CHICAGO JOURNALS

Information Mirages in Experimental Asset Markets
Author(s): Colin Camerer and Keith Weigelt
Reviewed work(s):
Source: *The Journal of Business*, Vol. 64, No. 4 (Oct., 1991), pp. 463-493
Published by: The University of Chicago Press
Stable URL: http://www.jstor.org/stable/2353290
Accessed: 04/01/2013 16:30

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at
http://www.jstor.org/page/info/about/policies/terms.jsp

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.

The University of Chicago Press is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Business.*

http://www.jstor.org

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

Colin Camerer
Keith Weigelt
University of Pennsylvania

Information Mirages in Experimental Asset Markets*

One explanation for the apparent volatility of asset prices is that people overreact to trades that are uninformative, creating self-generated information "mirages." We test whether mirages occur in experimental asset markets. There are insiders in only half the periods, so traders cannot be sure if the trades of others reveal information. We observed four clear miages in 47 periods without insiders. Mirages always occurred early in an experimental session; in later periods, traders learn whether there are insiders by observing nonprice information such as the speed of trading, and mirages occurred only temporarily.

I. Introduction

We investigate behavior in laboratory asset markets in which traders are sometimes informed of asset values. We test whether traders overreact to uninformative trades, mistakenly inferring information from them. The existence of price "mirages," caused by such mistakes, might explain why asset prices seem to be excessively volatile.

A. *Information Aggregation*

Empirical evidence that asset markets react swiftly and reliably to incorporate new information into prices is given an underpinning by theories of "information aggregation" (see, e.g., Grossman 1976, 1981; Grossman and Stiglitz 1980; and Jordan 1982). In these theories, traders have diverse information about the value of assets. If traders know how information affects prices, they can infer what information is causing observable price changes. Prices effectively re-

* The financial support of the Richard Paget Research Chair at Northwestern University's Kellogg Graduate School of Management and the New York University's Center for Entrepreneurial Studies are gratefully acknowledged. Helpful comments have been received from some anonymous referees, Jon Baron, Jim Jordan, David Kreps, Howard Kunreuther, Charles Plott, Shyam Sunder, and seminar audiences at New York University, University of Minnesota, University of Pennsylvania, and the Public Choice Society. Marcy Lynch and James Myers are also acknowledged for their research assistance.

(*Journal of Business*, 1991, vol. 64, no. 4)
© 1991 by The University of Chicago. All rights reserved.
0021-9398/91/6404-0002$01.50

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

veal the information, aggregating the diverse information of traders. In equilibrium, traders can learn nothing more from prices. Prices are said to fully reveal information if prices are those that would occur if traders pooled information openly.

Because the temporal distribution of traders' private information is usually not known, testing information-aggregation models with natural data is difficult. Joint hypothesis tests are possible if assumptions are made about the distribution of information (e.g., Schwert 1981; Huberman and Schwert 1985).[1] However, if the joint hypothesis is rejected, one can always blame the assumption about information distribution rather than the information aggregation model.

Because one can control the temporal distribution of information in market experiments, they are well suited to testing predictions about information aggregation. In previous experiments, prices did reveal information in simple settings (Plott and Sunder 1982; Friedman, Harrison, and Salmon 1984; cf. Ang and Schwarz 1985). In complex settings the results are mixed. Plott and Sunder (1988) and O'Brien and Srivastava (in press) found poor information aggregation. Forsythe and Lundholm (1990) and Camerer and Chernew (1987) found successful aggregation, and Copeland and Friedman (1987*a*) found mixed results. Successful information aggregation in complex settings seems to depend on traders having a common dividend structure and extensive experience (e.g., two or three experimental sessions), or on diverse information being distributed sequentially.

B. Information Mirages

In most previous experiments, uninformed traders knew with certainty that there were insiders. Their only mental challenge was inferring from price signals what insiders knew. In our sessions there are periods with no insiders; uninformed traders cannot be sure whether there are insiders or not in a particular period. Therefore, traders could falsely infer information from price signals when there are no insiders and prices do not convey information. Once one trader makes such an error, she may trade as if informed, causing other traders to infer mistakenly that she is an insider. We call the price path that results from these errors an information mirage because prices make traders see information which is not really there.

Many tests suggest stock and bond prices are too volatile to be consistent with rational reaction to news (see Camerer 1989; LeRoy

1. They test whether stock and bond prices respond to Consumer Price Index (CPI) announcements. Assuming that information on all the commodity prices in the CPI is known by consumers before the CPI is announced, stock and bond prices—which are affected by inflation (Fama 1981)—should not respond to the CPI announcement. They find that prices largely anticipate the announcement, but there is some market response when the announcement is made.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

1989). The most well-known evidence of excessive volatility comes from the "variance bounds" tests of LeRoy and Porter (1981), Shiller (1981, 1986), and others, but those tests are controversial (e.g., Marsh and Merton 1986). Better tests (Campbell and Shiller 1987; Roll 1988; West 1988) also suggest excess volatility.

Asset prices are also much more volatile when markets are open for trading than when markets are closed (Oldfield and Rogalski 1980; French and Roll 1986).[2] For instance, the per-hour variance of returns on all New York Stock Exchange (NYSE) and American Stock Exchange (AMEX) stocks during trading hours is about 70 times as large as the per-hour variance over weekends (French and Roll 1986). Trading-hour returns may be more volatile simply because more public or private information is generated during trading hours. Other evidence casts doubt on this appealing view.[3] French and Roll thus consider the possibility that "self-generating" trading causes part of the volatility during trading hours (see also Black 1986): "Imagine that the market opens in the morning and, just by chance, the first ten orders are on the sell side. Other traders, who have no information about the asset except that the first ten orders were to sell, might draw the reasonable conclusion that something bad has happened to the company. They might be induced to sell also and, if they are, the market price should fall. The price drop might attract the attention of other traders who extrapolate its recent path and sell, which attracts other traders who sell, and so on"(French and Roll 1984, p. 17).

That is precisely what we mean by an information mirage. Grossman (1989) advances a similar explanation for the October 1987 stock market crash: "Ex-post, we label an event a 'panic' when a group of investors has shifted out of equities for noninformational reasons, and this shift has caused substantial numbers of other investors to shift out of equities because they think that the price has moved for informational reasons" (Grossman 1989, pp. 7–8).

Cornell and Shapiro (1989) describe an unusual 2-month period in which some U.S. Treasury bonds appeared to be mispriced relative to comparable issues. The best explanation of the mispricing is that the bond price rose because Japanese investors bought many of the bonds

2. Price movements are not observed when markets are closed. But one can calculate an implied volatility by comparing the change from one day's closing price to the next day's opening price.

3. The most convincing data are from weeks in 1968 when the exchanges were closed on Wednesdays due to a paperwork backlog. Assuming the same amounts of public and private information are generated on these closed Wednesdays as on Wednesdays that were open for trading, the return variances of weeks containing closed Wednesdays should be the same as those of weeks containing open Wednesdays. The actual variances from closed-Wednesday weeks are 82.1% as large as those from open-Wednesday weeks. This suggests that trading is self-generating, but the sample of closed Wednesdays is too small to be conclusive. However, Hertzel, Kendall, and Kretzmer (1990) report corroborating results from currency markets.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

and squeezed American dealers, who had sold bonds short, when the dealers tried to buy back bonds to cover their short positions. But they also hint at a mirage-type explanation in their conclusion (p. 309).

The anecdotal evidence is suggestive, but searching more thoroughly for mirages in natural data is problematic. Short-lived mirages will not be detectable if data are weekly or daily. Even tests with trade-by-trade data do not identify mirages clearly unless one knows whether traders are reacting to trades or to information. In our sessions, we can tell whether changes are mirages because we control the flow of information: if we observe prices changing dramatically when there is no new information, we know we have seen a mirage.

II. Design of the Experiment

Experimental design parameters are summarized in table 1. Sessions had either 12 or 9 subjects ("traders"), recruited through solicitations in classes and posted sign-up sheets. Traders were Wharton undergraduate students (sessions 1 and 5), New York University (NYU) evening MBA students (sessions 2–4), and Northwestern University MBA students (sessions 6–7). Only the traders in session 1 had experience in previous experiments.

Traders were endowed with two assets at the beginning of each period and with 10,000 francs of working capital that was repaid at the end of the period. Traders could hold their assets or trade them for francs in a double-oral auction. At the end of the trading period, assets paid dividends to the traders who held them. After the dividends were paid, the assets expired.

Traders belonged to one of three possible dividend types. There were four traders of each type (except in sessions 6–7, where there were three). Traders' types determined the amount of dividends the asset would pay them in each state. (The differences in dividends across traders are like differences in taxes, wealth, risk tastes, or liquidity.) Traders knew their own dividends, but all they knew about others was that dividends may be different for different traders.

A. *Earnings*

All trading and earnings were in terms of francs, which were converted to dollars at the end of the session, at a rate of $.002 per franc. At the end of each trading period the state was announced, dividends were paid, and traders calculated their profits for the period. Profits were determined from summing dividend revenue and subtracting purchase prices.

