ECONOMETRICA
JOURNAL OF THE ECONOMETRIC SOCIETY

Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets
Author(s): Vernon L. Smith, Gerry L. Suchanek, Arlington W. Williams
Source: *Econometrica*, Vol. 56, No. 5 (Sep., 1988), pp. 1119-1151
Published by: The Econometric Society
Stable URL: http://www.jstor.org/stable/1911361

Econometrica, Vol. 56, No. 5 (September, 1988), 1119–1151

BUBBLES, CRASHES, AND ENDOGENOUS EXPECTATIONS IN EXPERIMENTAL SPOT ASSET MARKETS[1]

BY VERNON L. SMITH, GERRY L. SUCHANEK, AND ARLINGTON W. WILLIAMS

Spot asset trading is studied in an environment in which all investors receive the same dividend from a known probability distribution at the end of each of $T = 15$ (or 30) trading periods. Fourteen of twenty-two experiments exhibit price bubbles followed by crashes relative to intrinsic dividend value. When traders are experienced this reduces, but does not eliminate, the probability of a bubble. The regression of changes in mean price on lagged excess bids (number of bids minus the number of offers in the previous period), $\bar{P}_t - \bar{P}_{t-1} = \alpha + \beta(B_{t-1} - 0_{t-1})$, supports the hypothesis that $-\alpha = E(\tilde{d})$, the one-period expected value of the dividend, and that $\beta > 0$, where excess bids is a surrogate measure of excess demand arising from homegrown capital gains (losses) expectations. Thus, when $(B_{t-1} - 0_{t-1})$ goes to zero we have convergence to rational expectations in the sense of Fama (1970), that arbitrage becomes unprofitable. The observed bubble phenomenon can also be interpreted as a form of temporary myopia (Tirole, 1982) from which agents learn that capital gains expectations are only temporarily sustainable, ultimately inducing common expectations, or "priors" (Tirole, 1982). Four of twenty-six experiments, all using experienced subjects, yield outcomes that appear to the "chart's eye" to converge "early" to rational expectations, although even in these cases we get $\hat{\beta} > 0$, and small price fluctuations of a few cents that invite "scalping."

KEYWORDS: Rational expectations, stock market trading, price bubbles, experimental markets.

1. INTRODUCTION

A LONG STANDING THEORY of common stock valuation holds that a stock's current market value tends to converge to the (risk adjusted) discounted present value of the rationally expected dividend stream. If markets are efficient, then, *in equilibrium*, stock prices should change only when there is new information that changes investors' dividend expectations. We examine this rational expectations model in a laboratory environment in which we can control the dividend distribution, and traders' knowledge of it in a market with a finite trading horizon. From rational expectations theory we hypothesize that although deviations from risk adjusted dividend value might be temporarily sustainable by divergent individual expectations, such deviations cannot persist because of the uncertain profits that can be earned by arbitraging the asset's price against its expected dividend value. Consequently, individual adjustments will occur until any risk differences are compensated, and expectations become common and coincide with dividend value. However, current theory makes no prediction as to how long this will take, and whether, or in what form, this process can be characterized.

Three adjustment dynamics can be distinguished: the process that describes changes in the asset's dividend value, the evolution of agents' price expectations, and the asset's price adjustments. Unless agents' expectations are common

[1]We are grateful for support from the National Science Foundation to Indiana University (A. Williams, PI) and to the University of Arizona (V. Smith, PI), and for the research assistance of Shawn LaMaster. Hard copies of the experimental instructions are available at cost upon request.

and correspond to dividend value, the three dynamics will not coincide. Any differences may be due to a lack of common, not irrational, expectations. An important issue is whether the three dynamics converge as expectations become more homogeneous. Fama's (1970) criterion for market efficiency is that there exist no systematic price patterns that allow arbitrage to yield a positive expected net profit, while in the rational bubble literature agents are assumed to have common priors about the value of the asset. For example in Tirole (1982, p. 1163), "... we investigate the possibility of speculative behavior when traders *have* rational expectations. The general idea is fairly simple: *unless* traders have different priors about the value of a given asset... (the)... market does not give rise to gains from trade. Thus *speculation relies on inconsistent plans and is ruled out by rational expectations* (all italics ours)." But how do traders come to "have" rational expectations: i.e. "consistent plans": i.e. common priors? In an experiment we *cannot* control expectations when the theory provides no explicit implementable model of expectations,[2] but we can control the dividend structure and trader knowledge of it. Consequently, we *can* ask whether common knowledge of a common dividend payout is sufficient to induce common expectations. But there is no a priori basis for *assuming* that *initially* all traders will expect other traders to react in the same way to the same information. Each trader may be uncertain as to the *behavior* of others with the same information. Operationally, then, in testing market efficiency or rational expectations in multiperiod asset environments, the important issue is whether through learning within and across (experimental or natural) markets agents will come to "have" rational common expectations and thus produce a no-arbitrage equilibrium.

