¢}$. Consequently PRE forecast in Table IV for Het environments is for two Type 1 and one Type 2 traders to purchase at a price of about 39¢, the best rejected bid by the other Type 1 trader.

In the corresponding NRE calculations, both Type 1 and Type 2 traders have expected gains of 45¢, so the allocation forecast is a bit less precise, and the price forecast is slightly higher. Hom experiments involve a two state world, with $Z = \{GGG, BBB\}$. The signal sets then have at most 3 elements, signaling the first state or second state or a null signal. If a null signal is received, PRE forecast responses are the same as those for NRE; otherwise, PRE forecasts the same responses as FRE. The entries in Table IV presume that nonnull signals are received.

To ensure valid comparisons of predictive power, the forecasts in Table IV and the numerical forecasts for asset price and allocation were all computed before beginning the data analysis. In the next section we use all experiments with an information market (the last eight experiments listed in Table II) to test the competing model forecasts of information value (or price) and information purchase (or allocation). We use all experiments with no infor-

mation market (the first six in Table 2) to test the forecasts of asset price and asset allocation.¹⁹

IV. Results

A. Data Overview and Qualitative Results

Our main output files contain one line for each action (bid, ask, accept, cancel) taken by each participant in each rep, yielding over 2,000 lines per experiment which provide a continuous record of market prices and asset allocation. The input files contain all payout and message parameters. For the information market experiments we also recorded all bids by all subjects in all reps, the resulting information prices and allocations, and trading profits. This mass of raw data was reduced by stages. The most relevant aspects of the raw asset market data are summarized in time graphs as in Figure 3. The outcomes of the sealed-bid information auction are graphed in Figure 4. Further stages of data reduction will be described in the next section.

We begin our discussion with Figure 3, displaying the bid and ask prices (lower and upper lines) and transaction prices (stars) for a few trading periods in experiment Exp6. In the first rep (i.e., trading period), the first news message was sent to Type 1 traders at $t = 60$ sec and contained the "Bad" news that the lower payout applied that round; this is summarized by the 1B notation in Figure 3 near the upper end of the vertical line from the 600 tenth-second mark in Panel 1. The /D7 indicates that response to this news produced a downtick (*D*) and the lowest possible PRE signal (7), so all traders were able to rule out states *Gxx* to obtain reservation prices (and hence PRE price) \$1.40 for the second subperiod. The allocation at the time the news went out was 8 shares held by Type 1 traders, 11 by Type 2, and the remaining 8 by Type 3, as indicated near the bottom of the same vertical line. The news apparently had little immediate effect on the already depressed transaction prices, and the bid and ask prices maintained the same narrow trading range until a little after the next news event, indicated as "Good" news to Type 2 traders (2G). At that point the ask begins to move up, eventually followed by the bid, with relatively heavy transaction volume. The uptick was appropriate, but the PRE signal (6) remained too low. Evidently Type 2 traders did most of the purchasing over this time period as their holdings increased from 13 to 23 shares. The last news message of 3G

¹⁹The experiments with no information markets of course provide no evidence on information price and allocation. Asset data from the information market experiments exist but are much more difficult to analyze and interpret, chiefly because information monopolists are sometimes present (see Sunder (1988) for some discussion of this issue) and several new conventions have to be introduced to define the NRE and PRE forecasts. Since the evidence is consequently weaker and would require a more lengthy exposition, we decided to omit it. Nevertheless, in the interest of completeness we recently performed a basic analysis of the information market asset data. The results (available on request) are consistent with those reported in Section IV.

had no effect on the market, and indeed Type 2 traders purchased the remaining shares in the last subperiod at prices within a generally narrow bid-ask trading band around the final equilibrium price of \$1.65.

One can generate similar market commentary for all the other reps and can form impressions of the relative quality of the three forecasts by comparing actual transaction prices to those predicted in the middle two subperiods and by comparing the actual allocations to those forecast.²⁰ The dashed lines indicate the parametric PRE forecasts, while the NRE and FRE forecasts are indicated by dotted and dash-dotted lines, respectively. Before we turn to statistical tests in the next subsection, we note that the “interocular trauma test” seems to tell us that PRE usually yields better price forecasts than NRE and allocation forecasts of about the same quality. By default, PRE allocation forecasts are better than those of FRE, and the price forecasts usually seem of about equal quality.

