Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models

Charles R. Plott

California Institute of Technology

Shyam Sunder

University of Minnesota

The study reports on the ability of competing models of market information integration and dissemination to explain the behavior of simple laboratory markets for a one-period security. Returns to the security depended upon a randomly drawn state of nature. Some agents (insiders), whose identity was unknown to other agents, knew the state before the markets opened. With replication of market conditions the predictions of a fully revealing rational-expectations model are relatively accurate. Prices adjusted immediately to near rational-expectations prices; profits of insiders were virtually indistinguishable from noninsiders; and efficiency levels converged to near 100 percent.

I. Introduction

This study reports on the behavior of five markets created in a laboratory environment to explore some theoretical implications of insider information. Recently, principles of rational expectations have

been extended to theories of the allocative properties of securities markets. The resulting models deviate substantially from traditional demand and supply models and have important consequences for social policies regarding market organization and control. Because of the possibly subtle and complex role played by information and expectations in market activity, the models are hard to test, and many aspects of them remain controversial.

We have attempted to create laboratory markets which are sufficiently simple to allow a clearer glimpse of the inherently complex phenomena of information utilization and dissemination within a market. We do not attempt to evaluate the models as applied to the U.S. securities markets or any other complicated, naturally occurring market. Our goal is only to determine the appropriateness of various models for predicting the behavior of very simple markets with the hope that the understanding gained will be useful in ascertaining features of models which will have successful applications to more complex markets.[1]

Security markets have two important dimensions. The first is time in the sense that securities yield a stream of cash over time. An inquiry about the nature of this first dimension, by using laboratory techniques, was initiated by Forsythe, Palfrey, and Plott (forthcoming). The second dimension, uncertainty, is explored in this study by abstracting the time dimension and focusing on a security with only a one-period life. Several competing modes of thought can be applied to predict the price of this security. Because it can be traded for capital gains, theories of pure speculation (notably Keynes 1936; also Harrison and Kreps 1978) maintain that the prospect of speculative returns can dominate the underlying returns or intrinsic value of the security. The price might wander over a wide range of values or possibly "explode." At the other extreme are strong versions of rational-expectations theories which assert that market prices adjust instantaneously to equilibrium price levels which reflect at least all available information about the state of nature. In the case of the "crystal ball" hypothesis the equilibrium may reflect even more information than the sum of what is available to individual participants.

Differences among some models turn on the role played by price itself in conveying information about the state of nature. On one hand are fully revealing rational-expectations (RE) equilibria (see Lucas 1972 and Green 1973) in which prices in equilibrium reveal the

---

[1] We cannot claim that we have tested any of the models found in the literature. All models are accompanied by technical assumptions and qualifications which are not present in the simple markets we created. These markets may be used to test these models only if such assumptions are placed in an "as if" category and are not taken literally.

state because they must be consistent with the expectations individuals have about the state of nature when they face those prices. On the other hand are the prior-information (PI) equilibria which hold that individuals do not condition expectations upon price. Instead, expectations are exogenous to the price formation process with individuals utilizing whatever prior information they might have at their disposal. With expectations formed, prices are determined by a straightforward application of the principles of demand and supply as in a Walrasian system.

In the sections below we report on five markets. The results of the first experiment, in which some insiders were given less than certain information about the state of nature before trading began, provided little support for the RE model. Since our own expectations leaned in favor of prior-information equilibrium, we designed the second market to give the RE model the best chance against the PI model; insiders were given perfect information about the realized state before trading began, thus allowing traders more opportunity to gain experience over repeated occurrences of the state. The RE model performed well, with prices and allocations converging close to the fully revealing equilibrium. The next two markets involved replications without the experience with the certainty case, and the final market, 5, involved an increase in the number of states of nature from two to three. The RE model continued to predict well.

The paper is developed as follows. In Section II the design of markets is discussed. In Section III our hypotheses are stated, and in Section IV they are reviewed in light of the results. In Section V other implications of the results are discussed in light of current mathematical models and the statistical models applied to the U.S. securities industry. The last section is a summary of conclusions. Details of experimental procedures which are at the heart of the market are outlined in the Appendix.

II. Design of Markets

Each market experiment involved several periods. In each period securities which had one-period lives were traded. Each security paid a single dividend to its holder at the end of the period. The dividend from holding a security differed across individuals (which could be analogous to different tax brackets or risk preferences) and depended upon the (randomly drawn) state of nature.[2] Differences in dividends and possibly expectations led to the existence of gains from exchange

---

[2] Instead of being viewed as a security, the traded object could also be viewed as a resource being traded among the owners of alternative production technologies.

and market activity. The markets were organized as oral double auctions similar to those on the floor of the New York Stock Exchange.[3]

Subjects were students at the Graduate School of Business at the University of Chicago. In markets 1, 2, and 4 all subjects were inexperienced in that none had previously participated in a laboratory market. Subjects in markets 3 and 5 had participated in one or more of the earlier markets.

By following the accepted method for inducing preferences (Smith 1976; Plott 1979) each individual subject, i, was assigned a dollar redemption function of the form $R_i^t = \gamma_i[a_i + d_i(\theta_t)x_i^t + \Sigma_s P_s^{it} - \Sigma_p P_p^{it} + C_i^t]$, $a_i < 0$, $d_i(\theta) > 0$, $\gamma_i > 0$, $x_i^t \geq 0$, where

R_i^t = dollar earnings of individual i in period t,

x_i^t = units held by i at the end of period t (all short sales were prohibited so $x_i^t \geq 0$) is the sum of initial endowment of securities plus purchases less sales in period t,

$d_i(\theta)$ = dividend rate in francs for individual i and expressed as a function of the state of nature θ,

$\Sigma_s P_s^{it}$ = revenue from sales of securities during period t,

$\Sigma_p P_p^{it}$ = cost of securities purchased during period t,

$\theta \in \Omega$ = possible states of nature,

C_i^t = initial endowment of cash in francs,

a_i = fixed cost in francs, and

γ_i = conversion rate of francs into dollars.

As long as an individual has a positive utility for money, he or she would like R_i^t as large as possible. This motivation systematically induces values on the securities by virtue of derived demand theory. Such values can then be used as parameters on market models.

As is implicit in the formula above, at the beginning of each period each individual was given an initial endowment of working capital (C_i^t) which was sufficiently large never to serve as a binding constraint on purchases of securities. Each individual was also given an initial endowment of securities ($\bar{x}_i^t$) at the beginning of each period. Since short sales were never permitted except as a reduction of an individual's initial endowment, there was a fixed supply of securities, $\Sigma_i \bar{x}_i^t$.

The initial endowment in the form of working capital and securities was frequently of substantial value. A fixed cost each period, $a_i < 0$, was imposed to reduce the cost of the experiment.

All trading was in francs. Use of francs allows flexibility in the

---

[3] Any buyer (seller) is free at any time to make an oral bid (offer) to buy (sell) one unit of the security at a designated price. Such bids and offers are publicly announced and recorded. Only one (the last) bid and offer are outstanding at any time. Sellers (buyers) are free to accept any public bid (offer) they wish.

choice of parameters while avoiding the technical problems inherent in using large dollar amounts. Francs have been successfully employed in other experiments (Friedman 1967; Forsythe et al. forthcoming).

All investors were aware of the mechanism used to determine the state of nature each period and had experience with its operation (see Instruction Set 1 in the Appendix for a description and table 1 for prior probabilities used).[4] At the beginning of trading each period, prior information of noninsiders about the state consisted solely of this knowledge. Certain investors became insiders by receiving a clue card bearing information about the realized state. Distribution of blank clue cards to noninsiders prevented the identity of the insiders from being revealed.

The case where no individual received information about the realized state is called the "no-information" case. As shown in table 1, the first few periods of every market (the first four in 1, the first four in 2, the first two in 3, the first four and the last in 4, and the first three in 5) were conducted under the no-information condition. This condition served as both a training period and a period of calibration for experimental purposes, as will be discussed below.

In periods 5 and 6 of market 2, all investors were privately given clues in order to let agents learn about the price/state correspondence prior to their conducting operations in less informed environments. The results suggested that this type of training was unnecessary for the successful application of RE models. This was confirmed in markets 3, 4, and 5.

In most periods of all markets only one-half (two out of four) of the agents from each dividend (preference) type received information about the realized state of nature.[5] The design is displayed in table 1. This allowed simultaneous existence of insiders and noninsiders and the study of a market with asymmetric distribution of information.