Traders could not sell more shares than they owned,[4] and their net

4. Otherwise, traders could collectively generate infinite surplus since different traders have different dividends.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

TABLE 1 Experimental Design Parameters

Session	Length (Minutes) and Number of Trading Periods	Number of Subjects and Status	Number of Possible Insiders	Prior Probability of States		Trader Types	Dividend (Francs)		Expected Dividend
				G	*B*		Good State (*G*)	Bad State (*B*)	
1	(6)	12				I	375	100	265
	15	Wharton undergraduate	0,6	.6	.4	II	340	150	264
						III	225	175	205
2	(6)	12				I	375	100	265
	15	NYU MBA	0,6	.6	.4	II	340	150	264
						III	225	175	205
3	(6)	12				I	375	100	265
	16	NYU MBA	0,6	.6	.4	II	340	150	264
						III	225	175	205
4	(4)	12				I	375	100	265
	21	NYU MBA	0,6	.6	.4	II	300	100	220
						III	225	175	205
5	(4)	12				I	375	100	265
	20	Wharton undergraduate	0,6	.6	.4	II	300	100	220
						III	225	175	205
6	(6)	9				I	350	135	221
	16	Northwestern MBA	0,3,6	.4	.6	II	275	150	200
						III	100	160	136
7	(6)	9				I	375	100	210
	17	Northwestern MBA	0,3,6	.4	.6	II	240	150	186
						III	100	175	145

NOTE.—NYU = New York University.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

francs on hand could never be negative. At the end of the session, traders calculated their earnings and were paid in cash. Demand for assets is thus induced by demand for dollars (Smith 1976).

B. State Information

Except for a few "warm-up" periods (called state *W*), in which no state was determined, the state (*G* or *B*) was determined before trading began by a public drawing from a bingo cage, with replacement. A public drawing from a second bingo cage determined the number of insiders. The results of these draws were not announced until after the period was over.

Before trading began, each trader was shown a private clue card, based on the results of the draws. The clue cards had either *N* (no information) or *G* or *B*, indicating which state had been drawn. Traders with informative clue cards (*G* or *B*) are insiders. Which traders, if any, were insiders each period was randomly predetermined and varied each period.[5]

In sessions 1–5, there was a 50% chance in each period that all traders would get *N* cards (called state *N*), and there was a 50% chance that half the traders would get informative cards and half would get *N* cards. The number of insiders was always evenly distributed among dividend types. In sessions 6–7 there were equally likely to be zero, three, or six insiders.

Subjects met in a classroom and were read instructions (which are virtually identical to those in Plott and Sunder (1982). After several practice draws from the bingo cage used to determine the state, the market experiment began. Sessions lasted for approximately 3 hours. Trading periods were 6 minutes long (4 minutes in sessions 4–5).

The probabilities of the states and the number of insiders were common knowledge. Traders did not know the dividend values of other traders or which traders, if any, had inside information. At the end of each period, the number of insiders and the state were announced.

The market was organized as a double-oral auction: buyers shouted out bids to buy, and sellers shouted out offers to sell. Bids had to top outstanding bids, and offers had to undercut outstanding offers. A matching bid and offer was a trade, and a trade erased all previous bids and offers. Bids, offers, and trades were recorded on a transparency, along with the identifying numbers of the traders, so traders had a perfect history of the period's previous market activity. No history of previous periods of trading was posted.

5. We did not make the same traders insiders every period, as in some experiments (e.g., Plott and Sunder 1982), because insiders could then deduce whether there were any insiders from their own lack of information. Also Banks (1985) observed that information is fully revealed even when different subjects are insiders each period.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

III. Competing Hypotheses about Prices and Allocations

We distinguish two extreme hypotheses, *RE* and *PI*.

RE. Traders supplement private information with the information of others from price signals. Since the inside information is perfect in the experiment, prices fully reveal information.[6]

PI. Traders do not learn from price signals; they use only their private information about the state.

It is hard to be formal about theories intermediate between *RE* and *PI* (though see Copeland and Friedman 1987*a*, 1987*b*, 1991), but we note one example. In Jordan's (1982) dynamic "temporary equilibrium," traders make bids and offers based on private information before observing price signals; then, as they observe price signals, they learn others' information and they converge to *RE*. Plott and Sunder (1982), Copeland and Friedman (1987*a*, 1987*b*), and others have noticed that actual trading in experiments does seem to follow a "Jordan path" of convergence to *RE*, both across and within periods.[7]

Each trader's endowment of francs is large enough to buy virtually the entire market supply of assets, and the supply is fixed (by the initial endowment and the short-selling restriction). Thus, there is excess demand at any price less than the highest expected value; in competitive equilibrium, prices will be bid up to the highest expected value (or certainty equivalent, if risk neutrality is not assumed). Of course, there is no guarantee that competitive equilibrium will result because the double-oral auction is not Walrasian. But recent game-theoretic models (Friedman 1984; Wilson 1985) and hundreds of experiments with the double-oral auction (Smith 1982, pp. 944 ff.) suggest it does generally converge to competitive equilibrium.

Hypotheses *RE* and *PI* make differing predictions about what the highest expected values, hence prices, will be (table 2), and about what types of traders will hold units at those prices (see table 3). The predictions of *RE* are straightforward: Prices will equal the highest expected dividend value of any trader, conditioned on the union of all traders' information. Traders with the highest expected dividend (conditioned on all information) will hold all units.

6. Full revelation creates a well-known paradox: if traders can learn everything from prices, they have no incentive to invest in gathering information. To create such an incentive, prices must be "partially revealing" because of noise (Grossman and Stiglitz 1980; Hellwig 1980; Diamond and Verrecchia 1981) or because profitable trades reveal the information in a non-Walrasian setting (Dubey, Geanakoplos, and Shubik 1987) or when the informed trader is a monopolist (Kyle 1985). In our experiments, the paradox does not arise because information is costless to traders.

7. Jordan's model does not strictly apply to experimental data because he assumes trades do not take place at temporary prices (Jordan 1982, pp. 246–47). However, a Jordan path will result if traders ignore capital gains or are myopic and regard each period as the last (Kobayashi 1977).

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

TABLE 2 **Predicted Prices (in Francs) under Private Information (*PI*) and Rational Expectations (*RE*) Theories**

	No Information (State *N*)		Inside Information (State *G*)		Inside Information (State *B*)	
Session	*PI*	*RE*	*PI*	*RE*	*PI*	*RE*
1–3	265	265	375	375	265	175
4–5	265	265	375	375	265	175
6	221	221	350	350	221	160
7	210	210	375	375	210	175

TABLE 3 **Type of Trader Predicted to Hold Units by Each Theory, in Each State**

	State		
Theory	*G*	*B*	*N* or *W*
Sessions 1–3:			
RE	I*i*, I*u*	III*i*, III*u*	I, II
PI	I*i*	I*u*, II*u*	I, II
Sessions 4–7:			
RE	I*i*, I*u*	III*i*, III*u*	I
PI	I*i*	I*u*	I

NOTE.—I*u* denotes uninformed type I traders, I*i* denotes informed type I traders, etc.

Hypothesis *PI* predicts that prices will be the highest expected dividend value, conditioned only on private information; traders with the highest expected dividend (conditioned on private information) will hold all units.

A. Numerical Examples of Predictions

The predictions are complicated. Examples from sessions 1–3 might help explain them.

In states *N* and *W* there is no information, except the prior probabilities, about which state occurred. The highest expected value, 265,[8] will be the equilibrium price (assuming risk neutrality).

In state *G*, type I traders have the highest valuation (375). Under *PI*, only informed type I traders (I*i*) will hold units; uninformed type I traders (I*u*) will think the assets are worth their expected value of 265 and will sell them to type I*i* traders. Under *RE*, however, uninformed traders learn that the state is *G*, so *RE* predicts that both type I*i* and I*u* traders will hold assets. Note that both theories predict a

8. We assume the difference between 264 and 265 (.2 cents) is empirically meaningless.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

price of 375, but they differ in their predictions of which traders will hold assets (*RE* predicts I*u* and I*i*, *PI* predicts only I*i*).

We devised a novel test of predicted allocations.[9] Table 4 illustrates our test for sessions 1–3. Suppose the state is *G*. Then *RE* predicts type I*i* and I*u* traders will hold all units, and *PI* predicts only I*i* traders will hold units. Since both the *RE* theory and the *PI* theory predict that type I*i* traders will hold units, we call the I*i* traders *RE* + *PI* traders. Since *RE* + *PI* traders are predicted by both theories to buy units, their behavior provides no evidence about whether *RE* or *PI* is a superior theory (*RE* + *PI* behavior does provide evidence of how both theories predict in general). Since the only theory that predicts that type I*u* traders will buy units is the *RE* theory, we call type I*u* traders *RE*-only traders.

By defining the groups of traders this way, we can test the degree of rational expectations.[10] First we can ask whether *RE*-only traders are buyers or sellers, on average: do uninformed traders learn enough about the state to get on the correct side of the market and buy? Hypothesis *RE* predicts they will; *PI* predicts they will not.