However, PRE seems to produce occasional large errors. For example, in rep 10 of Exp6, the first news event is 3B (which should have little effect), but the subsequent observed price is \$1.84, so $s = 1$ signifying 1G! From Figure 3, one can see that this seemingly strong signal actually arises from a weak market: a wide bid/ask spread (bid $\approx$ \$1.50, ask $\approx$ \$1.80) persists with only one transaction (a buy) during the calibration interval $(t_1, t_1 + \varepsilon) = (60, 80)$. Had this transaction not occurred, or had ε been chosen to be 3 seconds longer, the signal would have been the neutral $s = 4$. Nevertheless, our mechanical PRE forecast treats the realized signal exactly the same as one that arose from a strong buying surge that carried transaction prices upward. Hence, Type 3 traders (who know they got bad news) and Type 1 traders (who know they didn’t get good news) are forecast to sell to badly fooled Type 2 traders who believe that Type 1 traders will eventually repurchase from them at a price approaching \$1.95. Of course, traders did not actually behave in this manner, so the PRE forecast is far from the mark in this and several similar cases.

The rest of this subsection examines the outcomes of the sealed-bid information market. Note from Table IV that the information allocation forecasts for PRE are essentially the same as those for NRE, and both differ from the vacuous FRE forecast of information allocation. The PRE and FRE models forecast a zero price of information in the Hom experiments (M7, M9, Info2 and Info3), but in the remaining Het experiments the PRE price forecasts are much closer to those of NRE.

Figure 4 shows the observed price of purchased information in the eight experiments that feature information auctions. Except perhaps for M7, a Hom:Sim experiment, the prices in the site M information auctions seem to converge to near NRE levels or (in the case of the Het experiments) to the very similar PRE levels. On the surface, then, NRE and PRE do about equally well in explaining the site M data. However, after only 12 reps it is

²⁰ Exp6 is perhaps especially useful in distinguishing PRE and FRE because it runs 20 reps all of the same regime. Time graphs for all reps of all experiments are available on request.




Figure 3. Price-Time graphs for selected trading periods (reps) in Experiment Exp6. The upper-step function is the market ask price, the lower is the market bid price, and stars indicate transaction prices. Dashed lines (----) indicate theoretical forecast prices assuming partial revelation of private information (PRE), dotted lines (.....) assume no revelation (NRE), and dash-dotted lines (-.-.-) assume full revelation (FRE). Vertical lines indicate news events, the message and signal (e.g., 1B/D7) noted at the top of a line, and the asset allocation at that time (e.g., 8, 11, 8 shares held by traders of Type 1, 2, and 3, respectively, at the time of the first news event) at the bottom of the vertical line.





Figure 3—Continued




Figure 4. The market value of information. The market value of information is measured as the clearing price (the 4th highest bid) in sealed-bid auctions in every period (rep) in every relevant experiment. Panel A shows prices for the Heterogeneous (Het) experiments M8, M10, Info1, and Info4. Panel B shows prices for the Homogeneous (Hom) experiments M7, M9, Info2, and Info3. Theoretical forecast values for information for the no revelation (NRE) and partial revelation (PRE) models are indicated by dashed (---) and dotted (····) lines, respectively. The model of full revelation of private information (FRE) forecasts an information price of zero in both environments. The PRE also forecasts zero price in the Hom environment. The names of the experiments (e.g., M7, Info2) are arbitrary labels.