The results of market 1 suggested that the markets may adjust to the general pattern of information in the market. But some replication may be necessary for each information pattern before equilib-

---

[4] Even though individuals were trained, we still had no way of knowing their subjective probabilities.

[5] In market 1 only one out of three investors of each type was an insider, and the information received by insiders was less than certain. The "clue" given to the insiders was a sample of 10 draws with replacement. The sample was taken from urn X containing balls marked "0" and "I" [$\mathrm{pr}(0 \mid X) = 4/5$, $\mathrm{pr}(I \mid X) = 1/5$] if the state randomly chosen was X and the sample was drawn from urn Y [$\mathrm{pr}(0 \mid Y) = 3/5$, $\mathrm{pr}(I \mid Y) = 2/5$] if the randomly chosen state was Y. The samples given to the three insiders are shown in table 1 together with the Bayesian posterior probabilities of the state, given the sample. If the insiders were perfect Bayesians, then their subjective probabilities of the state, given the information they received, would be those in the table.

TABLE 1
INFORMATION DESIGN

Market No.	Prior Prob. of X	Knowledge about Others' Dividends	Period	N Informed Agents	Precision of Information	Correlation of Info. across Investors	How Many	Who	What Informed Know	Info. on Card	Posterior Prob (X)	Actual State
							Common Knowledge about Informed Agents			**Information Held by Informed Agents**		
1	1/3	None	1	None	...	...	Yes (0)	...	Yes	No card	⅓	Y
			2	None	...	...	Yes (0)	...	Yes	No card	⅓	Y
			3	None	...	...	Yes (0)	...	Yes	No card	⅓	X
			4	None	...	...	Yes (0)	...	Yes	No card	⅓	Y
			5	3-one of each type	Imperfect (a sequence of 10 clues)	Perfect	No	No	Yes	0100101010	.15	Y
			6	3-one of each type		Perfect	No	No	Yes	0000000011	.555	X
			7	3-one of each type		Perfect	No	No	Yes	0100110100	.15	Y
			8	3-one of each type		Perfect	No	No	Yes	0000010000	.77	Y
			9	All		Perfect	No	No	Yes	1110000011	.062	Y
			10	All		Perfect	No	No	Yes	1010000011	.15	X
			11	All		Perfect	Yes	Yes	Identical	1111111001	.003	Y
2	1/3	None	1	None	...	...	Yes (0)	...	Yes	No card	⅓	X
			2	None	...	...	Yes (0)	...	Yes	No card	⅓	X
			3	None	...	...	Yes (0)	...	Yes	No card	⅓	Y
			4	None	...	...	Yes (0)	...	Yes	Y	0	Y
			5	All	...	...	No	No	Yes	Y	0	Y
			6	All	...	...	No	No	Yes	Y	0	Y
			7	6-two of each type	Perfect	Perfect	No	No	Yes	X	1	X
			8	6-two of each type	Perfect	Perfect	No	No	Yes	Y	0	Y
			9	6-two of each type	Perfect	Perfect	No	No	Yes	X	1	X
			10	6-two of each type	Perfect	Perfect	No	No	Yes	Y	0	X
			11	6-two of each type	Perfect	Perfect	No	No	Yes	Y	0	Y

668

3 .4 Different investors may have different dividends

								Posterior Prob.	
1	None	⋯	⋯	No	No	Yes	None	.4	Y
2	None	⋯	⋯	No	No	Yes	None	.4	Y
3	6-two of each type	Perfect	Perfect	No	No	Yes	Y	0	Y
4	6-two of each type	Perfect	Perfect	No	No	Yes	X	1	X
5	6-two of each type	Perfect	Perfect	No	No	Yes	Y	0	Y
6	6-two of each type	Perfect	Perfect	No	No	Yes	X	1	X
7	6-two of each type	Perfect	Perfect	No	No	Yes	Y	0	Y
8	6-two of each type	Perfect	Perfect	No	No	Yes	X	1	Y
9	6-two of each type	Perfect	Perfect	No	No	Yes	Y	0	Y
10	All	⋯	⋯	No	No	Yes	Y	0	Y
11	All	⋯	⋯	No	No	Yes	Y	0	Y
12	All	⋯	⋯	No	No	Yes	X	1	Y

4 .4 Different investors may have different dividends

								Posterior Prob.	
1	None	⋯	⋯	No	No	Yes	None	.4	X
2	None	⋯	⋯	No	No	Yes	None	.4	Y
3	None	⋯	⋯	No	No	Yes	None	.4	X
4	None	⋯	⋯	No	No	Yes	None	.4	Y
5	6-two of each type	Perfect	Perfect	No	No	Yes	Y	0	Y
6	6-two of each type	Perfect	Perfect	No	No	Yes	X	1	Y
7	6-two of each type	Perfect	Perfect	No	No	Yes	Y	0	Y
8	6-two of each type	Perfect	Perfect	No	No	Yes	X	1	Y
9	6-two of each type	Perfect	Perfect	No	No	Yes	Y	0	Y
10	6-two of each type	Perfect	Perfect	No	No	Yes	Y	1	Y
11	6-two of each type	Perfect	Perfect	No	No	Yes	X	0	X
12	6-two of each type	Perfect	Perfect	No	No	Yes	Y	1	Y
13	6-two of each type	Perfect	Perfect	No	No	Yes	X	0	X
14	None	⋯	⋯	No	No	Yes	None	.04	Y

5 $P(X) = .35$, $P(Y) = .25$, $P(Z) = .4$ Different investors may have different dividends

								Posterior Prob.			
								X	Y	Z	
1	None	⋯	⋯	Yes (0)	⋯	Yes	No card	.35	.25	.40	N
2	None	⋯	⋯	Yes (0)	⋯	Yes	No card	.35	.25	.40	X
3	None	⋯	⋯	Yes (0)	⋯	Yes	No card	.35	.25	.40	N
4	6-two of each type	Perfect	Perfect	No	No	Yes	X	1	0	0	X
5	6-two of each type	Perfect	Perfect	No	No	Yes	X	1	0	0	Y
6	6-two of each type	Perfect	Perfect	No	No	Yes	Y	0	1	0	N
7	6-two of each type	Perfect	Perfect	No	No	Yes	Z	0	0	1	N
8	6-two of each type	Perfect	Perfect	No	No	Yes	Z	0	0	1	Y
9	6-two of each type	Perfect	Perfect	No	No	Yes	Y	0	1	0	Y
10	6-two of each type	Perfect	Perfect	No	No	Yes	Y	1	0	0	X
11	6-two of each type	Perfect	Perfect	No	No	Yes	X	0	1	0	Y
12	6-two of each type	Perfect	Perfect	No	No	Yes	Y	1	0	0	N
13	6-two of each type	Perfect	Perfect	No	No	Yes	Z	0	0	1	N

669

rium is finally established for that information pattern. A reduction in
the number of different information patterns would thus result in a
reduction of periods necessary to see equilibrating behavior should it
exist. In all markets after market 1 (periods 5 through 11 in market 2,
3 through 12 in market 3, and so forth) the information contained on
the clue card received by the insiders indicated the realized state of
nature with certainty.

Information that everyone knows publicly (as opposed to privately)
is called "common knowledge" (Aumann 1976). Since the instructions
were read publicly, each agent knew what other agents knew from
that source. In particular, every agent knew that no one knew the
number or the identity of insiders. The exceptions are periods 1
through 4 of market 1, 1 through 4 of market 2, 1 through 3 of
market 5, in all of which the fact that no one had any information was
announced, and also in period 11 of market 1 in which the clue was
publicly announced. Additional common knowledge of possible im-
portance was that the clues of all insiders were identical. From the
nature of the instructions, agents could deduce in all but market 1
that the dividend values for every agent remained constant from
period to period. Agents did not know the number of agent types,
although in markets 3, 4, and 5 they were told that the dividends of
others may differ from their own (in 1 and 2 nothing was said), and
they did not know that the insiders were the same agents throughout
the relevant periods.

The dividend parameters for all markets are summarized in table 2.
Agents in each experiment were partitioned into three types (desig-
nated as I, II, and III) according to dividend returns. There were
four agents of each type except in market 1, where there were three.
This provided markets with 12 agents except market 1, in which there
were nine. Each period each agent had an initial endowment of two
certificates, giving a market supply of 24 units (18 for market 1). In
addition, each period each agent was given 10,000 francs in working
capital which was returned to the experimenter at the end of the
period by the imposition of a fixed cost of 10,000 francs. The conver-
sion rates of dollars per franc are in table 2.