The objectives of this study, as it developed, were to answer the following questions: (i) Will economic agents trade an asset whose dividend distribution is common knowledge? (ii) If so, can we characterize (empirically) the price adjustment process, and interpret it in terms of convergence to dividend value? (iii) Will we observe price bubbles and crashes as part of the adjustment process occurring in any or some of the experiments?

Before discussing the background for these questions we state briefly our principal finding: expectations (as measured by forecasts) and price adjustments

[2] Bidding theory articulates an explicit Nash version of rational expectations. Thus, an equilibrium bid function $b_i = \beta(v_i | I)$ relates agent i's message b_i to his environment (the item's value, v_i), and the institution, I. This is an equilibrium bid function if each agent i expects his $N-1$ rivals to also use this behavioral decision rule. In the case of the first price auction institution, and constant relative risk averse agents (with CRRA parameter $1-r_i$), we have $\beta(v_i | I_i) = (N-1)v_i/(N-1+r_i)$, with, say, $r_i \in (0,1]$. In this theory one can "control" expectations in an experimental implementation by letting each individual, i, bid against $N-1$ computerized bidders, and *informing* the subject that each computerized bidder bids a fixed fraction of his value $b_j = \hat{\beta}_j v_j$, where the v_j are drawn from the same distributions as v_i in each auction, and each $\hat{\beta}_j$ is drawn once for all auctions from some distribution on the interval $[(N-1)/N, 1)$. Here we "control" expectations by giving each bidder complete information on the bidding behavior of his rivals, where that behavior is defined by the Nash model of *equilibrium* bidding (Walker, Smith, and Cox, 1986). In the absence of a corresponding micro model of the individual agent in bubble theory, the experimenter cannot know what it means to induce common expectations.

are both adaptive, but the adaptation over time across experiments with increasing trader experience tends to a risk adjusted, rational expectations equilibrium.

In the next section we summarize previous experiments that are related to the asset trading environment described in Section 3. The design parameters and the interplay between market performance and the sequence of experiments are discussed in Section 4.

2. PREVIOUS EXPERIMENTS

Several double auction market studies have been characterized by some form of asset trading over time (Miller, Plott, and Smith, 1977; Williams, 1979; Plott and Agha, 1982; Williams and Smith, 1984). In a typical experiment, a constant stationary supply is induced on five agent sellers ("producers") and a two-period cyclically stationary demand is induced on five agent buyers. A third group of agents, who are asset "traders," have the exclusive right to buy in one period and sell in the next. Thus the environment is represented by cyclically stationary flow supply and demand conditions, with agent traders empowered to make asset carryover decisions.

With the exception of Williams and Smith (1984), in all of these experimental markets a pure replication of the environment is imposed by the experimental design. For example in Miller, Plott, and Smith (1977) and in Williams (1979), in addition to demand repeating a two-period cycle, traders can only buy in the low-price period and sell in the high-price period, and are required to close out their inventory positions by the end of each two-period cycle. In Williams and Smith (1984) traders can carry units across market cycles and the rate of convergence is retarded. All of these experimental studies report a significant treatment effect from the speculative action of traders: i.e. in the final market period, contracts tend to be nearer the intertemporal competitive equilibrium price than to either of the cyclical autarky theoretical equilibrium prices, or to the observed contracts in paired comparison cyclical autarky experiments. These results, and all of the many experimental studies of double auction markets without asset trading (see the summary by Smith, 1982) can be interpreted as supporting rational expectations theory as originally defined by Muth (1961, p. 316).