not clear where the price might eventually end up after more replications. Therefore, we ran the site U experiments for the maximum feasible time, 20 reps in most cases. One can see that the price of information in the site U Hom experiments (Info2 and 3) lies everywhere below the price in the corresponding Het experiments and does end up very close to 0 at 5–6¢ in the

final reps. On the other hand, the Het experiments yield information prices that persist well above zero. In Info1, the price seem stabilized at the 20–25¢ level, with one Type 2 and all three Type 1 traders purchasing in the last five reps. Apart from perhaps overly conservative conventions chosen in Section III. A, the bid data (available on request) seem to be in close agreement with the PRE forecast. Similarly, the bid data confirm the convergence of the PRE forecast of zero price and random allocation for the Info2 and Info3 information markets by the last few reps. Info4 shows signs of very slow convergence as information purchasers persistently earned lower net trading profits than nonpurchasers only in this experiment, suggesting a price that has further to fall. Nevertheless, the last two reps have outcomes (prices of 70¢ and 66¢, only one of six purchases by the “wrong” type of trader) that are roughly consistent with either NRE or PRE and quite inconsistent with FRE.

To summarize, it seem fair to say that the information market data are all reasonably consistent with PRE, while the Het data are definitely inconsistent with FRE and the Hom data are definitely inconsistent with NRE. PRE forecasts that the value of purchased information will fall to zero in the Homogeneous environment where trading behavior unambiguously signals private information but forecasts a positive value to information in the Heterogeneous environment because there the signal-message correspondence is not precise and private information is only partially revealed. In the Het case, specific trader types are forecast to purchase information. The data seem to bear out these PRE forecasts.

B. Statistical Analysis

An important preliminary question is the extent to which the PRE signals defined in Section III. A actually convey the information hypothesized. For example, when there is an uptick (more boughts than solds immediately after a news message), how often was it the case that the news actually was good (1G or 2G or 3G)? Basically one wants to see whether such “right” calls significantly outnumber the “wrong” calls (e.g., a downtick when news actually was good).²¹

Table V presents the performance of both the PRE signal and the underlying tick signal in all relevant experiments (those with Sequential information arrival and no information auction). For example, in Exp2 there were 32 events to be signalled: news for the second and third subperiods (denoted subperiod b and subperiod c below) for each of the 16 reps. The tick rule was right (up when news was good, down when news was bad) for 21 of these events, and wrong for only 4. There were also 7 instances of no signal, about

²¹One must also take into account the cases where the signal is null. Typically there are 2 to 4 transactions in the calibration interval (i.e., the initial $\varepsilon = 20$ seconds of the second or third subperiod), and buys and sells are about equally likely overall. Hence a priori one might expect the null signal about 25% of the time from the tick rule. The PRE distinguishes seven price intervals with the middle one (indicating a null signal) only \$0.06 wide, so one perhaps should expect null signals from PRE less frequently.

Table V
Signal Performance for PRE and for the Tick Rule

Right and wrong refer to the number of correct and incorrect predictions of private news messages by signals based on the parametric model of partial revelation of private information (PRE) and by signals based on the underlying tick rule. NP means no prediction, %*r* is the percentage of correct predictions and *z* is the binomial test statistic for the null hypothesis that the signal has no predictive power. For the price rule, a “bullseye” indicates that the signal-message correspondence was exact.

Experiment	Tick Rule					PRE					
	Right	Wrong	NP	% <i>r</i>	<i>z</i>	Bulls-eye	Right sign	Wrong	NP	% <i>r</i>	<i>z</i>
Exp2	21	4	7	84.0	3.40	10	14	7	1	77.4	3.05
Exp4	16	8	8	66.7	1.63	10	7	6	9	73.9	2.29
Exp5	25	5	8	83.3	3.65	35	0	2	1	94.6	5.43
Exp6	21	12	7	63.6	1.57	15	4	13	8	59.3	1.06
A3	18	9	5	66.7	1.73	5	10	11	6	71.4	0.78
A6	26	3	3	89.7	4.27	12	9	9	2	70.0	2.19
All	127	41	38	75.6	6.63	87	44	48	27	73.1	6.20

what one might expect. As an index of signal reliability, we compute *z* scores from the formula

$$z = \frac{r - 0.5n}{\sqrt{n(0.5)(0.5)}} \quad (1)$$

where *r* is the number right and *n* is the number either right or wrong, i.e., the number of nonnull signals.²² For example, the *z* value for the tick rule in Exp2 of 3.40 indicates acceptance that the signal is meaningful at approximately the 0.0003 level. A cruder but equally important performance index appears under the heading “% *r*”; it is simply (*r*/*n*)100%, the percentage of nonnull signals that were right. The calculations for PRE are the same, with the convention that both “bullseye” (e.g., signal says 1G when news was 1G) and “right sign” (e.g., signal says 1G when news was 2G) both count as being right, in order to maintain comparability with the less explicit tick signals.