Then we can ask whether *RE*-only traders buy as much as *RE* + *PI* traders: do uninformed traders learn well enough to buy as much as informed traders? This is a more stringent test of *RE*. Prices might fully reveal the state even if uninformed *RE*-only traders do not have the time or insight to buy units.

In *B* periods, the theories make completely different predictions. In sessions 1–3, *PI* predicts that uninformed type I*u* and II*u* traders will never learn the state is *B*, so they will buy all the units at the uninformed expected value price (265). Hypothesis *RE* predicts everyone will learn the state and type III traders will buy the units at a price of 175. In our test, Type I*u* and II*u* traders are *PI*-only traders, type III traders are *RE*-only traders, and there are no *RE* + *PI* traders (i.e., no traders are predicted by both theories to buy).

In *N* and *W* periods in sessions 1–3, both theories predict type I and II traders will buy units, so these traders are *RE* + *PI* traders and there are no *RE*-only and *PI*-only traders.

B. Predicted Prices and Allocations during Mirages

In *N* periods, nobody is informed, but traders may not know that. (If one trader's clue card says *N*, she cannot be sure if all others' cards

9. In most previous experiments (e.g., Plott and Sunder 1982), allocation predictions have been tested by counting the number of units held by traders predicted to hold them. For instance, in *G* periods one would count the number of units held by I*i* traders (to test *PI*) and the number of units held by I*i* and I*u* traders (to test *RE*). Such a comparison always favors *RE* over *PI* because the traders predicted by *PI* to hold units are a subset of the traders predicted by *RE* to hold units. Our test eliminates this bias.

10. The number of units these traders are predicted to buy or sell requires some calculation. These numbers are shown in an earlier draft of this article, available from us.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

TABLE 4 Allocation Predictions

Session	Trader Group	Purchase Predictions for State *G* If: *PI* True	Purchase Predictions for State *G* If: *RE* True	Purchase Predictions for State *B* If: *PI* True	Purchase Predictions for State *B* If: *RE* True	Purchase Predictions for States *N* and *W* If: *RE* + *PI* True
1–3	*RE* only	sell (I*u*)	buy (I*u*)	sell (III)	buy (III)	n.a.
	RE + *PI*	buy (I*i*)	buy (I*i*)	n.a.	n.a.	buy (I,II)
	PI only	n.a.	n.a.	buy (I*u*,II*u*)	sell (I*u*,II*u*)	n.a.
4–7	*RE* only	sell (I*u*)	buy (I*u*)	sell (III)	buy (III)	n.a.
	RE + *PI*	buy (I*i*)	buy (I*i*)	n.a.	n.a.	buy (I)
	PI only	n.a.	n.a.	buy (I*u*)	sell (I*u*)	n.a.

NOTE.—I*u* denotes uninformed type I traders, and I*i* denotes informed type I traders, etc. n.a. = no traders in a group; *RE* only (*PI* only) denotes traders only predicted by *RE*(*PI*) to buy units; *RE* + *PI* denotes traders predicted by both theories to buy units.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

say N, too, or if half the other traders are informed.) Thus, it is possible for some trades away from the expected-value equilibrium to spark a mistaken belief by traders with N cards that other traders know something—a mirage.

Mirages might drive prices to the G price (375 in our example sessions 1–3) or the B price (175). In a true G mirage, only type I traders will be holding assets because traders will have mistakenly inferred that the state was G, and type I traders have the highest expected value conditioned on that mistaken inference. Thus, if we observe high prices (near the G price of 375) but type I and type II traders are both holding assets, the high prices may not be due to a mirage. Similarly with B mirages: In a true B mirage, prices will be low and type IIIs will hold units. If prices are below 265 (as they often are) but type Is and IIs hold, we can guess the low prices are due to risk aversion or something other than a mirage. Thus, we use both prices and allocations to judge whether mirages occurred.

IV. Results: *RE* versus *PI*

First, we will compare the *PI* model and the *RE* model. Most of these comparisons replicate earlier experiments with similar designs. (Readers interested only in the existence of mirages should skip ahead to Section V.)

Several tests show that *RE* predicts better than *PI*, especially in later market periods when traders have some experience.

A. Analysis of Prices

Figures 1–7 show the time series of prices in each experimental session. Each point represents one trade. Two lines show the different prices predicted by *PI* and *RE* (in B periods). One line shows the common price prediction of *PI* and *RE* (in G, W, and N periods). The period number and state are shown at the bottom of each graph.

Prices converge roughly to equilibrium levels, especially after traders have some experience. In B periods with insiders, prices fall to the *RE* equilibrium level later in the session, but not early on.

The accuracy of the theories can be measured by the mean squared error (MSE) of each theory's price predictions, summarized in table 5. To measure convergence within each session, as subjects learn, the set of periods of each type was divided in half. The first half was called "early" and the second half was called "late."

Table 5 shows that, when *RE* and *PI* predict the same prices (in N and G periods), errors decline by almost half between early and late

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

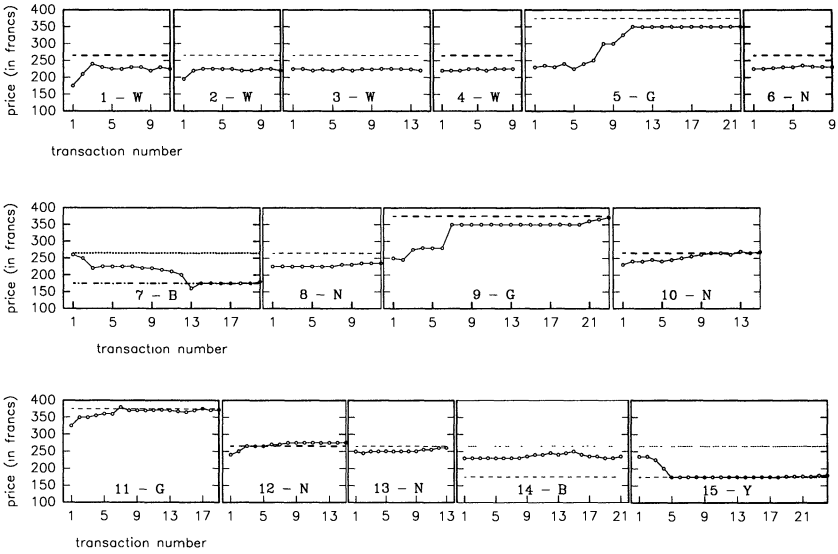

FIG. 1.—Time series of transaction prices, session 1. ○ Transaction price; ——— *RE* + *PI*; –·—·— *RE* only; ·––––· *PI* only.

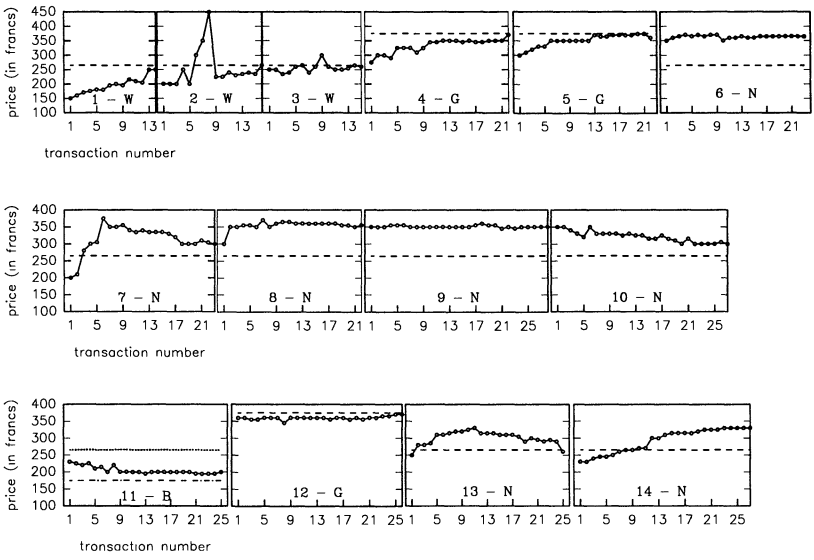

FIG. 2.—Time series of transaction prices, session 2. ○ Transaction price; ——— *RE* + *PI*; –·—·— *RE* only; ·––––· *PI* only.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

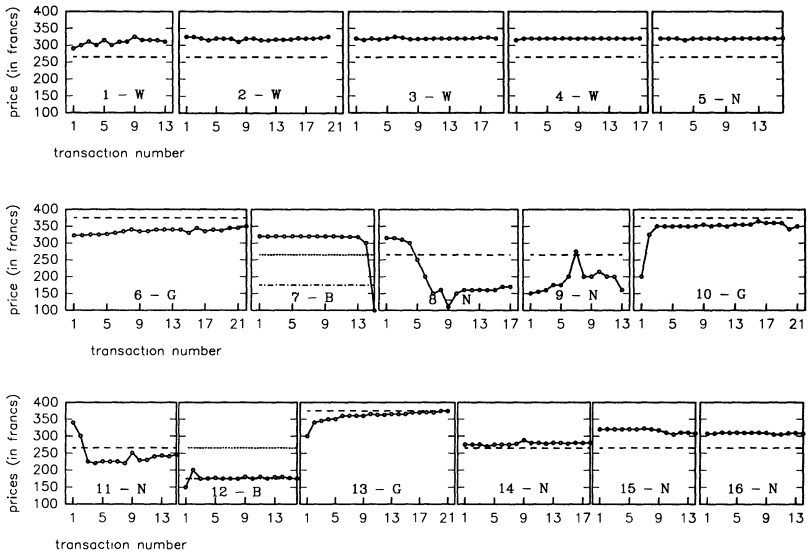

FIG. 3.—Time series of transaction prices, session 3. ○ Transaction price; ——— *RE* + *PI*; –·—·— *RE* only; ·––––· *PI* only.