Experiments in which the item traded is an asset proper, in the sense that the environment generates dividends for asset holders at the end of each trading period, were originated by Forsythe, Palfrey, and Plott (1982) and continued in Plott and Sunder (1982), and Friedman, Harrison, and Salmon (1984). Although these important contributions shifted the experimental environment to that of pure asset trading, they maintained two characteristics of the earlier cited "speculation" studies: (i) A two-period A-B cycle (three-period A-B-C cycle in the Friedman, Harrison, and Salmon, 1984, experiments) in private (dividend) values is induced on the item traded, which is repeated over a trading horizon of several cycles. These dividend values differ for different groups of agents creating the same type of (induced value) gains from exchange as in the earlier experiments. (ii) The inventories (shares and money) of traders are *reinitialized* at the beginning of each cycle as a means of achieving a pure replication of the cyclical environment. Within this framework, these asset market experiments are interpreted as yielding prices tending to converge over time toward levels consistent with the rational expectations hypothesis. This is because agents bid for assets initially in period A on the basis of their private information, but slowly learn, across replicating cycles, to adjust their contracting to account for additional information concerning the period B *market* value of the asset. In these experiments agents are observed to engage in very little trading for capital gains in spite of the repetitive pattern of price increases. This may be a consequence of the short capital gains horizon.

Our immediate objective in the present series of experiments was to determine whether agents would actively trade an asset when all investors faced identical uncertain dividend payout schedules. The previous cited asset experiments pay different dividends to different investors on the grounds that investors have different opportunity costs. But if this is so, subject agents ought to have their own homegrown differences in opportunity cost (as in field environments). Consequently it is an open question whether artificially inducing different dividend values on subject investors is a *necessary* condition for observing trade. If our agents are not observed to trade this supports the strong version of the theory in which risk neutral agents have common initial expectations (induced, presumably, by contemplating the implications of a common dividend distribution). Our second objective, given that agents are observed to trade in this environment, is to characterize the observed price adjustments. Do we observe convergence to the rational expectations equilibrium as in previous asset market experiments? Do subjects' forecasts of the mean price (collected in nine experiments), taken as a measure of their price expectations at the time of interrogation, reveal adaptive or rational expectations?

Because of our concern that there might be insufficient divergence in subjective expected values to observe trading, and/or that the finiteness of our market horizon might frustrate any possibility of observing bubbles, we introduced a random valued buyout condition in the first series of experiments in an effort to enhance the possibility of a bubble. As it happens, these ex ante concerns were not supported. Bubbles (relative to the dividend value of the asset) are observed in most of our experiments with inexperienced and to a lesser extent experienced subjects. Moreover, eliminating the random buyout does not eliminate bubbles.

3. THE ASSET MARKET MECHANISM

The trading procedure employed in this study is an enhanced version of the PLATO computerized double-auction mechanism described by Williams and Smith (1984) for commodity markets with intertemporal speculators. The basic trading mechanics for asset-market speculators are identical to those for the commodity-market speculators in the Williams and Smith study. Figure 1 pro-

WEEK 1 — TRADING PERIOD (columns)

RECORD SHEET for TRADER 3	1	2	3	4	5
Unit 1 selling price	3.ØØ		2.4Ø		2.5Ø
Unit 1 purchase price	Ø.ØØ		Ø.ØØ		Ø.ØØ
Profit	3.ØØ		2.4Ø		2.5Ø
Unit 2 selling price					2.5Ø
Unit 2 purchase price					Ø.ØØ
Profit					2.5Ø
Unit 3 selling price					2.5Ø
Unit 3 purchase price					2.ØØ
Profit					Ø.5Ø
Unit 4 selling price					
Unit 4 purchase price					
Profit					
Unit 5 selling price					
Unit 5 purchase price					
Profit					
Dividend Earnings	Ø.8Ø	2.ØØ	Ø.48	Ø.48	
Total Earnings for Period	3.8Ø	2.ØØ	2.88	Ø.48	

Period purchased	1	3	3
Purchase price	2.45	2.3Ø	2.35

Inventory= 3, Working capital=$1Ø.84, Dividend per unit=$????