For the most part, the data in Table V confirm the general reliability of both signals and also show that null signals are not especially common. The worst performance occurs in Exp6, a demanding but important 20-rep Het environment. In this experiment reservation prices in subperiod *b* move up or down by \$0.250 and \$0.100 in response to news received by Type 1 and 2 traders, respectively, but only by \$0.025 for news to Type 3 traders, with

²² Under the null hypothesis that the signal is uncorrelated with the news message, *z* will asymptotically have the unit normal distribution. Thus, for large *n*, the probability that *z* exceeds 0, 1, 2, . . . under the null hypothesis is 0.50, 0.16, 0.02, Hence, large positive values of *z* indicate acceptance of the alternative hypothesis that the signal is positively correlated with the appropriate news at high levels of confidence.

more complex cases but similar magnitudes for subperiod *c*. It turns out that 10 of the 12 wrong tick rule signals arose from the relatively unimportant Type 3 traders' news, the other two arising from Type 2 traders' news. All 14 of the signals produced by Type 1 traders' news, the most important sort, were correct. Thus, on closer examination, the signals seemed quite useful even in this experiment.

The bottom line of Table V indicates overall that nonnull signals were transmitted 80–90% of the time. Of these, the tick signals were right over 75% of the time, and the PRE signals (which, if correct, are usually more informative) were right 73% of the time. The corresponding total *z* scores of 6.69 and 6.20 indicate that even at microscopic confidence levels one can conclude that the signals are meaningful.

We now compare the abilities of the three models to forecast asset market outcomes. Our hypothesis tests on asset prices are based on mean squared forecast errors in transaction prices (MSE), defined as follows. Let $p_k(X)$ denote the subperiod *k* price forecast by model *X* ($X = \text{NRE, FRE, or PRE}$), and let $P_1, \dots, P_{N_k}$ denote the actual transaction prices in that subperiod.²³

Then

$$\text{MSE}_{X,k} = \frac{1}{N_k} \sum_{i=1}^{N_k} (P_i - p_k(X))^2.$$

An alternative definition of asset price forecast error considers the trading opportunities implicit in the bid and ask prices rather than the actual transaction prices. It also has the advantage that it is always defined even in subperiods when no transactions were consummated. If changes in bid or ask prices occur at times $t_1 \leq \dots \leq t_{n_k}$ in subperiod $k = (t_0, t_{n_k+1})$ then for A_i denoting the average of the bid and ask prices in the interval (t_i, t_{i+1}) we define the time-weighted mean squared error as

$$\text{TWMSE}_{X,k} = \frac{1}{(t_{n_k+1} - t_0)} \sum_{i=0}^{n_k} (t_{i+1} - t_i) [A_i - p_k(X)]^2.$$

Our asset allocation tests are based on misallocation percentages (MAP), the fraction of the 27 shares held at the end of each subperiod by traders who are not forecast to hold shares in that subperiod under the relevant model. Recall the FRE makes no allocation forecasts for the second and third subperiods when traders are heterogeneously informed.

We apply four test statistics to the reduced data (MSE, TWMSE, and MAP) to assess the relative performance of the competing forecasts. The first is the traditional *F*-ratio with (K-1, K-1) degrees of freedom, where price forecasts *X* and *Y* are compared over the subperiods $k = 1, \dots, K$ in which they differ

²³Recall that the relevant subperiod for PRE begins 20 seconds after the news arrival time, so subperiod *b* is typically the time interval $60 + 20 = 80$ to 120 seconds, and subperiod *c* is 140 to 180 seconds. This short subperiod convention was used for all subperiod *b* and *c* data in Table V. Subperiods *a* and *d* are not used for statistical tests since all forecasts coincide for these subperiods.