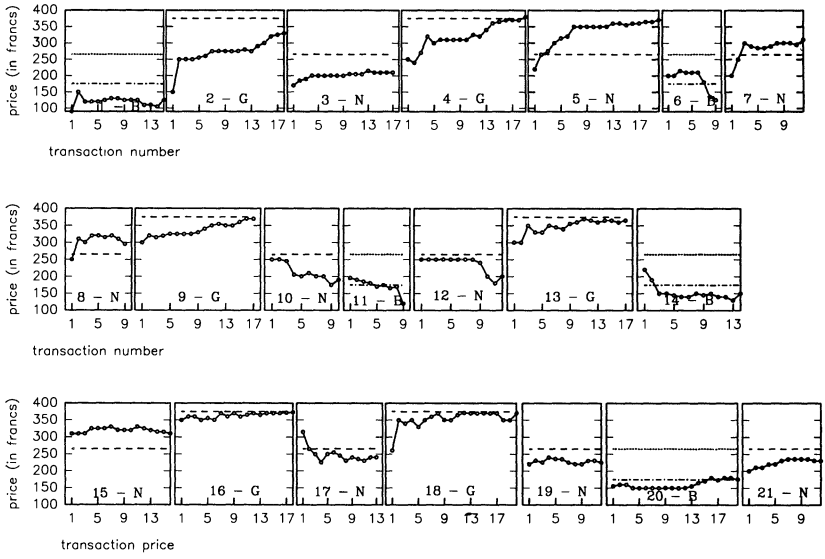

FIG. 4.—Time series of transaction prices, session 4. ○ Transaction price; ——— *RE* + *PI*; –·—·— *RE* only; ·––––· *PI* only.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

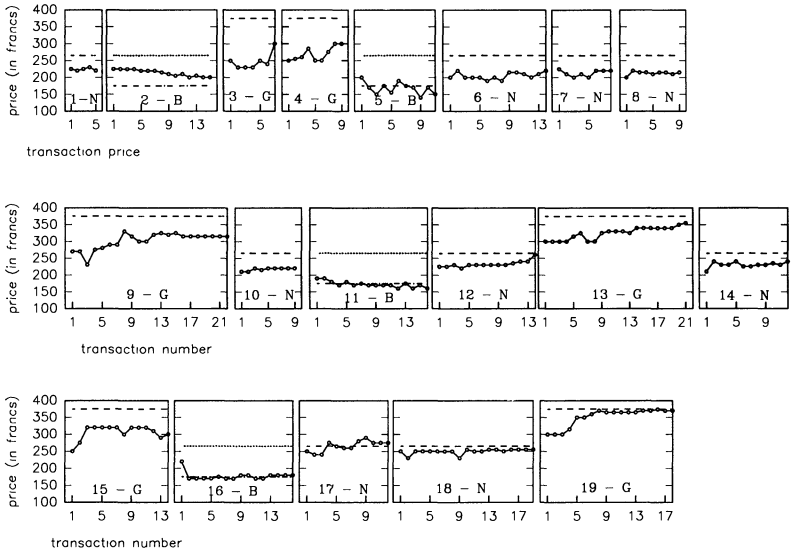

FIG. 5.—Time series of transaction prices, session 5. ○ Transaction price; ——— *RE* + *PI*; –·—·— *RE* only; ·––––· *PI* only.

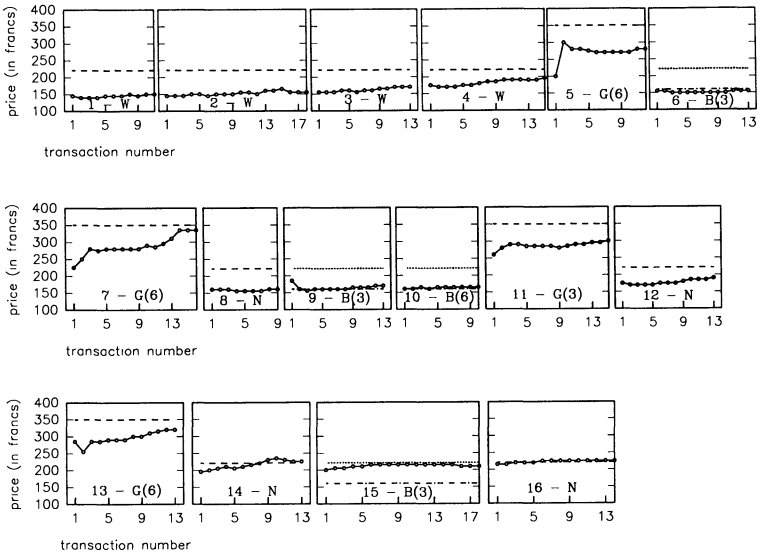

FIG. 6.—Time series of transaction prices, session 6. ○ Transaction Price; ——— *RE* + *PI*; –·—·— *RE* only; ·––––· *PI* only.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

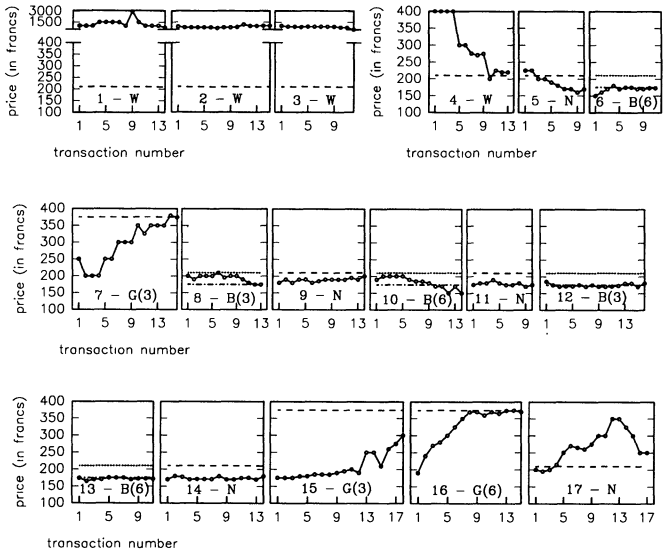

FIG. 7.—Time series of transaction prices, session 7. ○ Transaction price; ——— *RE* + *PI*; –·—·— *RE* only; ·–––– *PI* only.

periods. In *B* periods, *RE* errors get much smaller from early to late (36.0–13.7) while *PI* errors get larger (66.3–70.7).[11]

In *G* periods, *RE* and *PI* predict the same price, so we distinguish the two with a novel profitability test. The test examines the trading behavior of uninformed type I traders (I*u*) because *RE* predicts an I*u* trader will buy units for 375 while *PI* predicts he will sell for 265. Using this difference in predictions, we can compare a trader's profit per trade under the two theories. For example, in sessions 1–3, if an I*u* trader sold a unit at 300, his profit under *PI* would be 300 − 265, or 35 francs. Under *RE*, his profit would be 300 − 375, or −75 francs. Similarly, if an I*u* trader bought a unit at a price of 300 francs, his profit under *PI* would be 265 − 300, or −35 francs, while under *RE*, profit would be 375 − 300, or 75 francs.