BUYER 2 BIDS $2.48 SELLER 8 OFFERS $2.6Ø

ENTER OFFER

>$ SELLING -Data→switch

ACCEPT BID

CONFIRM CONTRACT

Last 9 contracts: 2.5Ø,2.5Ø,2.5Ø,2.55,2.6Ø,2.61,2.4Ø,2.65,2.53

Trading Period 5 now in progress. SECONDS REMAINING: 1Ø8

Ø of 12 people have voted to end period 5: -LAB→ vote to end

FIGURE 1.—Screen display for asset market trader.

vides a participant's screen display for our asset market. All agents (referred to as traders) are able to switch between buying mode and selling mode by pressing a key labeled DATA. Traders are free to enter a price quote to buy (or sell) one asset unit by typing their entry and then touching the rectangular area on their screen display labeled "ENTER BID" (or "ENTER OFFER"). Traders are likewise free to accept any other trader's bid to buy (or offer to sell) by touching a screen area labeled "ACCEPT BID" (or "ACCEPT OFFER"). The acceptor must then touch an area labeled "CONFIRM CONTRACT" at which time a binding contract is formed and the exchange information is recorded in the buyer's and seller's private record sheets.

Price quotes must progress so as to reduce the bid-ask spread. Only the highest bid to buy and the lowest offer to sell are displayed to the entire market and are open to acceptance. Price quotes that violate this rule are placed in a "rank queue"; after a contract occurs the rank queue automatically enters the best (highest) queued bid and best (lowest) queued offer as the new bid-ask spread. Smith and Williams (1983) have shown that this version of the double auction tends to outperform three alternative versions in terms of allocative efficiency and the speed of convergence to a competitive equilibrium.

Trading occurs over a sequence of 15 (or 30) market periods, each lasting a maximum of 240 seconds. Market participants can bypass this stopping rule by unanimously voting to end a period. Registering a vote to end a period does not affect a trader's ability to participate actively in the market. The number of seconds remaining and the current vote to end the period are presented as shown at the bottom of the Appendix display. Screen displays are updated approximately every second.

At the beginning of the experiment, each trader is given an asset endowment and a cash endowment. A trader's cash holding (referred to as "working capital") at any point will differ from his/her cash endowment by: (i) accumulated capital gains (or losses) via market trading, and (ii) accumulated dividend earnings via asset units held in inventory at the end of each trading period. At the experiment's conclusion, participants are paid in cash the amount of their final working capital. It is worth re-emphasizing that traders' asset and cash holdings are endogenous to the experiment beyond the beginning of trading period 1. We do not "reinitialize" the market at any time as has been done in the cited studies of experimental asset markets, with the exception of Williams and Smith (1984).

Traders are informed in the instructions of the probabilistic nature of the dividend structure that they will encounter and the total number of trading periods in the experiment. Specifically, they know all the possible (per-unit) dividend values that might be drawn (i.i.d.) and the probability associated with each potential dividend value. They do not, however, know the actual dividend that will be awarded at the end of any trading period until that period's conclusion, at which time they are informed of their dividend earnings for that period. Prior to *each* period, traders are reminded of the dividend distribution, and informed of the "average," minimum, and maximum possible dividend earnings for each unit held in their inventory for the remainder of the experiment. All participants are verbally informed that the dividend structure and actual dividend draws are the same for everyone in the market. At the end of each period, market participants are also given access to a table displaying the average, maximum, and minimum contract price, as well as the dividend awarded in all previous periods.

When a trader buys an asset unit the price and the period purchased are recorded in the trader's inventory table (see the Appendix). (Endowed asset units are recorded as being purchased in period 0 at a price of 0.) Traders can continue to buy asset units as long as their working capital is sufficient to cover the purchase price. There is also a (rarely binding) maximum inventory size of seven units due to the horizontal space limitations of the display screen. Traders can sell off inventory units at any time. However, short sales are not permitted. For record-keeping purposes, inventories are automatically maintained on a *first in first out* basis. When a unit is sold, the sale price, purchase price, and resulting capital gain or loss are recorded in the trader's record sheet.

In some of the experiments reported below, all asset units were automatically purchased by the experimenters at the market's conclusion. (The default mode is to award only the period 15 dividend.) The "buy-out price" equals the sum of the dividend draws over all 15 periods plus or minus a constant (each with .5 probability). The buy-out option keeps the expected value of an asset unit from falling to the expected value of a single dividend draw during the final trading period. When the buy-out option is utilized, the information presented to subjects regarding the maximum, expected, and minimum dividend earnings associated with holding an asset unit for the remainder of the market is automatically adjusted to account for the buy-out.