by the formula

$$F_{X,Y} = \frac{\sum_{k=1}^K \text{MSE}_{X,k}}{\sum_{k=1}^K \text{MSE}_{Y,k}}$$

Large F values favor Y over X , and their associated probabilities can be interpreted as confidence levels associated with rejecting the null hypothesis that X and Y are equally good forecasts in favor of the alternative hypothesis that Y is better (i.e., has smaller forecast errors). Exactly the same formula and interpretation can be applied to the alternative measure TW MSE. For testing allocation forecasts, one can apply the F -ratio formula with MSE replaced by MAP².

Since the distributional assumptions underlying the F -ratio are not even approximately satisfied by our data, we also employ the nonparametric Wilcoxon rank-sum T - statistic. One rank-orders the combined price or allocation data for the two forecasts X and Y over the subperiods $k = 1$ to K in which they differ. One computes S , the sum of the ranks for the X forecast, and normalizes by the formula $T = (S - \mu)/\sigma$, where $\mu = K(K + 2)/2$ and $\sigma^2 = K^2(2K + 1)/12$ are the mean and variance of S under the null hypothesis that the X and Y data come from the same distribution. Under this null hypothesis T has the unit normal distribution asymptotically for large K . Again large values of T favor forecast Y over forecast X , and the test applies to the MSE, TW MSE (or equivalently $\text{RMSE} = \sqrt{\text{MSE}}$, $\text{RTW MSE} = \sqrt{\text{TW MSE}}$), and MAP data.

Potentially more powerful tests exploit the natural pairing of X and Y forecasts for the same subperiod. The paired differences $d_k = \text{RMSE}_{X,k} - \text{RMSE}_{Y,k}$ have mean 0 and unknown variance under the null hypothesis, so using the sample variance $s^2 = (K - 1)^{-1} \sum (d_k - \bar{d})^2$, where $\bar{d}$ is the sample mean, one obtains the simple t -test statistic $t = \bar{d}/s$ with the usual interpretation. Again, the Normality assumption may not be valid, so we consider the best single indicator of relative performance to be a signs test using the statistic z as defined in equation (1), with $n = K$ and $r =$ the number of k such that $d_k > 0$. Both the t and z statistics apply equally to the paired differences in RTW MSE and MAP data.

Table VI presents our statistical results. Entries all show test results for X as indicated in the column heading and $Y = \text{PRE}$, so positive entries favor PRE over the given alternative forecast. The evidence for PRE over NRE is overwhelming for asset price forecasts: in every one of the six experiments, using either RMSE or the time-weighted version, every one of the test statistics favors PRE (F values greater than 1, other test statistics positive). The pooled results, for all Hom and all Het reps, and for both sites, are highly significant. The grand pooled z value, based on 104 paired comparisons, is over 4.5, corresponding to a confidence level of less than 0.0001. The other test statistics for asset prices are even more decisive. The evidence on the

Table VI

Hypothesis Tests Comparing PRE to Alternative Models of No (NRE) and Full (FRE) Information Revelation

Variable K is the number of subperiods where PRE forecasts differ from NRE or FRE. F , T , t , and z are significance tests of PRE versus the alternative; the F -test, the Wilcoxon T -test, Student's t , and the binomial test, respectively. MSE and TW MSE are the mean-squared error statistics and time-weighted mean-squared error statistics for price predictions, respectively, and MAP is the percentage of shares misallocated relative to the forecast. Observations in which NRE and PRE give the same value of MAP are omitted from the z -test and the number of observations remaining appears in parentheses, e.g., 63 observations overall. Positive entries for the T , t - and z -statistics favor PRE.