The dividends, paid at the end of the period, differed by agent type
and according to the state of nature. Reading table 2 we can see that in
market 1 there are two possible states, X and Y. An agent of type I
receives 150 francs for every certificate held if the state is X and 350
francs per certificate if the state is Y. The dividends for other types
and other markets are determined similarly. For example, in market 5
there are three possible states (X, Y, and Z), and an agent of type I
receives dividends per certificate of 120, 170, or 320 francs, depend-
ing upon whether the state is X, Y, or Z.

TABLE 2
DIVIDEND PARAMETERS

MARKET AND INVESTOR TYPE	INVESTORS (N)	INITIAL ENDOWMENT Certificate	INITIAL ENDOWMENT Francs	FIXED COST*	DOLLAR PER FRANC	DIVIDENDS* X	DIVIDENDS* Y	DIVIDENDS* Z	PRIOR PROBABILITY X	PRIOR PROBABILITY Y	PRIOR PROBABILITY Z	EXPECTED DIVIDEND
1:												
I	3	2	10,000	10,000		150	350	...				283.3
II	3	2	10,000	10,000	.002	250	300	...	1/3	2/3	...	283.3
III	3	2	10,000	10,000		300	100	...				166.7
2:												
I	4	2	10,000	10,000		100	350	...				266.7
II	4	2	10,000	10,000	.002	200	300	...	1/3	2/3	...	266.7
III	4	2	10,000	10,000		240	175	...				196.6
3:												
I	4	2	10,000	10,000		400	100	...				220
II	4	2	10,000	10,000	.003	300	150	...	.4	.6	...	210
III	4	2	10,000	10,000		125	175	...				155
4:												
I	4	2	10,000	10,000		375	100	...				210
II	4	2	10,000	10,000	.003	275	150	...	.4	.6	...	200
III	4	2	10,000	10,000		100	175	...				145
5:												
I	4	2	10,000	10,000		120	170	320				212.5
II	4	2	10,000	10,000	.003	155	245	135	.35	.25	.4	169.5
III	4	2	10,000	10,000		180	100	160				152

* In francs.

671

The prior probabilities of states and the dividends were chosen so that the prices and allocations predicted by competing models would reasonably be separated. The final column gives the expected value of a certificate for a given type of agent based on the prior probabilities.

III. Models

We focus on two competing ideas. The first is the PI equilibrium (Walrasian) hypothesis and the second is the RE equilibrium hypothesis. For completeness, however, a third set of three models will be mentioned.

The PI equilibrium hypothesis has been the traditional vehicle for incorporating uncertainty into market and other economic models. Each individual, i, is provided with endowment of securities, $\bar{x}_i$, a priori information, $p^i(\theta)$, about the state of nature θ. Demand functions, $x_i = D^i(p, \bar{x}_i)$, are then generated from the theory that individuals attempt to

$$\max_{x_i} \sum_\theta u^i(x_i, \theta) p^i(\theta) \text{ subject to } p(x_i - \bar{x}_i) = 0.$$

Prices are determined by the law of supply and demand

$$\sum_i D^i(p, \bar{x}_i) - \sum_i \bar{x}_i = 0.$$

For the market settings and parameters described for the experiments above, the application of this model is straightforward. In the absence of risk aversion, at any price below the expected dividend value, agents will demand as many units as their working capital will buy. At prices above the expected dividend value, agents will sell the two units initially endowed and would sell more if short sales were permitted. In the absence of risk aversion, market demand will thus be "horizontal" until the capital constraint becomes binding as shown in figure 1, and the demand price will equal the expected dividend value of the agent type with the maximum expected dividend. Supply is limited by the initial endowments.

The predictions of this model for various states and information conditions imposed across all experiments, under the assumption of risk neutrality, are summarized in table 3. These figures are taken from table 2 by determining the type of agent who has the maximum expected value. The model predicts both price and final allocations.
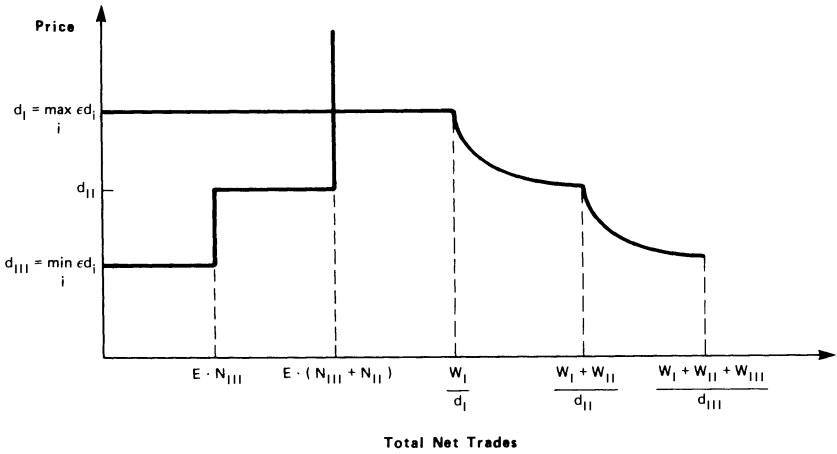
FIG. 1.—Supply and demand functions: W_i = initial working capital of investors of type i; N_i = number of agents in the market of dividend type i; E = initial endowment of securities per agent; d_i = dividend of agents of dividend type i; ϵ = mathematical expectation with respect to the prior probability distribution of the states of nature.

For example, in market 2 the price will be 266 in state X because uninformed agents of both types I and II have an expected value of 266. All other agents have lower expected values and should sell to these four agents who have a horizontal demand at that price.

The second model is of RE equilibrium in which individuals condition their expectations about the underlying state of nature upon the equilibrium values of the endogenous variables of the system. The equilibria themselves must be consistent with this type of adaptive behavior. A variety of models exist, depending upon the information conveyed to agents by the equilibrium values, the inference process of agents, etc. The setting we chose for the experimental markets is one in which agents are assumed to be risk neutral and prices are fully revealing.

The RE equilibrium prices given states $X, Y,$ and Z are not equal to each other. Therefore, individuals who have no external information about the state of nature but who observe all transactions should, according to the RE equilibrium models, infer the states of nature and adjust their behavior accordingly. The resulting predictions are also given in table 3.

The critical differences between the price predictions of the two models occur when the expected value of the security based upon the prior information and dividends of the uninformed individuals lies above the RE equilibrium price. The RE expectations equilibrium models predict that such individuals will revise their expectations

TABLE 3

Price and Allocation Predictions: Prior Information (PI) versus Rational Expectations (RE)* with Risk Neutrality

Market and Model	Price				Information of Insiders	Type of Agent Holding Certificates in Equilibrium		
	None	X	Y	Z	None	X	Y	Z
2:								
PI	266	266†	350	…	I and II	I and II uninformed†	I insiders†	…
RE	266	240†	350	…	I and II	III†	I†	…
3:								
PI	220	400	220†	…	I	I insiders†	I uninformed†	…
RE	220	400	175†	…	I	I†	III†	…
4:								
PI	210	375	210†	…	I	I insiders†	I uninformed†	…
RE	210	375	175†	…	I	I†	III†	…
5:								
PI	212	212†	245	320	I	I uninformed†	II insiders†	I insiders†
RE	212	180†	245	320	I	III†	II†	I†

* In market 1 information given to insiders was probabilistic. Predictions are not given here in order to save space.
† The two predictions differ here.

674

downward, based on market-generated information, and allow the price to fall.

The pattern of holdings predicted by the two models almost always differs. The RE model predicts that the uninformed individuals will behave the same as the insiders since market price here is fully revealing. Thus, the PI equilibrium predicts that only the insiders of a given type can hold, while the RE model predicts that all members of this type can hold.

For convenience of exposition the competing predictions can be summarized by the following two hypotheses:

> HYPOTHESIS 1: Prices converge to the prior-information equilibrium values given in table 3.
> ALTERNATE HYPOTHESIS 1: Prices converge to the rational-expectations equilibrium values given in table 3.
> HYPOTHESIS 2: Security holdings converge to the prior-information equilibrium values given in table 3.
> ALTERNATE HYPOTHESIS 2: Security holdings converge to the rational-expectations equilibrium values given in table 3.