4. OVERVIEW OF DESIGN PARAMETERS, MARKET PERFORMANCE, AND THE SEQUENCE OF EXPERIMENTS

We report the findings from 27 experiments using the design parameters listed in Table I. In every experiment there were three endowment classes of agents (see columns 2–4), each consisting of three (four) subjects in the nine (twelve) trader experiments. This design permits pure expansions in the size of a market to be effected without altering its per capita structure. Designs 1 and 3 were used mostly for inexperienced subjects. All subjects had participated in a previous double auction experiment with induced flow supply and demand conditions. Consequently, all experiments with an "x" suffix in Table I and charted in the figures used subjects who had been in at least two previous experiments.

Many of the experiments we report were directly motivated by questions and puzzles posed by the results of earlier experiments. The research program developed as a continuing dialogue between hypothesis and empirical results in an effort to increase our understanding of the trading patterns that emerged in these markets. This historical theme allows the reader to appreciate the experiments we scheduled at each stage in the development of our tentative conclusions. The reader is cautioned that this narrative, and the associated price charts, may give the initial impression that rationality is "grossly" violated. The empirical analysis in Section 5, however, reveals that the predominating characteristic of these experiments is the tendency for expectations and price adjustments to converge to intrinsic value across experiments with increasing subject experience.

Our first pilot experiments (not reported) used subjects with no previous double auction experience of any kind, and the expected dividend or holding value of a share was computed and reported to the subjects only for the first period. Because prices in these experiments deviated by a wide margin from $E(\tilde{D}_t^T)$, we decided to increase the experience level and information state of our subjects to eliminate the possibility that our results were sensitive to these factors.

TABLE I

Design	Endowment[a] Class I	Endowment[a] Class II	Endowment[a] Class III	Dividend $\tilde{d}$, cents ($p = 1/4$)[b]	Expected Dividend per Period, $E(\tilde{d})$, cents	Intrinsic (Dividend) Value per Share Period 1, $E(\tilde{D}_1^T)$[c]	Experiments[d]
1	(\$2.80; 4)	(\$7.60; 2)	(\$10.00; 1)	(0, 4, 8, 20)	8	\$2.40 (Including Buyout)	(5; 12) (7; 12) (12xn; 9, 3c) (17; 12) (23pc; 12)
2	(\$2.25; 3)	(\$5.85; 2)	(\$9.45; 1)	(0, 4, 14, 30)	12	\$3.60 (Including Buyout)	(6x; 9) (9x; 9) (10; 9) (16; 9) (18; 9) (19x; 9) (20xpc; 9)
3	(\$2.80; 4)	(\$7.60; 2)	(\$10.00; 1)	(0, 8, 16, 40)	16	\$2.40	(26; 12) (41f; 12)
4	(\$2.25; 3)	(\$5.85; 2)	(\$9.45; 1)	(0, 8, 28, 60)	24	\$3.60	(25x; 9) (28x; 9) (30xsf; 9) (36xx; 9) (39xsf; 9) (43xnf; 9) (46f; 9) (48xnf; 9) (49xnf; 9) (50xxf; 9) (90f; 9) (124xxf; 9)
5[e]	(\$2.25; 3)	(\$5.85; 2)	(\$9.45; 1)	(0, 8, 28, 60)	24	\$7.20	(42xf; 9)

[a] In experiments with 9(12) traders, 3(4) traders are assigned to each class.

[b] Each dividend outcome occurs with probability 1/4 in each period.

[c] Each period's expected dividend value, $E(\tilde{D}_t^T)$, $t = 1, 2, \ldots, T$, is computed and displayed to each trader before the beginning of the period. In designs 3–5 (no buyout), $E(\tilde{D}_t^T) = E(\tilde{d})(T - t + 1)$, $t = 1, 2, \ldots, T$. In designs 1–2 (with buyout), $E(\tilde{D}_t^T) = \Sigma_{\tau=1}^{t} \hat{d}_\tau + 2E(\tilde{d})(T - t + 1)$, since the buyout at T is $\Sigma_{t=1}^{T} \hat{d}_t \pm 0.50$, probability 1/2, in design 1 ($\Sigma \tilde{d}_t \pm 1.00$, probability 1/2 in design 2), $\hat{d}_\tau$ refers to the realized dividend at the end of τ.

[d] (5; 12) means experiment number 5 using 12 subjects. x means experienced. xx means superexperienced. s means subjects were trained in a sequence of independent single period asset markets. n means some novice (inexperienced) subjects were combined with experienced subjects. f means subjects were asked to forecast next period's mean price. pc means price controls were set at $E(\tilde{D}_t^T) \pm 0.10$ for $t = 1, 2, 3$. In experiment (12xn; 9, 3c), 3 of the 12 traders were confederates.