The *RE* and *PI* ways of calculating profits are similar to competing accounting schemes. The difference in the accounting schemes is how the assets are valued. We assume that traders are motivated to earn profits. It turns out that traders' profits are consistently negative when

11. To measure convergence within a period, we counted the fraction of times that the MSE declined between the first and second halves of the period (Friedman, Harrison, and Salmon 1984). These fractions are given in the bottom of table 5. The hypothesis that within-period improvements in *RE* predictions are Bernoulli trials (with $p = .5$) is rejected in late periods.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

TABLE 5 Mean Square Error (MSE) from Predicted Price Level for Each Theory

		N Periods		*G* Periods		*B* Periods			
	W Periods	Early Periods	Late Periods	Early Periods	Late Periods	Early Periods		Late Periods	
Session	*RE* + *PI*	*RE* + *PI*	*RE* + *PI*	*RE* + *PI*	*RE* + *PI*	*PI*	*RE*	*PI*	*RE*
1	(4) 43.5	(3) 29.9	(2) 12.3	(2) 67.7	(1) 18.6	(2) 47.8	50.7	(1) 83.3	20.8
2	(3) 52.5	(4) 84.6	(3) 54.0	(2) 40.0	(1) 16.1	(1) 60.5	32.6		
3	(4) 52.2	(4) 66.8	(3) 36.0	(2) 42.6	(1) 22.6	(1) 67.0	139.9	(1) 80.3	9.1
4		(5) 56.5	(5) 40.4	(3) 72.5	(3) 27.1	(3) 108.1	37.1	(2) 110.0	10.3
5		(5) 51.0	(4) 25.4	(3) 105.0	(3) 53.5	(2) 75.5	29.3	(2) 90.8	10.6
6	(4) 61.2	(2) 53.7	(2) 8.4	(2) 75.3	(2) 61.4	(2) 63.2	8.4	(2) 33.8	28.0
7	(4) 703.7	(3) 31.1	(2) 55.8	(2) 135.3	(2) 73.8	(3) 34.6	14.2	(2) 37.3	4.5
Mean	(19) 189.6	(26) 55.1	(21) 34.6	(16) 78.4	(12) 41.3	(14) 66.3	36.0	(10) 70.7	13.7
Fraction of periods with second-half MSE < first-half MSE	.684*	.654*	.714**	1.00**	1.00**	.214**	.642	.600	1.00**

NOTE.—Early (late) denotes periods in the first (second) half of the periods in each session. The number of periods in each cell is in parentheses.
* $p < .10$.
** $p < .05$ by two-tailed binomial test of $p = .50$.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

accounted for in the *PI* way and positive when accounted for in the *RE* way. We conclude they are valuing assets using the *RE* value rather than the *PI* value (i.e., we suspect they have learned *RE*).

B. Analysis of Trading Patterns

Table 6 shows average net purchases per trader across groups of sessions with identical designs. (The types of traders predicted to buy by the *RE* and *PI* theories are shown in table 4.)

Trading data do not support convergence to *RE* as strikingly as the price data do. For instance, in *G* periods, the *RE*-only traders are net buyers (contrary to the *PI* prediction that they will sell), but they buy only a unit or so while the *RE* + *PI* traders buy about six units each. When the state is *B*, *RE*-only traders sell and *PI*-only traders buy in early periods—evidence for *PI* over *RE*—but the pattern is reversed in late periods. An index of the degree of information aggregation indicates trading activity is roughly halfway between *PI* and *RE*.[12]

It is curious that prices indicate full information revelation while trading patterns indicate only partial revelation. Copeland and Friedman (1987*a*) observed the same discrepancy. Perhaps uninformed (*RE*-only) traders learned the state too late in the period to make profitable trades.

C. Analysis of Allocative Efficiency

Since trading patterns do not strongly support *RE*, we now examine whether misallocation of units was costly in terms of traders' expected dividends. We do this by calculating allocative efficiency,[13] defined generally (Plott and Smith 1978) as the fraction of available surplus earned by traders.

Call the total dividends earned according to each theory $D(RE)$ and $D(PI)$, call actual dividends $D(\text{actual})$,[14] and call dividends earned if there is no trading $D(NT)$. Our measure of allocative efficiency for *RE* is

$$\text{Eff}(RE) = [D(\text{actual}) - D(NT)]/[D(RE) - D(NT)],$$

and for *PI* is

$$\text{Eff}(PI) = [D(\text{actual}) - D(NT)]/[D(PI) - D(NT)].$$

12. The index takes the number of purchases by each trader minus the number predicted by *PI*, divided by the difference between the number of units predicted by *RE* and *PI*. The index varies from zero (if *PI* is true) to one (if *RE* is true). In our data, the index ranges between .3 and .6. Experiments by others yield comparable results.

13. Note that allocative efficiency is different from market efficiency, which measures the speed and reliability with which prices incorporate information.

14. In *N* periods we used expected dividends instead of actual dividends.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

TABLE 6 Mean Net Purchases per Trader in Each Trader Group

Session	State	Number of Periods	Early Periods			Number of Periods	Late Periods		
			RE + *PI*	*RE* Only	*PI* Only		*RE* + *PI*	*RE* Only	*PI* Only
1–3	*W*	11	.62			0			
	N	11	.57			8	.93		
	G	6	7.10	−.60		3	8.17	.17	
	B	4		−.31	1.63	2		3.00	−1.00
4–5	*W*					0			
	N	10	1.08			9	2.20		
	G	6	4.59	−.67		6	5.92	1.00	
	B	5		.35	.60	4		3.00	−1.56
6–7	*W*	8	1.34			0			
	N	5	.33			4	2.25		
	G	4	5.67	3.50		3	5.25	3.38	
	B	5		−.87	5.00	4		−.08	2.88
Means	*W*	19	.92			0			
	N	26	.72			21	1.73		
	G	16	5.80	.40		12	6.32	1.39	
	B	14		−.27	2.47	10		1.77	.33

NOTE.—*RE* only (*PI* only) denotes traders only predicted by *RE* (*PI*) to buy units; *RE* + *PI* denotes traders predicted by both theories to buy units.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

These efficiencies will be zero if there is no trading, and one if trading is exactly as predicted by the theory. Negative efficiencies indicate trading in the opposite direction of that predicted by the theory. Note that since *RE* always predicts the largest possible dividends earnings, Eff(*RE*) is also a measure of overall allocative efficiency.

Table 7 summarizes the mean efficiencies in early and late periods of all sessions. In *N* and *G* periods, where Eff(*RE*) = Eff(*PI*), efficiencies always improve from early to late and are quite high, between .75 and 1.00. Even though *RE*-only traders are not buying as many units as *RE* predicts (as noted in the previous section), their failure to do so is not harming allocative efficiency much.

In *B* periods, *RE* is worse than *PI* in early periods and much better than *PI* in late periods.

In sum, we have compared the accuracy of *RE* and *PI* on prices, trading patterns, and allocative efficiencies. When *PI* and *RE* agree, their predictions are generally accurate. When *RE* and *PI* disagree, *PI* is accurate in early periods, but as traders gain experience, *RE* becomes better. Specific results are remarkably close to those in other experiments.[15]

V. Results: Mirages

The main purpose of our experiment is to test whether information mirages occur in a controlled setting.

To identify mirages, we examined the holdings data and price path in each *N* period. Examining mirages that are apparent in hindsight may be fruitful, even if these mirages are rare and could not be forecast (by subjects or by us).

We looked for mirages in *N* periods in which the maximum or minimum price deviated substantially from the expected-value price, and the holdings pattern was consistent with the pattern predicted in a *G* or *B* period.

We call mirages "sustained" if trading activity (prices and holdings) resembled activity in insider periods throughout the mirage period, and "temporary" if activity resembled insider-period activity during only part of the period.

Table 8 shows the four sustained mirages we observed in 47 *N* periods, along with some descriptive statistics. (A table for temporary mirages is available from us.) Sustained mirages do occur, but they are not common.

15. In their *B* periods, Plott and Sunder (1982) found that efficiency (*TE*) averaged 46%, while in late periods ours averaged 41%. The mean squared errors of *RE* predictions were about 41 francs in our late *G* periods, and 14 francs in late *B* periods. The analogous errors in Friedman, Harrison, and Salmon's (1984) session 5 were about 4 francs and 14 francs, but their periods were 15 minutes long.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

TABLE 7 **Mean Allocative Efficiencies as a Fraction of Maximum Gains from Exchange under Each Theory**

Sessions	State	Number of Periods	Early Periods			Number of Periods	Late Periods		
			RE	*RE + PI*	*PI*		*RE*	*RE + PI*	*PI*
1–3	*W*	11		.597		0			
	N	11		.700		8		.935	
	G	6		.889		3		.959	
	B	4	−.090		−.002	2	.686		−.550
4–5	*N*	10		.342		9		.596	
	G	6		.363		6		.930	
	B	5	.022		−.325	4	.746		−.100
6–7	*W*	8		.538		0			
	N	5		.295		4		.718	
	G	4		.846		3		1.00	
	B	5	−.334		.273	4	−.067		.028
Means	*W*	19		.572		0			
	N	26		.484		21		.748	
	G	16		.681		12		.955	
	B	14	−.137		−.019	10	.409		−.499

NOTE.—The efficiency is the same in states *W*, *N*, and *G* because *RE* and *PI* have the same allocations in these periods.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