Two additional measures, profit and efficiency, are relevant for separating the PI and RE models. If the market data are completely revealing of the state, then insider information is of no value. Thus, under the RE hypothesis insiders should make no more profit than uninformed individuals. In some respects this may be a better indicator of market-generated information than the final holdings. Risk-averse, uninformed individuals may reap the benefits of almost full information by selling to informed individuals at near equilibrium prices and avoid the risk of holding in the absence of certainty. The holdings data could then be used to *reject* the RE model even though the uninformed acted with very little uncertainty.

> HYPOTHESIS 3: Profits of insiders are greater than the profits of the uninformed agents as implied by the prior-information equilibrium prices and allocations given in table 3.
> ALTERNATE HYPOTHESIS 3: Profits of insiders and the uninformed agents converge to equality as implied by the rational-expectations equilibrium prices and allocations given in table 3.

Plott and Smith (1978) used the fraction of consumer plus producer surplus (total profits in these markets) exploited as the measure of efficiency. Application of the concept to these markets is complicated by two considerations. First, under uncertainty the measure must be

ex ante with respect to the information which exists in the market. The following generalization seems to be appropriate:

$$\text{efficiency } (E) = \frac{\substack{\text{total expected returns to allocation, } A,\\ \text{conditioned on available information}\\ \text{in the market}}}{\substack{\text{total expected return of rational}\\ \text{expectations allocation, } \hat{A}, \text{ conditioned}\\ \text{on information in the market}}}$$

$$= \frac{\mathscr{E}R(A \mid I)}{\mathscr{E}R(\hat{A} \mid I)}.$$

If the RE allocation yields the maximum which can be reasonably expected, this measure of efficiency is simply the expected value of the actual allocation taken as a percentage of that maximum.

The second problem involves initial endowments. In the above sense the markets we studied are reasonably efficient even if no trade takes place. In order to emphasize this fact we have constructed a measure which is zero if no trading takes place. Let $\overline{X}$ be the no-trade allocation (i.e., the initial endowments):

$$\text{trading efficiency}(TE) =$$
$$\frac{\mathscr{E}R(A \mid I) - (\text{expected returns of } \overline{X} \text{ conditioned upon information in the market})}{\mathscr{E}R(\hat{A} \mid I) - (\text{expected returns of } \overline{X} \text{ conditioned upon information in the market})}.$$

Because different models sometimes predict different allocations, the efficiency measure of the actual allocation when compared with the predicted efficiency can be used as a measure of the relative accuracy of competing models. Of course this measure defines the efficiency of the RE equilibrium allocation to be 100 percent. The efficiency of the PI equilibrium allocation differs from 100 only when certain events occur. Column 2 of table 4 lists the efficiency of the PI equilibrium allocation (both E and TE) for all information conditions in which such events occur.

> HYPOTHESIS 4: Measures of efficiency (E and TE) converge to the values given in column 2 of table 4 which are implied by the prior-information model.
> ALTERNATIVE HYPOTHESIS 4: E and TE measures of efficiency converge to 100 percent in all cases as implied by the rational-expectations model.

TABLE 4

EFFICIENCY MEASURES WHEN ALLOCATIONS PREDICTED BY TWO MODELS DIFFER

Market and Efficiency Measure	Under PI Equilibrium Allocations (1)	Under Actual Allocations (Listed for Each Period Chronologically) (2)	Under RE Equilibrium Allocations (3)
2 (state X):			
E	42	57, 70	100
TE	−133	−72, −10	100
3 (state Y):			
E	57	79, 88, 89, 98, 99	100
TE	−125	13, 38, 47, 88, 94	100
4 (state Y):			
E	57	92, 95, 93, 94, 94	100
TE	−125	59, 72, 66, 69, 69	100
5 (state X):			
E	67	82, 94, 100	100
TE	−112	−15, 65, 100	100
Mean efficiency:			
E	57	88.3	100
TE	−123	45.5	100

NOTE.—Efficiency measures for market 1 which had probabilistic sample information are not given in order to save space. E = efficiency measure; TE = trading efficiency measure; PI = prior information; RE = rational expectations.

Three ideas must be listed for completeness. First, Keynes (1936) suggests that speculation and the possibility of capital gains can cause prices to be unrelated to underlying returns and information. Taken literally this means that there is no systematic relationship between the state, the underlying returns, and information on one hand and the price and allocations on the other. Because all four pairs of hypotheses discussed above imply systematic relationships, we interpret Keynesian ideas as a rejection of all these hypotheses.

Beja (1976) and Milgrom (1979) recognized that when price is a sufficient statistic for the state individuals who have any uncertainty at all about the underlying state can ignore all their private information, thereby yielding all prices in the range of the equilibrium-price correspondence as equilibria for all states of nature. This theoretical indeterminacy is critically dependent on the existence of some uncertainty in all agents and, unlike the Keynes hypothesis, restricts the range of prices to the equilibrium set. If this indeterminacy is encountered, prices should be at one of the RE equilibrium prices, but beyond that there need be no relationship between the state and the market variables such as price and allocations.

The second is the strong version of the efficient-markets hypothesis

which suggests that prices adjust instantaneously to all available information. Replication of time periods is unnecessary to establish equilibria.

The third is the crystal-ball hypothesis which maintains that market adjustments reflect information beyond that held by agents. It is tempting to assert that the market will somehow know and then apply a model as if it does. This hypothesis would have markets converging to fully revealing equilibrium even when no agent has information beyond prior probabilities.

IV. Results

The time series of prices for all periods of all markets are shown in figures 2 through 6. Price predictions of the PI model when different from those of the RE equilibrium model are shown by the horizontal
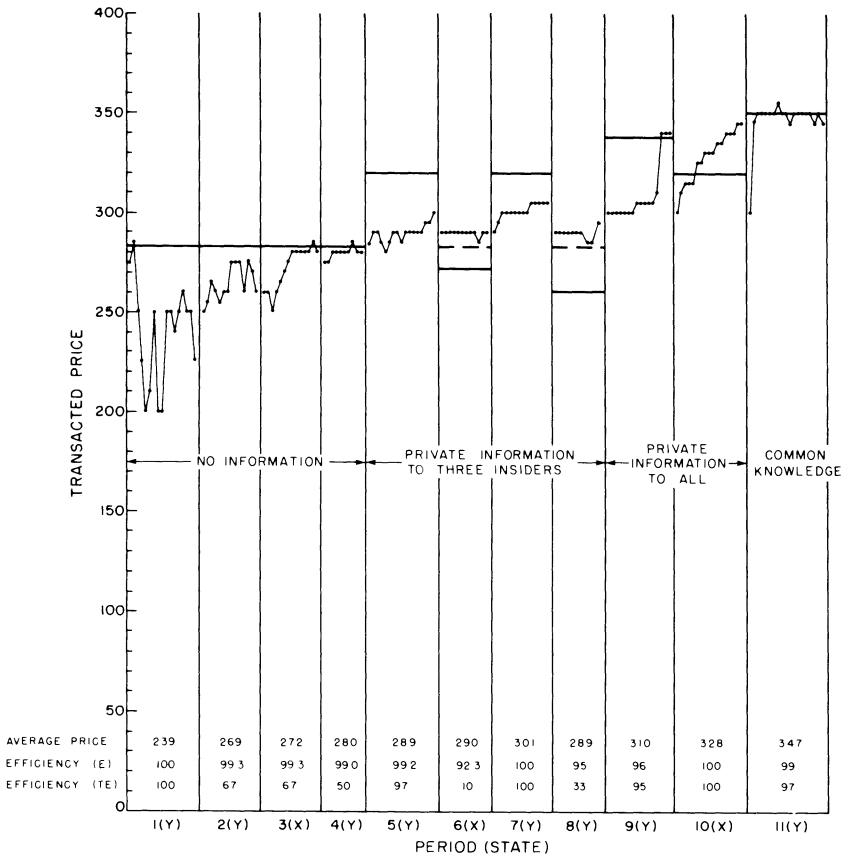

FIG. 2.—Market 1. Time series of contract prices

FIG. 3.—Market 2. Time series of contract prices

dotted line. Predictions of the RE equilibrium model are shown by the horizontal solid line. Each dot represents one trade at the indicated price in chronological order.

Hypotheses about prices, 1 and alternate 1, are tested in two steps, first against each other and then the winner against the random price behavior. Hypothesis 1 as well as the random-behavior hypothesis is rejected in favor of alternate hypothesis 1—that prices converge to the RE predictions.