Experiment	Test	MSE		TWMSE		MAP
		NRE	FRE	NRE	FRE	NRE
Exp2	K	14	15	14	16	17
	F	1.544	0.487	1.651	0.413	0.615
	T	1.539	-1.887	1.539	-2.696	1.519
	t	2.49	-5.65	2.49	-5.30	-1.74
	z	1.603	-3.873	1.603	-4.000	-0.534(14)
Exp4	K	19	14	19	15	12
	F	2.506	0.735	2.435	0.276	0.557
	T	1.023	-0.103	0.557	-2.903	1.997
	t	3.26	-0.20	1.86	-4.02	-2.12
	z	0.688	-1.069	0.688	-2.840	-0.577
Exp5	K	20	0	20	0	1
	F	9.130	-	11.966	-	2.914
	T	4.181	-	5.288	-	-
	t	6.39	-	8.49	-	-
	z	2.683	-	4.024	-	1.000
Exp6	K	19	18	19	18	12
	F	2.448	0.603	2.477	0.502	0.887
	T	1.956	-0.174	1.518	-1.218	0.000
	t	2.43	-0.90	1.91	-1.68	-0.34
	z	1.605	-0.942	1.147	-0.942	1.000(9)
M3	K	14	18	14	19	16
	F	2.076	0.664	2.334	0.649	0.592
	T	1.608	-0.773	1.952	-1.051	2.212
	t	3.14	-3.07	3.43	-2.72	-2.65
	z	1.603	-2.357	2.138	-2.523	-1.603(14)
M6	K	18	17	18	17	15
	F	2.464	0.619	2.635	0.587	0.790
	T	2.119	-1.567	2.420	-1.446	0.644
	t	5.35	-2.29	5.41	-2.31	-0.72
	z	2.828	-2.667	2.828	-2.667	-0.832(13)
All Hom	K	53	32	53	34	32
	F	3.294	0.528	3.965	0.456	0.712
	T	4.565	-1.600	5.573	-2.594	1.586
	t	9.24	-3.95	9.35	-4.95	-1.70
	z	4.258	-3.53	5.357	-4.116	0.00(30)

Table VI—Continued

Experiment	Test	MSE		TWMSE		MAP
		NRE	FRE	NRE	FRE	NRE
All Het	<i>K</i>	51	50	51	51	41
	<i>F</i>	1.627	0.629	1.705	0.549	0.662
	<i>T</i>	2.218	−0.927	2.155	−2.292	2.187
	<i>t</i>	4.11	−2.87	3.56	−4.71	−2.76
	<i>z</i>	2.100	−3.394	2.100	−4.060	−1.566(33)
All Site U	<i>K</i>	72	47	72	49	42
	<i>F</i>	2.288	0.532	2.539	0.404	0.685
	<i>T</i>	4.261	−1.117	4.676	−3.367	1.988
	<i>t</i>	7.00	−3.28	6.41	−5.98	−2.34
	<i>z</i>	3.299	−3.354	3.771	−4.428	0.00(36)
All Site M	<i>K</i>	32	35	32	36	31
	<i>F</i>	2.363	0.642	2.554	0.617	0.684
	<i>T</i>	2.423	−1.576	2.880	−1.661	2.002
	<i>t</i>	6.00	−3.62	6.25	−3.48	−2.06
	<i>z</i>	3.182	−3.549	3.5534	−3.666	−1.732(27)
All	<i>K</i>	104	82	104	85	73
	<i>F</i>	2.315	0.586	2.545	0.509	0.384
	<i>T</i>	4.933	−1.709	5.574	−3.399	2.722
	<i>t</i>	8.99	−4.74	8.53	−6.80	−3.13
	<i>z</i>	4.510	−4.859	5.099	−5.748	−1.133(63)

allocation forecasts of PRE versus NRE is mixed: few of the test statistics for individual experiments are significant at even the 0.05 level, and of those that are, the signs are positive about as often as negative. The grand pooled results are also mixed, with *F* and matched *t*-tests indicating significant evidence in favor of NRE and the Wilcoxon test indicating the opposite. We believe that the matched *z*-test is the most reliable indicator, and it is insignificant at −1.13. We conclude that NRE and PRE do about equally well in predicting asset allocation, with aberrant PRE forecasts (such as the previously discussed rep 10 of Exp6) roughly cancelling out the more refined forecasts provided in other cases.