TABLE 8 Ex Post Description of Mirages in *N*-State Periods

Session, Period	Number of Units Held, Number of Traders Holding Any Units, by Dividend Type			Prices				Location of Maximum (Minimum) Price in Trading Period	
	I	II	III	Open	Mean	Close	Maximum (Minimum)	Elapsed Time (Minutes)	Fraction of Previous Trades
2, 6	19,3*	2,1*	3,1	350	363	365	370	2.0–3.0	.29
Means, nonmirage *N* periods ($N = 3$)	11,2.7*	12.7,2.7*	.3,.7	310	317	287	342	1.5–2.3	.32
Means, insider *G* periods ($N = 3$)	22.7,3*	1.0,1.0	.7,.7	312	350	360	371	5.0–6.0	.96
3, 8	3,3*	4,1*	17,3	315	200	170	(110)	3.0–4.0	.53
Means, nonmirage *N* periods ($N = 5$)	8.2,2.4*	13.6,3.6*	2.2,1.2	312	293	292	(283)	2.4–3.4	.48
Means, insider *B* periods ($N = 2$)	3.5,1.5	6.5,1.5	14.0,3.5*	235	241	138	(125)	2.5–3.2	.53
4, 5	22,4*	0,0	2,2	220	335	370	370	3.5–4.0	1.00
Means, nonmirage *N* periods ($N = 7$)	14.1,3.9*	5.6,2.7	4.3,1.9	229	246	244	269	1.9–2.4	.57
Means, insider *G* periods ($N = 6$)	21.0,3.5*	1.5,0.5	1.5,0.8	268	333	365	366	2.9–3.5	.70
6, 8	2,1*	9,3	7,2	60	158	160	(155)	1.0–2.0	.44
Means, nonmirage *N* periods ($N = 2$)	16.5,2.5*	1.5,1.0	0.5,0.5	195	201	208	(193)	1.0–1.0	.04
Means, insider *B* periods ($N = 4$)	9.5,1.5	5.0,1.7	3.5,1.5*	175	173	175	(168)	1.0–2.0	.10

* Trader type predicted to hold units by *RE* and *PI*.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

For example, in mirage period 6 of session 2 the mean price was 363 (the fully revealing *G* price was 375). The maximum price of 370 came between the second and third minutes of trading, after 29% of the trades in the period had been made.[16] Nineteen of the 24 units were held by type I traders (three different type I traders held units), and only two units by type IIs. Since the mean holding by type I traders was 11 units in nonmirage periods, and 22.7 units in *G* periods, the holdings pattern in the mirage period suggests type I traders believe the state is *G*.

Note that there are mirages which resemble both states, *G* and *B*. The restriction against short selling, which might be thought to force prices upward and contribute to *G* mirages, cannot explain *B* mirages.

The four sustained mirages always occurred in the first or second *N* period in a session and never occurred twice in a session. We think sustained mirages in these sessions are disequilibrium phenomena resulting from inexperience with *N* periods.

It is useful to compare the price paths during sustained mirages with those in insider periods and in nonmirage *N* periods. The dotted line in figure 8 shows the time series of trade prices from one apparent mirage (period 5 of session 4), with the time of trades on the *x*-axis. Solid lines show the lower envelope of price paths in all *G* periods, and the upper envelope of *B* prices from the same session.[17] Unconnected dots represent trades in nonmirage *N* periods of session 4. Note that most trades in *N* periods take place within the *G* and *B* envelopes.

Subjects in session 4 could reasonably conclude that prices should be at least as high as the upper envelope (upper solid line) if the state was *G* and as low as the lower envelope (lower solid line) if the state was *B*. Bayesian subjects might then decide whether the state was *G* or *B* or *N*, based on whether the observed price path was above or below either boundary. Viewed this way, a mirage is simply an error (perhaps an unsystematic one) in Bayesian inference (cf. Aoki and Friedman 1986).

The paths in figure 8 show how the mirage was sustained after it began. After the first minute of trading, prices moved into the range of high prices that were typical of *G* periods with insiders (the dotted line crossed the upper solid line). Subjects then guessed the state was *G*—prices had "revealed" it!—and began bidding even higher. By the

16. Reaching a maximum in the middle of the period is evidence against a sustained mirage since prices usually rise or fall monotonically when insiders are present. However, in period 6 of session 2, the drop from the maximum of 370 to the closing price of 365 is small.

17. To construct the *G* boundary, we plotted the price paths in all the *G* periods and took the lower bound at each point in time as the composite boundary. The *B*-period boundary is the same, except upper bounds were used (and we ignored period 7 in constructing fig. 8 because it was so unusual).

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

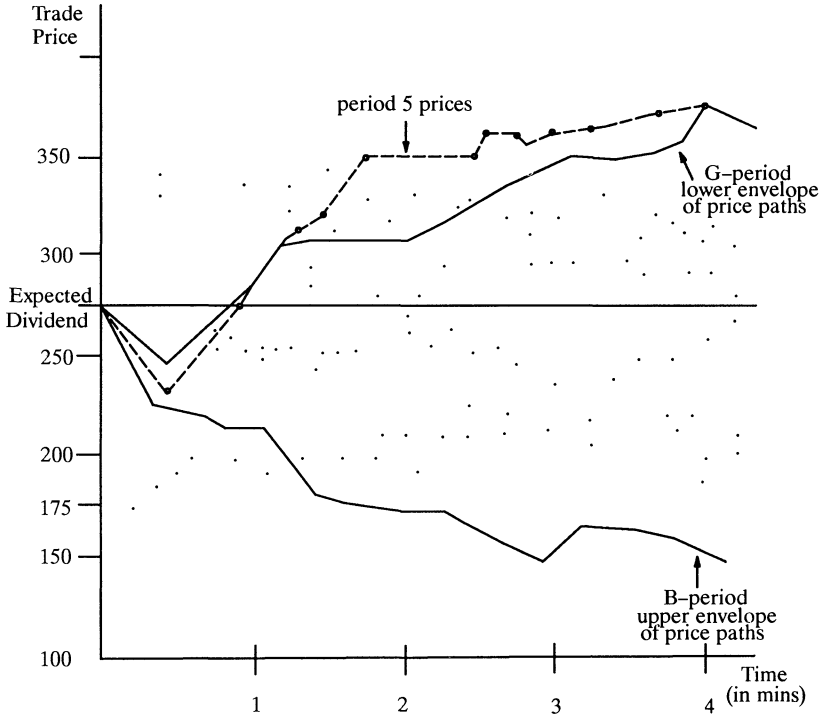

FIG. 8.—A sustained mirage price path (period 5, state *N*, in session 4). Unconnected points indicate time and price of trades in *all N* periods (except period 5) in session 4.

second minute of trading, the mirage price path was well above the typical *N*-period prices (unconnected dots). Traders grew more and more confident that there were insiders who knew the state was *G*.[18]

The dotted line in figure 9 shows the price path of a sustained mirage in period 8 of session 3. Between the second and third minutes, prices fell outside the usual range of *N*-period prices (unconnected points) into the range of *B*-period prices. Subjects then inferred, wrongly, that insiders knew the state was *B* and prices fell further.

The trading that began these mirages is mysterious. Careful study of the bids, offers, and trades early in the mirage path did not reveal any stylized facts or patterns explaining why mirages begin in some periods but not in others.

Since many mirages occurred in the first *N* period, a simple psychological theory is that mirages are caused by traders overgeneralizing

18. Note that mirage period 5 was the first *N* period in the session. Since traders had not yet seen the data from other *N* periods (unconnected dots), their Bayesian inference problem is especially difficult.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

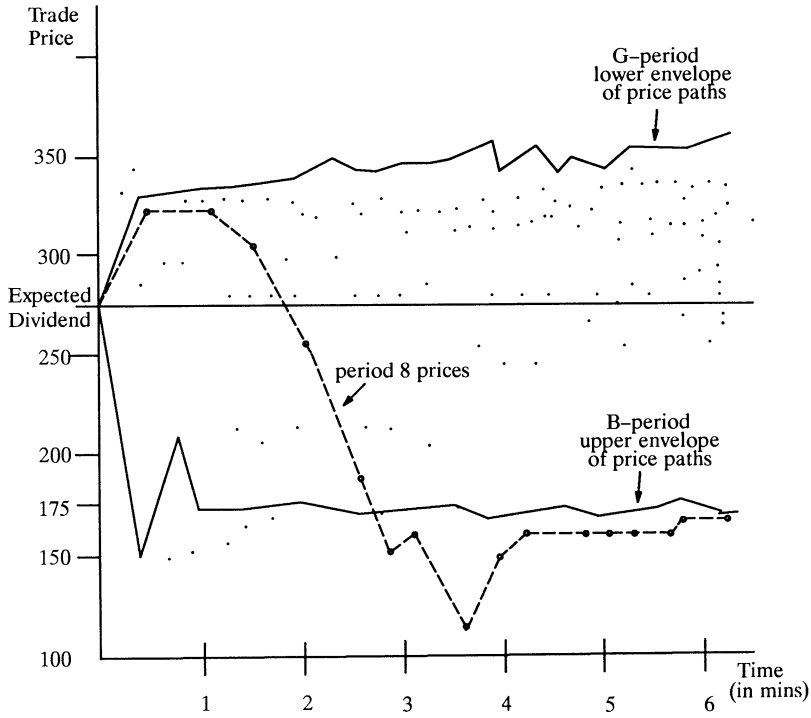

FIG. 9.—A sustained mirage price path (period 8, state *N*, in session 3). Unconnected points indicate time and price of trades in *all N* periods (except period 8) in session 3.

from their limited experience in preceding periods with insiders. The price series in figure 3, for instance, show the possible influence of period 7 (a *B* period with insiders) on the mirage that occurred in period 8. The *B* information in period 7 was revealed in the last minute, by an uninterrupted series of offers plummeting from 310 to 100. In period 8 there was a similar series of offers between the second and the third minutes, from 350 to 200, which falsely signaled *B* information. The similarity (or "representativeness") of that price drop to the drop in the previous period may have caused traders to underweigh the prior probability (.2) that period 8 was a *B* period and overestimate the chance that it was a *B* period (Arrow 1982; Tversky and Kahneman 1982; Camerer 1987, 1990).