In step 1, across all periods of all markets, the price predictions of the competing models differed 17 times.[6] In all but four periods average price is closer to the RE equilibrium than to the PI equilibrium. Similarly, in 13 of the 17 periods the average absolute and

[6] Market 1: periods 6, 8; market 2: periods 7, 9; market 3: periods 3, 5, 7, 8, 10; market 4: periods 5, 7, 8, 10, 12; market 5: periods 4, 5, 11.

| AVERAGE PRICE | 234 | 217 | 189 | 333 | 163 | 173 | 364 | 165 | 396 | 167 | 177 | 395 |
| EFFICIENCY (E) | 95 | 98 | 79 | 100 | 88 | 89 | 100 | 98 | 100 | 99 | 100 | 100 |
| EFFICIENCY (TE) | 54 | 82 | 13 | 100 | 38 | 47 | 100 | 88 | 100 | 94 | 100 | 100 |

PERIOD (STATE)

FIG. 4.—Market 3. Time series of contract prices

squared deviations from the RE equilibrium are smaller than the corresponding deviations from the PI equilibrium. Further, and in general, the relative advantage of the RE over the PI model increases upon replication. The few cases where the PI model does better than the RE occur in the early states of experience acquisition.

Market 5 was designed as a check on the robustness of the RE equilibrium model. Because the three-state situation is inherently more complicated, new theories can be developed to compete with the theory of fully revealing prices. In particular, we thought the market might reveal only part of the available information. That is, it may reveal that nature has not chosen some state as opposed to revealing exactly which state was chosen. In the two-state world, revelation of one state which is not chosen is equivalent to complete state revelation. The three-state world is different. The maximum expected

FIG. 5.—Market 4. Time series of contract prices

681

FIG. 6.—Market 5. Time series of contract prices

TABLE 5

NUMBER OF CERTIFICATES IN WRONG HANDS

RATIONAL EXPECTATIONS VERSUS PRIOR INFORMATION HYPOTHESIS

Market	Total Number Certificates		Period												
			1	2	3	4	5	6	7	8	9	10	11	12	13	14
1	18	Realized State	Y	Y	X	Y	Y	X*	Y	Y*	Y	X	Y			
		Rational Expectations	0	2	2	3	1	12	0	8	0	2	1			
		Prior Expectations					8	4	7	4						
2	24	Realized State	X	X	Y	Y	Y	Y	X*	Y	X*	Y	Y			
		Rational Expectations	1	2	0	-1	0	0	22	0	12	0	0			
		Prior Expectations							4	5	23	5	7			
3	24	Realized State	X	Y	Y*	X	Y*	Y*	X	Y*	X	Y*	Y	X		
		Rational Expectations	10	11	18	0	10	9	0	2	0	2	0	0		
		Prior Expectations			15	8	19	19	11	24	6	24				
4	24	Realized State	X	Y	Y	X	Y*	X	Y*	Y*	X	Y*	X	Y*	X	Y
		Rational Expectations	13	16	19	19	7	0	3	5	0	4	0	4	0	17
		Prior Expectations					21	2	21	21	2	21	0	21	2	
5	24	Realized State	Z	X	Z	X*	X*	Y	Z	Z	Y	Y	X*	Y	Z	
		Rational Expectations	11	1	0	13	4	10	0	0	0	0	0	0	0	
		Prior Expectations				11	20	13	5	0	7	9	24	11	14	

*Allocations predicted by the two models are disjoint.   See Table 3.

values of uninformed agents will be 192 for the type II uninformed investors if the information is "not Z," 222 and 262 for type I investors if the information is "not Y" and "not X," respectively. As is shown on the price time series, prices never approach these values in periods in which they might be predicted. The RE model applies to the three-state environment as readily as to the two-state environment.

In the second step, alternate hypothesis 1 was compared with the random-price-behavior hypothesis. Out of 398 nonzero price changes in the five markets, 284 were in the direction of RE equilibrium price. Since under random behavior only one-half of these changes would be expected to move the price in the direction of RE price, the probability of the observed behavior being generated by random chance is practically zero (8.5 SD away from the mean). Similar statistics are obtained by testing for convergence of mean squared error and mean absolute deviation from RE price toward zero upon replication.

Hypothesis 2 can be rejected in favor of its alternative. Each model predicts the type of agents who will hold the certificates. Table 5 lists the number of certificates in the *wrong* hands from the point of view of each model. The RE equilibrium model is clearly superior to the PI model. The models' predictions differ in 36 of the 61 periods in the price markets. In 29 of these 36 periods error from allocations predicted by the RE model is smaller. In 18 of the 36 periods the RE model made no errors at all. The PI model made zero errors in only

TABLE 6

NET PROFIT OF INFORMED AGENTS AS A PERCENT OF
NET PROFIT OF UNINFORMED AGENTS*

Market	Realized State	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1	X			68			147				-53				
	Y	108	88		110	118		118	91	109		98			
2	X	85	105					6208		267					
	Y			107	115	100	103		107		109	105			
3	X	76			100			103		101			100		
	Y		97	230		129	137		102			93	99		
4	X	111			152		178			142		130		117	
	Y		108	109		128		123	118		119		113		90
5	X		121		147	115						105			
	Y						141			102	101		100		
	Z	69		100				157	135					97	

*There was no difference in information provided to various individuals during periods covered
by hatched boxes. Figures are given for control.

two out of these 36 periods. In 17 periods predicted allocations were disjoint (marked by *), and in 12 of the 17 the error rate for the RE model was lower and all five exceptions occurred during the early adjustment periods.

An examination of the trading activities of the uninformed agents in periods when others were informed is particularly useful. When the two models make contradictory predictions about their trading activity, behavior of the uninformed tends to conform to the PI model in early periods, but upon replication, as the experience is gained, the RE model dominates the PI.[7] Note that the behavior of uninformed individuals in contradiction to the PI model implies that these agents were so sure of having learned the state from the endogenous variables that they chose to hold these positions even though they could have avoided all uncertainty by trading, usually very close to the RE equilibrium price.

The profits of insiders relative to the uninformed agents, summarized in table 6, reject hypothesis 3 in favor of its alternative. The table shows, for each period of each market, the percentage ratio of average realized profit per insider to the average realized profit per uninformed agent earned during the period. In all cases the ratios approach 100 as the experience in the market accumulates. The early "jumps" reflect first experiences with a new state. During initial pe-

[7] A table of these data is available from the authors.

riods insiders have the advantage. After replication, however, the advantages of inside information vanish completely.

Both efficiency measures are shown for all periods of all markets in figures 2 through 6. To test the fourth set of hypotheses we concentrate our attention on the set of periods for which the PI equilibrium price is different from the RE equilibrium price. This is a necessary and sufficient condition for the PI equilibrium allocations to be different from the RE equilibrium allocations across agents classified by dividend distributions.

The data in table 4 reject hypothesis 4 in favor of its alternative. The mean efficiencies are much closer to the RE prediction (100) than to the PI prediction of 57 for efficiency and -123 for trading efficiency. Chronologically, as the traders gain experience the efficiency measures move consistently from close to PI prediction toward the RE predictions. Out of 11 changes in efficiency measures repeated in table 4, nine changes are away from PI toward the RE prediction. The probability of obtaining nine out of 11 changes of the same sign by random chance is only 0.039.

We can safely reject the strong RE model, the Keynes model, and the crystal-ball model as general models. Markets take time to adjust to structural changes, so the strong RE model can be rejected. After the RE equilibrium is established, prices adjust instantaneously in response to different realizations of state, but a structural change such as a new and unknown state of nature will not be instantaneously adjusted to. The necessity of adjustment time was also a finding in Forsythe et al. (forthcoming). The early periods of no information adjust closely to the expected value and not to the perfectly informed equilibria, so the crystal-ball hypothesis is discredited.

Our version of Keynes can be rejected because the RE model does so well. If we treat any deviation from RE equilibrium price as evidence of support for the Keynes model, then the Keynes model deserves a closer look. During the no-information periods of the first three markets, prices converged to near the maximum expected value. In market 4 this did not happen. Prices remained well below the maximum expected value for the first four periods while the no-information condition was in effect and returned to the low levels in period 14 when the no-information condition was imposed again. The same phenomena occurred during the first three periods, the no-information condition of market 5. In both cases prices converge to an RE equilibrium price, but it is the wrong price, given the information. The Beja-Milgrom criticism of RE which becomes most relevant when no one is perfectly informed rests on the possible existence of exactly this type of phenomenon.