The asset price evidence clearly favors FRE over PRE. For example, all four of our test statistics for the grand pooled MSE data favor FRE, and three of them are highly significant. The fourth, a marginal Wilcoxon *T* of -1.709 again suggests the role played by occasional gross errors in the PRE signals.

C. Summary of Experimental Results

The evidence in Table V indicates that good PRE signals are available in our experimental markets. Additional evidence suggests that traders are able to use these partially revealing signals to enhance market efficiency.

The FRE, a strong-form hypothesis, continues to significantly outperform the alternatives with regard to asset price forecasts. It fails badly, however, when we focus on other features of the market, namely the price and allocation of purchased information and the allocation of traded assets. The mechanical PRE model does much better in these dimensions. The PRE model outperforms both FRE and NRE in explaining the price of purchased information (Figure 3.) The FRE model cannot make any asset allocation forecasts during the relevant (second and third) trading subperiods, but PRE does as well as NRE in this dimension. Finally, PRE makes significantly better asset price forecasts than does the more naive NRE model.

V. Discussion

The results of our 14 experiments generally support partial revelation of information over the nonrevelation and full revelation alternatives we tested. A specific parametric model of partial revelation (defined in Section III. A and subsequently referred to as PRE) clearly outperformed the alternative models in predicting the price and allocation of purchased information. The PRE provided correct signal-message correspondences most of the time; it dominated NRE in forecasting asset prices; it differed insignificantly from NRE in forecasting asset allocations; and it dominated FRE in forecasting asset allocation. Its only important failure was as second place performance to FRE in forecasting asset prices.

It is worth emphasizing that our version of PRE was specified a priori and not modified after analyzing the data. As explained in Section IV. A, a wide bid/ask spread and/or approximately equal numbers of boughts and solds can produce unrealistic signals in the current version of PRE. We believe that a patient specification search would disclose some version consistent with the general model of PRE presented in Section II A that would do as well as or better than FRE in explaining asset prices while maintaining its dominance in other dimensions.²⁴ Such a modified version of PRE could not legitimately be tested on our data, so we leave the task to some future time when new experiments (employing Seq:Het or some similar sort of information arrival process) can be performed.²⁵

The intriguing connection of our PRE framework to the arcane art known as “technical analysis” has not escaped our attention. According to sympathetic textbook accounts (e.g. Gitman and Joehnk (1988, Chap. 8), the logical basis for technical analysis is that shifts in asset supply or demand are not instantaneous but rather take some time to complete themselves. (See also

²⁴One possibility is the “excess bids” indicator (new unaccepted bids less new unaccepted asks in a calibration interval) of Smith, Suchanek, and Williams (1988). This indicator could be used alone or in conjunction with some adjusted version of the tick or PRE rule.

²⁵Leamer (1983) eloquently reminds us that classical statistics are invalid for testing hypotheses obtained by a specification search which uses the same data for searching as for testing. An important advantage of experimental techniques is that (budget permitting) new data can be generated to permit valid tests of hypotheses obtained by a specification search.

Copeland (1976). Technical indicators are intended to detect such shifts while they are still underway. For instance, price increases or decreases supposedly presage more of the same if accompanied by heavy trading volume, but such "momentum" can be reversed when a "support" or "resistance" level is encountered. The empirical evidence is at best weak that such analysis is profitable at the margin in contemporary asset markets, but the prevalence of its practitioners suggests that there may be some inframarginal gains. In our laboratory we found that the tick rule, defined as accepted asks less accepted bids (a momentum indicator of sorts), and the parametric PRE rule, defined in terms of observed price relative to benchmark prices (implicitly involving support or resistance levels), do convey information.

A final speculation concerns minibubbles. If traders rationally attempt to extract information from imperfect market signals, then it seems likely that they will sometimes be misled. For example, a price rise due to noise may sometimes be misinterpreted as the arrival of positive private information, and traders' response may appear to confirm the interpretation, provoking further price rises. Such "minibubbles" or informational mirages were first observed in the laboratory by Camerer and Weigelt (1986) and might arise in contemporary asset markets as well. Hence those who pursue more refined versions of partial information revelations should be alert to possible connections to excess volatility and technical indicators in asset markets.