Indeed, every sustained mirage follows an earlier period with insiders with a similar price path,[19] which lends support to the representativeness theory.

19. See periods 5 and 6 in session 2, periods 4 and 5 in session 4, and periods 6 and 8 in session 6.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

A. *The Role of Nonprice Information*

Sustained mirages never occurred in later periods. We conjecture that as traders get experience, they learn to distinguish insider and nonsider periods using nonprice information. One kind of nonprice information, nonverbal cues, has been ruled out by computerized experiments that yield similar patterns of information revelation; see DeJong et al. (1987).

For instance, trading usually occurs more swiftly and with less preliminary bidding and offering when there are insiders. Since insiders are not monopolisits and the trading period is finite, their incentives are those of a Prisoner's Dilemma: Each wants to trade first because initial trades with uninformed traders are the most profitable but leak information (which makes later trades less profitable). Insiders will thus trade early in the period and with less time-consuming bidding and offering (relative to the amount of trading).

Table 9 summarizes trading intensity—the number of bids and offers per trade—and the average number of trades in each time interval. It is easy to distinguish insider and noninsider periods using these data. When there are insiders, trading is less intense and more frequent early in the period (especially in the first half-minute).

In the first half-minute trading intensity was usually greater than 10 in N periods, and less than 10 in G or B periods. A trader who intuited this rule could count bids, offers, and trades and figure out whether there were insiders or not, thus preventing mirages. But for some reason, the rule does not work in sessions 2 and 4; in those sessions, more mirages occurred (two sustained and five temporary mirages in 17 N periods).[20]

In sessions 4–5 we tried to test the theory that nonprice information prevents mirages by reducing the length of periods from 6 minutes to 4 minutes. We thought forcing traders to trade more quickly in all periods might harm their ability to distinguish between noninsider and insider periods and cause more mirages. The change worked in session 4 (with NYU MBAs), but not in session 5 (with Wharton undergraduates).

One can imagine other design changes to reduce useful nonprice information and cause mirages or enhance information and get rid of them. For example, a referee suggested inducing early trading in noninsider periods by paying subjects a commission for trading quickly. Would their frantic trading create mirages?

von Borries and Friedman (1988) ran experiments in which only one

20. In all other sessions there were two sustained and four temporary mirages in 30 N periods. The difference in proportions of sustained or temporary mirages, 7/17 and 6/30, has a z-statistic of 1.59 ($p = .056$) in a two-sample binomial test, normally approximated.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

TABLE 9 Pace and Intensity of Trading

Session	State	Number of Periods	Trading Intensity (Minute Intervals)			Pace of Trading (Minute Intervals)			Average Number of Trades in Period
			0–.5	.5–1	1–2	0–.5	.5–1	1–2	
1	*N*	6	19.18	14.64	3.28	.33	.33	3.00	13.50
	G or *B*	6	4.38	3.84	4.37	2.17	2.67	2.67	21.67
2	*N*	7	6.10	3.99	4.60	1.57	1.96	3.29	24.71
	G or *B*	5	6.00	2.39	4.05	1.60	3.30	4.40	24.40
3	*N*	7	18.60	10.06	5.26	.43	.21	3.29	15.14
	G or *B*	5	5.57	3.78	3.55	1.40	1.30	4.40	19.00
4	*N*	10	10.41	6.25	4.81	1.20	1.30	3.60	13.20
	G or *B*	9	10.00	5.48	4.61	1.50	2.50	5.89	15.91
5	*N*	9	20.21	5.22	5.07	.33	1.00	3.22	11.33
	G or *B*	10	5.31	4.50	4.36	1.60	2.20	3.60	14.90
6	*N*	4	18.00	4.71	3.46	.50	1.75	3.25	12.25
	G or *B*	8	5.31	4.42	3.61	2.00	1.50	3.50	13.75
7	*N*	5	16.50	5.00	5.20	.40	1.20	2.00	13.20
	G or *B*	8	8.71	6.75	4.29	.88	1.00	3.00	14.50

NOTE.—Trading intensity = (number of bids and offers)/(number of trades); pace of trading = number of trades per time interval.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

trader was informed (insiders were monopolists). Since monopolists are not forced to trade quickly, more mirages should result: there were more mirages (and less information revelation) compared to a baseline session with several insiders.

Sunder (1991) ran experiments in which subjects bought inside information so the number of insiders was endogenous. In periods with few active insiders, some mirages occurred.

B. Temporary Mirages

In many N periods, there were "temporary" mirages, in which the price path veered toward the G or B price, then corrected itself and ended near the uninformed expected value. We counted nine temporary mirages in 47 periods.[21]

A typical path for temporary mirages is a series of noisy trades away from the expected value, toward either the G or B price, followed by a peak (or bottom) in prices and movement back toward the expected value at the end of the period.

A good example is period 17 of session 7 (see fig. 7). In this period, trader 7, a type I, bought units at escalating prices up to 350, then sold a unit at 325. The drop in price is an important signal: since insiders typically buy at monotonically increasing prices throughout a period (or sell at falling prices), a fall after a rise indicates the rise was uninformative. Furthermore, if others recognized that trader 7 was selling a unit at 325 that he had bought earlier, his sale clearly marked him as a speculator rather than an insider who knew the state was G. That realization caused prices to fall further, almost to the expected value.

VI. Conclusion

There is much ongoing debate over whether stock prices are too volatile to be rational estimates of intrinsic value. Prices are also more volatile when markets are open than when they are closed.

One explanation of excess volatility is that trading is "self-generating": in inferring information from the trades of others, traders sometimes err and their errors cause others to overreact, creating price paths that falsely reveal information that no one has. We call these information mirages.

It is hard to know the distribution of information in natural settings and thus hard to know whether trading is self-generating. Therefore,

21. Temporary mirages occurred in periods 7–9 (session 2), 9 (session 3), 10 and 15 (session 4), and 14 and 17 (session 7). In many of these periods, prices were consistently away from the uninformed expected value, and did not veer back toward it. We consider them temporary mirages because the holdings data suggested that subjects did not end the period believing there was inside information, which is one criterion for a sustained mirage.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

we tested for self-generating trading in experimental markets where we control the flow of information.

In the sessions, subjects trade a 1-period asset that pays a state-dependent dividend. In some periods, some subjects have perfect inside information about the state. In other periods there are no insiders. Uninformed subjects cannot be sure whether others are insiders or not. We can test whether uninformed traders learn the inside information when there are insiders and whether there are mirages in periods without insiders.

In our seven sessions, uninformed traders did learn the information of insiders (as in most earlier experiments).

In 47 periods without insiders, we observed four mirages that were sustained for an entire trading period; subjects traded, mistakenly, as if they had learned inside information from others. (We also counted nine temporary mirages, which started and ended within a trading period.) So mirages did occur, but they were not common.

Our results are related to other theories and evidence. For instance, Copeland and Friedman (1991) hypothesize a process by which traders form price expectations based on market signals. When signals are imperfect (i.e., trading is "noisy"), traders can wrongly infer the true state.

In experimental markets for long-lived assets, Smith, Suchanek, and Williams (1988) observed persistent departures from expected-value pricing ("bubbles"). They argue that while individual traders may know the intrinsic value of assets, if they are not sure that others know, then speculation may appear profitable (cf. Keynes [1936] 1964). Our mirages are caused by uncertainty about the information of others; bubbles seem to be caused by uncertainty about the rationality of others (or perhaps by irrationality).

Others have studied models in which "noise traders" act irrationally and have asked whether rational traders can eliminate the noisy effects (DeLong et al. 1989; Cutler, Poterba, and Summers 1990). In our sessions, traders who spark mirages by trading away from the uninformed (expected value) price act like noise traders.

Sustained mirages always occurred in early trading periods of our sessions. These mirages can be thought of as errors in Bayesian inference of information from prices. In early periods, traders have not yet learned the typical price paths in insider and noninsider periods. Noise trading then generates a price path that resembles the path in a previous insider period, to which other traders overreact.

But mirages did not occur very often in later periods because traders could identify the presence of insiders by the speed at which trading took place. Two facts support this conjecture: sustained mirages only occurred early in the sessions, before traders could gather enough data to distinguish between insider and noninsider periods by the pace of

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

trading; and in the two sessions where the speed of trading was equal in insider and noninsider periods, there were substantially more mirages.

Because traders can use the speed of trading to judge the existence of insiders, sustained mirages are unlikely to occur in experiments with experienced subjects. But temporary mirages often occurred late in the sessions; they could happen with experienced traders in natural markets, too, if noise traders made Bayesian inference difficult.

The importance of nonprice information in limiting mirages suggests one policy implication. Financial economists have sometimes argued (e.g., Manne 1966) that restricting insider trading inhibits the ability of prices to incorporate information quickly. Our results add a new twist to this argument. Since restricting insider trading makes it less frequent or slower, restrictions inhibit the ability of uninformed traders to detect whether there are insiders or not. While reducing volatility when there are insiders, these restrictions may cause more mirages when there are not insiders (as in our sessions 2 and 4), which is certainly undesirable if prices are supposed to reflect information.