We suspect the failure to adjust to the maximum expected value

resulted from the fact that no one knew that no one was informed in market 4. This in essence adds at least two subjective states, giving a total of four (X and no one knows, Y and no one knows, X with insiders, and Y with insiders) and thus would increase the number of total replications necessary for the convergence behavior. The large number of holdings on the "wrong" side of the market supports this hypothesis. For the first three periods of market 5 all agents should have known that no one was informed since no clue card was passed out as had been the case in the previous markets in which they had participated, nor was the method of giving clues even explained. We suspect, however, some possible confusion, as subjects speculated on the possible sources of others' information, thereby increasing the state space. Notice that the holdings converged to the RE pattern even though price had not adjusted. In summary, we suspect these uninformed periods provide support for a theory of temporary equilibria as opposed to the possible existence of Keynesian-type phenomena. Given the state of the data and theory, we cannot draw any firm conclusions.

The other period of interest in this regard is period 10 of experiment 1. Notice that price is increasing substantially above the expected value prediction. In all 12 trades above the expected value the buyer was the same agent, and he was an insider. This individual was simply overly optimistic given his information. The fact that this unwarranted (in theory and in fact) optimum did not serve as a lure for uninformed individuals does not reflect well on the Keynesian model. This was a good opportunity for a "bubble," but one did not develop.

V.   Implications for Related Issues

Formation of Rational Expectations

Since no theoretical models predict the path (as opposed to the end point) of the learning process in asset markets, we cannot conduct formal tests of theory. In this section we provide some descriptive material in the hope that generation of such data may help promote development of dynamic models of learning in competitive asset markets.

In our markets, at least two kinds of learning are identifiable. At the beginning of the first period of a market uninformed agents do not know the state, and no agent knows what the market price would be under any given state. From studying the static models one could easily infer that agents must first learn about the equilibrium price and net trade correspondence (the price given the state) and then

learn about the state from market conditions (e.g., price). Applied to the markets studied here this reasoning suggests that an agent must learn first about the structure of the dividends under various states and then learn about the state given the market conditions if (s)he is not an insider.

In the experimental markets the two types of learning do not seem to occur sequentially. Instead they occur simultaneously. The profit data can be used to make the point. The ratio of profits of agents predicted by the RE model to be buyers to those of the predicted sellers reflects the degree of knowledge about the state-price correspondence. Insiders and uninformed agents are represented in equal numbers on both sides of the market (except in market 1 where three out of nine are insiders). Furthermore, initial endowments of (predicted) sellers should be transferred to (predicted) buyers at prices approximating the value of the initial endowments of sellers and which afford buyers very little profit on the transaction. Therefore, profits of buyers and sellers might reasonably be expected to be equal if the equilibrium-price correspondence has been revealed and understood. Thus, as the ratio of profits of buyers to sellers goes to one, we can say agents have learned the state–equilibrium-price correspondence.

Learning about the state, given market conditions, is reflected in the differences between profits of insiders who are perfectly informed about the state (with the exception of market 1, where information is imperfect) and the uninformed agents. If the difference is zero (if the ratio is one), then the informed agents (insiders) and uninformed agents are equally informed.

For each experiment the ratio of buyer and seller profits for each period is shown in table 7. The ratio of profits of insiders and uninformed agents is shown in table 6. As can be seen, both measures converge to near 100 as replications occur and all agents become informed about both correspondences. Notice that the convergence occurs simultaneously. Thus, given these measuring devices, we conclude that learning about the price, given the state, and learning about the state, given the market, occur simultaneously and not sequentially, as the comparative-static models usually assume.

Sources of Agent Information

While these markets are simple relative to naturally occurring markets, behavior in these markets is not simple. The results indicate that agents are receiving accurate information from some source, but the exact source could not be determined. Bids, offers, and contracts were all made publicly by voice interaction. All were written on the

TABLE 7

RATIO OF REALIZED PROFITS OF INVESTORS PREDICTED TO BE BUYERS
BY THE RATIONAL EXPECTATIONS MODEL TO THE PROFIT OF PREDICTED SELLERS

Market	Type of Buyers	Period													
		1	2	3	4	5	6	7	8	9	10	11	12	13	14
1	I					158		149	139	141		103			
	II														
	III						153				530				
2	I					144	112		120		120	114			
	II														
	III							466		172					
3	I				164			113		103			103		
	II														
	III			160		136	129		120		114	109			
4	I						212			190		156		128	
	II														
	III					131		133	127		122		118		137
5	I							269	134					125	
	II						140			108	106		106		
	III				139	116						106			

blackboard. The timing of activities occurred in real time, and on occasion curses, laughter, or other signs of emotion may have been detectable in spite of efforts to control it.

From a formal point of view much more information was available than equilibrium prices. Thus, more information was available than is called for by the efficient-markets hypothesis (Grossman 1978; Jordan 1979). But it was not clear what information was used.

In an attempt to narrow the possibilities, a questionnaire which inquired about the extent to which individuals could identify insiders was circulated after the markets. Recall that the method of distributing inside information did not reveal the identity of insiders. Generally, insiders were better than the uninformed in their ability to identify other insiders. Most uninformed correctly guessed at least one insider if their options were limited to two. We *conjecture,* however, that the ability to identify the insiders is not a necessary condition for convergence in these markets.

The second attempt involved an analysis of bids and offers. Notice in markets 3, 4, and 5 that after some replications the opening contract is near the RE equilibrium price. Somehow information about the state was revealed before any trades took place. In table 8, the proportion of all first-market action, an opening bid or offer, which involved an insider is given for each market for those periods in which insiders exist. The proportion of times an insider is involved with the second, third, sixth, and eighteenth market actions is also given. The cumulative proportions are also given.

In four of the five markets relative activity of insiders decreases

TABLE 8

Relative Activity of Insiders in Early Stages of Markets

Market	Proportion of Nth Market Action Involving Insiders						Cumulative Proportion of Market Actions Involving Insiders Up to Nth Action				
	(1)	(2)	(3)	(6)	(18)		(1)	(2)	(3)	(6)	(18)
1	.75	.00	.50	.25	.00		.75	.39	.42	.42	.39
2	.20	.40	.20	.40	.00		.20	.30	.27	.40	.41
3	.83	.67	.67	.17	.50		.83	.75	.72	.69	.52
4	.89	.89	1.00	.67	.78		.89	.89	.93	.80	.74
5	.80	.50	.80	.44	.60		.80	.65	.70	.67	.61

NOTE.—Market actions shown in parentheses.

with time. It seems as though competition, evidenced by competing bids and offers among insiders during the opening stages of a period, reveals the state to the uninformed.

The exception to the pattern, market 2, is instructive and indicates the limitations of the measure we present. The opening offers in market 2 were made by uninformed individuals. They were also "safe" in the sense that they would have been reasonable sales given that the high-price state existed (recall that in market 2 agents had experience with the environment under conditions of complete information). Similarly, uninformed buyers can make safe opening bids. Information is revealed when someone (presumably an insider attempting to take advantage of the information) drops (increases) the offer (bid) substantially below (above) the safe levels. Thus, it is not simply bids and offers which convey the information, but certain critical bids and offers seem to convey it. Unfortunately, at this time we cannot identify these operationally.

The discussion of bids, offers, trader identification, and other endogenous sources of information suggests that the trading institutions themselves may be important in determining the applicability of the RE models. Institutions can dictate the type of information available to participants. For example, a computerized market which masks bids, offers, and trader identity or even volume may not operate as efficiently as one which does not. The closed book of the specialist is certainly an institution worthy of consideration in this regard.

Statistical Analysis of Price Changes

Even though only single-period securities are traded in these markets, their price data do have some properties typical of stock markets. The price series becomes a fair game as RE equilibrium is established. Further, the transaction-to-transaction log price returns have a serial correlation very close to zero, and their density function is leptokurtic (fat tailed relative to the normal density function).

The fair-game efficient-market model implies that it is not possible to devise trading rules based on past prices (weak-form efficiency) that will earn abnormally high returns on a consistent basis (Fama 1970). We tested five trading rules: (1) buy and hold, (2) trend filter, (3) 1-franc filter, (4) 5-franc filter, and (5) 25-franc filter.[8] For the

---

[8] *Buy and hold:* Buy one certificate at opening transaction price of each period; liquidate at closing transaction price of the period. *Trend filter:* Observe transaction price trend from opening to current price; if positive, buy if necessary to hold one certificate; if negative, sell if necessary to maintain a short position of one certificate. Liquidate at closing transaction price. y-*franc filter:* If transaction price goes up by y or more francs, buy if necessary to hold one certificate until the price goes down by y or more francs, at which time sell if necessary to maintain a short position of one certificate until the price goes up again by y or more francs. Liquidate at closing price.