REFERENCES

- Allen, B., 1981, Generic existence of equilibrium for economies with uncertainty when prices convey information, *Econometrica* 49, 1173-1199.
- Camerer, C. and K. Weigelt, 1986, Testing for information 'mirages' in experimental asset markets, DDS Working paper, University of Pennsylvania.
- Copeland, T., 1976, A model of asset trading under the assumption of sequential information arrival, *The Journal of Finance* 31, 1149-1168.
- and D. Friedman, 1987, The effect of sequential information arrival on asset prices: An experimental study, *The Journal of Finance* 42, 763-797.
- , 1988, The market value of information: Some experimental results, Working paper #176, UC Santa Cruz Applied Economics.
- and D. Galai, 1983, Information effects on the bid-ask spread, *Journal of Finance* 38, 1457-1469.
- Diamond, D. W. and R. E. Verrecchia, 1981, Information aggregation in noisy rational expectations economy, *Journal of Financial Economics* 9, 221-238.
- Easley, D. and J. O. Ledyard, 1986, Theories of price formation and exchange in oral auctions, Social Sciences working paper #661, California Institute of Technology.
- Fama, E., 1970, Efficient capital markets: A review of empirical work, *The Journal of Finance* 25, 383-417.
- Forsythe, R., T. Palfrey and C. Plott, 1982 Asset valuation in an experimental market, *Econometrica* 50, 537-567.
- Friedman, D., 1984, On the efficiency of double auction markets, *American Economic Review* 74, 1984 60-72.
- , 1987, A simple testable model of double auction markets, Manuscript *Journal of Economic Behavior and Organization*, Forthcoming.
- , G. Harrison, and J. Salmon, 1984, The informational efficiency of experimental markets, *Journal of Political Economy* 92, 349-408.

- Gitman, L. J. and M. D. Joehnk, 1988, *Fundamentals of Investing*, 3rd ed. (Harper & Row, New York).
- Glosten, L. and P. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *The Journal of Financial Economics* 4, 71-100.
- Grossman, S. and J. Stiglitz, 1976, Information and competitive price systems, *American Economic Review* 66, 246-253.
- Hellwig, M. F., 1980, On the aggregation of information in competitive markets, *Journal of Economic Theory* 22, 477-498.
- Jordan, J. S., 1983, On the efficient market hypothesis, *Econometrica* 51, 1325-343.
- Kyle, A. S., 1986, Informed speculation with imperfect competition, Memorandum #66, Princeton University Financial Research Center.
- Latham, M., 1985, Defining capital market efficiency, Finance working paper #150, IBER, University of California at Berkeley.
- Leamer, E. E., 1983, Let's take the con out of econometrics, *American Economic Review* 73, 31-43.
- Plott, C. and S. Sunder, 1982, Efficiency of experimental security markets with insider information: An application of rational expectations models, *Journal of Political Economy* 90, 663-698.
- Rubinstein, M., 1975, Securities market efficiency in an Arrow-Debreu economy, *American Economic Review* 65, 812-824.
- Smith, V., 1976, Experimental economics: Induced value theory, *American Economic Review* 66, 274-79.
- , 1989, Theory, experiment and economics, *Journal of Economic Perspectives* 3, 151-169.
- G. Suchanek, and A. Williams, 1988, Bubbles, crashes and endogenous expectations in experimental asset markets, *Econometrica* 56, 1119-1152.
- Sunder, S., 1988, Market for information: Experimental evidence, Working paper, University of Minnesota.
- Thaler, R., 1987, Anomalies: The January effect, *The Journal of Economic Perspectives* 1, 197-201.
- Verrecchia, R. E., 1982, Information acquisition in a noisy rational expectations economy, *Econometrica* 50, 1415-1430.
- von Borries, A. and D. Friedman, 1988, Monopolist insiders in computerized asset markets: A note on some experimental results, Working paper No. 178, University of California at Santa Cruz.
- Wilson, R., 1987, Equilibrium in bid-ask markets, G. Feiwel, ed.: in *Arrow and the Ascent of Economic Theory* (McMillan Press, London).