References

Ang, James S., and Schwarz, Thomas. 1985. Risk aversion and information structure: An experimental study of price variability in the securities markets. *Journal of Finance* 40 (July): 825–44.

Aoki, Masanao, and Friedman, Daniel. 1986. Asset price bubbles from poorly aggregated information—a parametric example. *Economics Letters* 21:49–52.

Arrow, Kenneth J. 1982. Risk perception in psychology and economics. *Economic Inquiry* 20 (January): 1–9.

Banks, Jeffrey S. 1985. Price-conveyed information versus observed insider behavior: A note on rational expectations convergence. *Journal of Political Economy* 93:807–15.

Black, Fischer. 1986. Noise. *Journal of Finance* 41 (July): 529–43.

Camerer, Colin F. 1987. Do biases in probability judgment matter in markets? Experimental evidence. *American Economic Review* 77 (December): 981–97.

Camerer, Colin F. 1989. Bubbles and fads in asset prices. *Journal of Economic Surveys* 3:3–41.

Camerer, Colin F. 1990. Do markets correct biases in probability judgment? Evidence from market experiments. In L. Green and J. H. Kagel (eds.), *Advances in Behavioral Economics,* vol. 2. Norwood, N.J.: Ablex.

Camerer, Colin F., and Chernew, Michael. 1987. Is trading self-generating? Experimental evidence. Working paper. Philadelphia: University of Pennsylvania, Department of Decision Sciences.

Campbell, John Y., and Shiller, Robert. 1987. Cointegration and tests of present value models. *Journal of Political Economy* 95 (October): 1062–88.

Copeland, Thomas, and Friedman, Daniel. 1987*a*. The effect of sequential information arrival on asset prices: An experimental study. *Journal of Finance* 42 (July): 763–97.

Copeland, Thomas, and Friedman, Daniel. 1987*b*. The market value of information: Some experimental results. Working paper. Los Angeles: UCLA Graduate School of Management, January.

Copeland, Thomas, and Friedman, Daniel. 1991. Partial revelation of information in experimental asset markets. *Journal of Finance* 46 (March): 265–95.

Cornell, Bradford, and Shapiro, Alan C. 1989. The mispricing of U.S. treasury bonds: A case study. *Review of Financial Studies* 2:297–310.

Cutler, David; Poterba, James; and Summers, Lawrence H. 1990. Speculative dynamics and the role of feedback traders. *American Economic Review* 80 (May): 63–68.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

DeJong, Douglas V.; Forsythe, Robert; Lundholm, Russell L.; and Watts, Susan G. 1987. Do prices convey information? Futher experimental evidence. Working paper no. 87–10. Iowa City: University of Iowa, Department of Economics, May.

DeLong, J. Bradford; Shleifer, Andrei; Summers, Lawrence H.; and Waldmann, Robert J. 1989. The size and incidence of the losses from noise trading. *Journal of Finance* 44 (July): 681–96.

Diamond, Douglas W., and Verrecchia, Robert E. 1981. Information aggregation in a noisy rational expectations economy. *Journal of Financial Economics* 9 (September): 221–35.

Dubey, Pradeep; Geanakoplos, John; and Shubik, Martin. 1987. The revelation of information in strategic market games: A critique of rational expectations equilibrium. *Journal of Mathematical Economics* 16:105–37.

Fama, Eugene F. 1981. Stock returns, real activity, inflation and money. *American Economic Review* 71 (September): 545–65.

Forsythe, Robert, and Lundholm, Russell. 1990. Information aggregation in an experimental market. *Econometrica* 58 (March): 309–47.

French, Kenneth R., and Roll, Richard. 1984. Is trading self-generating? Working Paper no. 121. Chicago: University of Chicago Center for Research in Security Prices, February.

French, Kenneth R., and Roll, Richard. 1986. Stock return variances: The arrival of information and the reaction of traders. *Journal of Financial Economics* 17 (September): 5–26.

Friedman, Daniel. 1984. On the efficiency of experimental double auction markets. *American Economic Review* 74:60–72.

Friedman, Daniel; Harrison, Glenn; and Salmon, Jon. 1984. The informational efficiency of experimental markets. *Journal of Political Economy* 92 (June): 349–408.

Grossman, Sanford J. 1976. On the efficiency of competitive stock markets where traders have diverse information. *Journal of Finance* 31:573–85.

Grossman, Sanford J. 1981. An introduction to the theory of rational expectations under asymmetric information. *Review of Economic Studies* 48:541–59.

Grossman, Stanford J. 1989. *The Informational Role of Prices*. Cambridge, Mass.: MIT Press.

Grossman, Sanford J., and Stiglitz, Joseph E. 1980. On the impossibility of informationally efficient markets. *American Economic Review* 70:393–408.

Hellwig, Martin. 1980. On the aggregation of information in capital markets. *Journal of Economic Theory* 22:477–98.

Hertzel, Michael G.; Kendall, Coleman S.; and Kretzmer, Peter E. 1990. The volatility of asset returns during trading and nontrading hours: Some evidence from the foreign exchange markets. *Journal of International Money and Finance* 9:335–43.

Huberman, Gur, and Schwert, G. William. 1985. Information aggregation, inflation, and the pricing of indexed bonds. *Journal of Political Economy* 93:92–114.

Jordan, James S. 1982. A dynamic model of expectations equilibrium. *Journal of Economic Theory* 26 (April): 235–54.

Keynes, John Maynard. [1936] 1964. *The General Theory of Employment, Interest, and Money*. London: Harcourt Brace.

Kobayashi, T. 1977. A convergence theorem on rational expectations equilibrium with price information. Institute for Mathematical Studies of Social Sciences Working Paper no. 79. Stanford, Calif.: Stanford University, January.

Kyle, Albert F. 1985. Continuous auctions with insider trading. *Econometrica* 53 (November): 1315–33.

LeRoy, Stephen F. 1989. Efficient capital markets and martingales. *Journal of Economic Literature* 27 (December): 1583–1621.

LeRoy, Stephen F., and Porter, Richard D. 1981. The present-value relation: Tests based on implied variance bounds. *Econometrica* 49 (May): 555–74.

Manne, Henry G. 1966. *Insider Trading and the Stock Market*. New York: Free Press.

Marsh, Terry A., and Merton, Robert C. 1986. Dividend variability and variance bounds test for rationality of stock market prices. *American Economic Review* 76 (June): 483–98.

O'Brien, John, and Srivastava, Sanjay. In press. Dynamic stock markets with multiple assets: An experimental analysis. *Journal of Finance*.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions

Oldfield, George S., and Rogalski, Richard J. 1980. A theory of common stock returns over trading and non-trading periods. *Journal of Finance* 35 (June): 729–51.

Plott, Charles R., and Smith, Vernon L. 1978. An experimental examination of two exchange institutions. *Review of Economic Studies* 45 (February): 133–53.

Plott, Charles R., and Sunder, Shyam. 1982. Efficiency of experimental security markets with insider information: An application of rational expectations models. *Journal of Political Economy* 90 (August): 663–98.

Plott, Charles R., and Sunder, Shyam. 1988. Rational expectations and aggregation of diverse information in laboratory security markets. *Econometrica* 56 (September): 1085–1118.

Roll, Richard. 1988. R^2. *Journal of Finance* 43 (July): 541–66.

Schwert, G. William. 1981. The adjustment of stock prices to new information about inflation. *Journal of Finance* 36:15–30.

Shiller, Robert. 1981. Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review* 71 (June): 421–36.

Shiller, Robert. 1986. The Marsh-Merton model of managers' smoothing of dividends. *American Economic Review* 76 (June): 499–503.

Smith, Vernon L. 1976. Experimental economics: Induced value theory. *American Economic Review* 66:274–79.

Smith, Vernon L. 1982. Microeconomic systems as an experimental science. *American Economic Review* 72 (December): 923–55.

Smith, Vernon L.; Suchanek, Gerry L.; and Williams, Arlington W. 1988. Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica* 56 (September): 1119–51.

Sunder, Shyam. 1991. Market for information: Experimental evidence. Working paper. Pittsburgh: Carnegie Mellon University, Graduate School of Industrial Adminsitration.

Tversky, Amos, and Kahneman, Daniel. 1982. Judgments of and by representativeness. In P. Slovic, A. Tversky, and D. Kahneman (eds.), *Judgment under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.

von Borries, Alexander, and Friedman, Daniel. 1988. Monopolist insiders in computerized asset markets: A note on some experimental results. Organized Research Activity in Applied Economics Working Paper no. 178. Santa Cruz: University of California.

West, Kenneth. 1988. Bubbles, fads and stock price volatility tests: A partial evaluation. *Journal of Finance* 43 (July): 639–55.

Wilson, Robert. 1985. Incentive efficiency of double auctions. *Econometrica* 53 (September): 1101–15.

This content downloaded on Fri, 4 Jan 2013 16:30:09 PM
All use subject to JSTOR Terms and Conditions