TABLE 9

PROFIT FROM MECHANICAL TRADING RULES: TESTS OF FAIR-GAME HYPOTHESIS

	Median	Mean	SD	t-Statistic of Excess Profit over Buy and Hold
Buy and hold	10.0	14.9	53.9	...
Trend filter	4.0	16.9	39.5	.36
1-franc filter	3.0	10.5	37.7	−.75
5-franc filter	4.0	11.1	37.4	−.64
25-franc filter	0.0	2.1	33.8	−2.16
Perfect information about RE price	10.0	26.1	40.3	2.09

NOTE.—The data for period 1 of market 4 were excluded in calculating the means.

single-period security used in these markets, equilibrium return over time is zero. The naive buy-and-hold strategy yields a median return of 10 francs per period and a mean return of 14.9 francs per period, which is significantly positive (see table 9). This return is generated during the early periods and declines to zero as RE equilibrium is approached. Ordinary least-squares regression of buy-and-hold returns (BH_t) earned each period on the absolute difference between equilibrium and average trading price (d_t) yields the following estimates:

$$BH_t = -1.86 + 0.609d_t, R^2 = .14;$$
$$(-0.22) \quad (3.12)$$

D-W statistic = 1.94 and t-statistics are given in parentheses. The results suggest that the profits generated by the buy-and-hold strategy are largely the product of disequilibrium trading.

Trend and 1- and 5-franc filters perform about as well as the buy-and-hold strategy; the 25-franc filter performs worse than the buy-and-hold. As equilibrium is approached, the returns from these filters also approach zero.

Returns from using the perfect advance knowledge of RE equilibrium price[9] are given in the last row of table 9. These returns are equal to or greater than the buy-and-hold strategy and indicate that the knowledge of equilibrium price is useful for trading in this market even though this price is not always attained. According to the Keynes model this knowledge would not be useful for a single investor.

The first-order serial correlation of log price changes $(\log_e P_t/P_{t-1})$ is very close to zero. Serial correlation in data for individual trading

[9] Any time current price is below the equilibrium price hold a long position of one certificate by buying at the current price if necessary; any time the current price is above the equilibrium level hold a short position of one certificate by selling, if necessary.

periods ranged from -0.65 to 0.60 with a median of 0.00 and mean of -0.022 (SD $= 0.30$). The magnitude of serial correlation does not seem to be affected by the existence of disequilibrium trades, as is indicated by the following regression of estimated serial correlation for each period ($\hat{\rho}_t$) on the absolute deviation of average trading price from the RE equilibrium price d_t:

$$\hat{\rho}_t = -0.0357 + 0.00038d_t;$$
$$\qquad(-0.83)\qquad(0.75)$$

t-statistics are given in parentheses, the coefficient of determination is 0.01, and D-W statistic is 1.91.

The frequency distribution of relative price changes in our markets is leptokurtic (fat tailed) relative to the normal distribution. Sample kurtosis (fourth moment around mean divided by squared variance) in all five markets, 13.7, 19.3, 52.9, 56.3, and 30.3, respectively, is much greater than the kurtosis of normal distribution, which is 3. Relatively large proportions of observations in the peak of the distribution of price changes derive from periods in which trading is at or close to the equilibrium price, while the larger price changes in the tails of the distribution come mostly from trading away from equilibrium price. This observation might lead one to speculate that the leptokurtosis of security-price returns is generated by trading at prices far from the equilibrium price and that trading closer to the equilibrium price would tend to bring the distribution closer to the normal. However, preliminary tests on our data do not support such a conclusion. For each market we identified periods when much of the trading was and was not close to equilibrium price by visual examination of price series and calculated the kurtosis of price changes observed in each class of periods. In three of the five markets the kurtosis of price changes observed in close-to-equilibrium trading periods alone was greater than that for far-from-equilibrium trading periods.

VI. Conclusions

Given time and replications these markets behave substantially as predicted by RE equilibrium models. There seems to be no doubt that variables endogenous to the operation of these markets served to convey accurately the state of nature to otherwise uninformed agents. We can conclude that the RE models must be taken seriously as not universally misleading about the nature of human capabilities and markets. Of course, only additional research will determine to what extent the experiences acquired in these simple markets can be extended to more complicated naturally occurring markets.

The exact variable(s) which serve(s) to inform the uninformed about the underlying state may differ according to the structure of the market institutions. Market institutions differ widely according to the nature of the endogenous variables observable by participants. Thus, institutional features might have implications for both the existence of RE equilibria and the speed with which they are attained. In these double oral auctions the key variables are not yet identified, but we suspect a knowledge of unaccepted bids and offers, in addition to price, is of primary importance.

Appendix

Markets were conducted in three steps: (1) training with the mechanism used to draw states of nature, (2) explanation of procedures and rules of the market, (3) conduct of markets for several periods.

Step 1:   Training with Mechanisms Used to Draw the States of Nature
Instruction Set 1 was distributed and read out loud. On the table between the subjects and the experimenters was kept a bingo cage with the appropriate number of balls. Subjects had the opportunity to observe the operations of the devices for many draws. Following this, subjects were asked to predict the outcome, with the incentive structure described in the instructions, for about 10–20 draws until most, though not all, subjects predicted the state with the highest relative frequency based on the bingo-cage numbers. No mention was made of probabilities.

Instruction Set 1
Each year we draw a ball from a bingo cage containing thirty-six balls numbered one through thirty-six. If the ball drawn is numbered one through twelve, outcome of the draw is called X; if a ball numbered thirteen through thirty-six is drawn, the outcome is called Y. [10]
You have to predict the outcome of each draw before it is announced. If your prediction is correct, you win \$0.25; if wrong, you lose \$0.10. Before the first draw is made, record your prediction by circling either X or Y in the first row of the enclosed sheet. After you have encircled one letter, the outcome will be announced and you should record the announced outcome in the blank space on the same row of the table. If your prediction is correct, circle the amount shown in the Win column, otherwise circle the amount shown in the Lose column.
Once you have recorded your prediction you must not make a change; any erasure will invalidate your prediction. At the end, add up your total winnings and losses and record the difference (net winnings or losses) at the bottom right corner of the sheet (see fig. 7).

Step 2:   Explanation of Procedures and Rules of the Market
Instruction Set 2 was distributed and read aloud. The experimenter illustrated a sequence of hypothetical transactions on the blackboard so each subject would understand how transactions were to be recorded on the record

---

[10] The numbers in this first paragraph were altered appropriately for each experiment; see table 2 for parameters.

Subject No._____

Number	Circle One Decision		Outcome X or Y	Win ($)	Lose ($)
1	X	Y	_____	0.25	-0.10
2	X	Y	_____	0.25	-0.10
.					
.					
.					

Total winnings _____

Total losses _____

Net winnings/losses _____

FIG. 7

sheet and how his/her profit would be calculated. The example was designed to minimize its normative effect on subsequent bidding behavior. Importance of accurate records of all transactions was emphasized.

Instruction Set 2

General.—This is an experiment in the economics of market decision making. Various research foundations have provided funds for this research. The instructions are simple, and if you follow them carefully and make good decisions, you might earn a considerable amount of money which will be paid to you in cash.

In this experiment we are going to simulate a market in which you will buy and sell certificates in a sequence of market years. Attached to the instructions you will find a sheet, labeled information and record sheet, which helps determine the value to you of any decisions you might make. You are not to reveal this information to anyone. It is your own private information.

The type of currency used in this market is francs. All trading and earnings will be in terms of francs. Each franc is worth *$0.003* to you. Do not reveal this number to anyone. At the end of the experiment your francs will be converted to dollars at this rate, and you will be paid in dollars. Notice that the more francs you earn the more dollars you earn.

Specific Instructions

Your profits come from two sources—from collecting certificate earnings on all certificates you hold at the end of the year *and* from buying and selling certificates. During each market year you are free to purchase or sell as many certificates as you wish, provided you follow the rules below. For each certificate you hold at the end of the year you will be given one of the two numbers of francs listed on row 19 of your information and record sheet. Note that earnings may be different for different investors.[11] The method by which one of the two numbers is selected each year is explained later in these instructions. Compute your total certificate earnings for a period by multiplying the earnings per certificate by the number of certificates held. That is, (number of certificates held) × (earnings per certificate) = total

[11] In instructions for market 2, this sentence was eliminated. In market 1, it had been replaced by "Notice that these amounts may differ from period to period."

certificate earnings. Suppose, for example, that you hold five certificates at the end of year 1. If for that period your earnings are 100 francs per certificate (i.e., the number selected from row 19 is 100) then your total certificate earnings in the year would be $5 \times 100 = 500$ francs. This number should be recorded on row 19 at the end of the year.

Sales from your certificate holdings increase your francs on hand by the amount of the sale price. Similarly, purchases reduce your francs on hand by the amount of the purchase price. Thus you can gain or lose money on the purchase and resale of certificates. At the end of each year all your holdings are automatically sold to the experimenter at a price of 0.

At the beginning of each year you are provided with an initial holding of certificates. This is recorded on row 0 of the year's information and record sheet. You may sell these if you wish or you may hold them. If you hold a certificate, then you receive "earnings per certificate" at the end of the year. Notice therefore that for each certificate you hold initially, you can earn during the year *at least* the amount shown as "earnings per certificate." You earn this amount if you do not sell that certificate during the year.

In addition, at the beginning of each year you are provided with an initial amount of francs on hand. This is also recorded on row 0 of each year's information and record sheet. You may keep this if you wish or you may use it to purchase certificates.

Thus at the beginning of each year you are endowed with holdings of certificates and francs on hand. You are free to buy and sell certificates as you wish according to the rules below. Your francs on hand at the end of a year are determined by your initial amount of francs on hand, earnings on certificate holdings at the end of the year, and by gains and losses from purchases and sales of certificates. All francs on hand at the end of a year in excess of *10,000* francs are yours to keep. These are your profits for the year.

Information about Dividends

Whether the dividend you receive from the certificates you hold is the X-dividend shown on row 19 or the Y-dividend on row 19 is determined by the experimenter at the beginning of the year by drawing a ball from a bingo cage containing forty balls numbered one through forty. If the ball drawn is numbered one through sixteen, X-dividend is paid; if the ball drawn is numbered seventeen through forty, Y-dividend is paid.[12]

At the beginning of each year, before trading starts, each investor will receive a clue card which will carry one of the following three: (i) X, (ii) Y, (iii) a blank. If your clue card carries an X, the dividend paid at the end of that year will be the X-dividend; if your clue card carries a Y, Y-dividend will be paid. A blank card tells you nothing about whether the X or the Y dividend will be paid.[13]

Trading and Recording Rules
(1)   All transactions are for one certificate at a time. After each of your sales or purchases you must record the TRANSACTION PRICE in the appropriate column depending on the nature of the transaction. The first

_____

[12] States and numbers in this paragraph were altered according to parameters for each experiment given in table 2.

[13] This paragraph describing the clue cards was not used in markets 1 and 2 and was altered to include state Z in market 5.

Trader No. _____

INFORMATION AND RECORD SHEET

YEAR _____

	Transaction Number	Transaction Price Sale	Purchase	Certificates on Hand	Francs on Hand
Beginning of the Year Holdings	0	/////////////////////////////// ///////////////////////////////			
	1				
	2				
	⋮				
	18				
X-Dividend _____	19	Total Certificate Earnings Dividend Rate on Hand at the End of the Year			
Y-Dividend _ _					
	20	Total Francs on Hand at the End of the Year			
	21	Less: Fixed Cost			
	22	End of Year Net Profit			

Transfer this amount
to your Profit Sheet

FIG. 8

transaction is recorded on row (1), and succeeding transactions are recorded on subsequent rows (see fig. 8).

(2)   After each transaction you must calculate and record your new holdings of certificates and your new francs on hand. Your holdings of certificates may never go below zero. Your francs on hand may never go below zero.

(3)   At the end of the year record your total certificate earnings in the last column of row 19. Compute your end of period totals on row 20 by listing certificate holdings and adding total certificate earnings to your francs on hand.

(4)   At the end of the year, subtract from your francs on hand the amount listed in row 21 and enter this new amount on row 22. This is your profit for the market year and is yours to keep. At the end of each market year, record this number on your profit sheet (see fig. 9).

(5)   At the end of the experiment add up your total profit on your profit sheet and enter this sum on row 15 of your profit sheet. To convert this number into dollars, multiply by the number on row 16 and record the product on row 17. The experimenter will pay you this amount of money.

Market organization.—The market for these certificates is organized as follows. The market will be conducted in a series of years. Each period lasts for seven minutes. Anyone wishing to purchase a certificate is free to raise his or her hand and make a verbal bid to buy one certificate at a specified price, and anyone with certificates to sell is free to accept or not accept the bid. Likewise, anyone wishing to sell a certificate is free to raise his or her hand and make a verbal offer to sell one certificate at a specified price. If a bid or offer is accepted, a binding contract has been closed for a single certificate, and the

PROFIT SHEET

Row	Market Year	Profit
1	1	
2	2	
⋮	⋮	⋮
14		
15	Total profit (in francs)	
16	Dollars per franc	0.003
17	Total dollars profit	

Name _____

FIG. 9

contracting parties will record the transaction on their information and record sheets. Any ties in bids or acceptance will be resolved by random choice. Except for the bids and their acceptance, you are not to speak to any other subject. There are likely to be many bids that are not accepted, but you are free to keep trying. You are free to make as much profit as you can.

Step 3:   Conduct of Markets
Seven minutes were permitted for each period, with warnings at five minutes, six minutes, and six-and-a-half minutes. In markets 1, 2, and 5 information cards were not distributed in no-information periods. In markets 3 and 4 blank information cards were distributed in no-information periods. These information cards were distributed at the beginning of each period. The bingo cage was rotated in full view of the subjects, a ball was drawn, and the appropriate information cards (prepared in advance) were distributed. A log of bids, offers, and transactions of the latest four or five was kept on the blackboard. In market 5 a photocopy of the log of each period was distributed to subjects at the end of each period.

References

Aumann, Robert J. "Agreeing to Disagree." *Ann. Statis.* 4, no. 6 (1976): 1236–39.
Beja, A. "The Limited Information Efficiency of Market Processes." Research Program in Finance Working Paper no. 43, Univ. California, Berkeley, 1976.
Fama, Eugene F. "Efficient Capital Markets: A Review of Theory and Empirical Work." *J. Finance* 25 (May 1970): 383–417.
Forsythe, Robert; Palfrey, Thomas R.; and Plott, Charles R. "Asset Valuation in an Experimental Market." *Econometrica*, forthcoming.
Friedman, James W. "An Experimental Study of Cooperative Duopoly." *Econometrica* 35 (July–October 1967): 379–97.
Green, Jerry R. "Information, Efficiency and Equilibrium." Discussion Paper no. 284, Harvard Univ., Harvard Inst. Econ. Res., March 1973.

Grossman, Sanford J. "Further Results on the Informational Efficiency of Competitive Stock Markets." *J. Econ. Theory* 18 (June 1978): 81–101.

Harrison, J. Michael, and Kreps, David M. "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations." *Q.J.E.* 92 (May 1978): 323–36.

Jordan, James S. "On the Efficient Markets Hypothesis." Mimeographed. Minneapolis: Univ. Minnesota, 1979.

Keynes, John Maynard. *The General Theory of Employment, Interest and Money.* New York: Harcourt Brace, 1936.

Lucas, Robert E., Jr. "Expectations and the Neutrality of Money." *J. Econ. Theory* 4 (April 1972): 103–24.

Milgrom, Paul. "Rational Expectations, Information Acquisition, and Competitive Bidding." Discussion Paper no. 406, Northwestern Univ., Center Math. Studies Econ. and Management Sci., 1979.

Plott, Charles R. "The Application of Laboratory Experimental Methods to Public Choice." In *Collective Decision Making: Applications from Public Choice Theory*, edited by Clifford S. Russell. Washington: Resources for the Future, 1979.

Plott, Charles R., and Smith, Vernon L. "An Experimental Examination of Two Exchange Institutions." *Rev. Econ. Studies* 45 (February 1978): 133–53.

Smith, Vernon L. "Experimental Economics: Induced Value Theory." *A.E.R. Papers and Proc.* 66 (May 1976): 274–79.