[e] $T = 30$ in experiment (42xf; 9): otherwise $T = 15$.

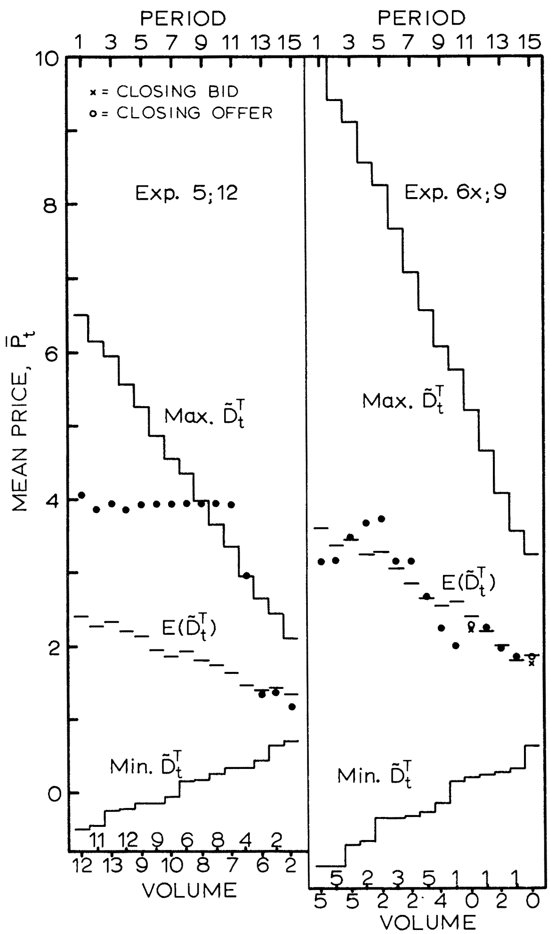

FIGURE 2

Our first two experiments, 5 and 6x, are charted in Figure 2.[3] The nine traders in 6x were a subset of the twelve who participated in 5. Because of the large volume of transactions we chart only the mean price by period. Experience appeared to be an important determinant of trading patterns. Both inexperienced and experienced subjects had a sense of the asset's intrinsic worth. For example, subjects in the inexperienced group asked the experimenters "why the buying panic?" and "shouldn't it sell near dividend value?" However, the subject who perceived a "buying panic" accumulated net inventory through period 11 and suffered a capital loss in period 12! Indeed, if you expected prices to hold steady

[3]Under the influence of RE theory and the strong previous experimental evidence favoring it in replicated environments, we hypothesized initially that allowing asset holdings to float (without reinitialization) might not be sufficient to yield observations that deviated much from the intrinsic value (dividend) rational expectations hypothesis, $E(\tilde{D}_t^T)$, over the ostensibly "short" horizon of $T = 15$ trading periods.

for many periods, it was rational to buy (or hold), collect the dividends, and plan to sell later at the inflated price. The top earning subject approximated this strategy. Experiments 6x, and 9x using experienced subjects from experiment 7, appeared to confirm our conjecture that with experience and full (calculated) information, prices would converge to the intrinsic value, $E(\tilde{D}_t^T)$, although 9x (Figure 3) suggests that behavior consistent with risk aversion may be observed in the first several periods.

In each of the first four experiments (5, 6x, 7, 9x) the mean price in period 2 was "close" to the mean in period 1. Therefore, we conjectured that expectations might be sensitive to the initial contracts, and that if we could induce initial trading at prices near $E(\tilde{D}_t^T)$, then the market might follow this path. We tested this conjecture by recruiting seven experienced subjects and two inexperienced subjects at Indiana University for experiment 12xn. Three experimenters at the University of Arizona participated as confederate "insiders" in a 12 trader market (most experiments were conducted multisite). Our plan was for the insiders to trade so as to maintain the price in a range within 10 cents of $E(\tilde{D}_t^T)$ for two trading periods, using period 3 to adjust the total insider share inventory to the level of the initial endowment, and then become inactive. Since both previous experienced trader markets had opened below $E(\tilde{D}_t^T)$, we guessed that insider activity in experiment 12xn would have to concentrate on buying support. Hence, the strategy was for the two insiders with the largest cash endowments (7.60 and 10.00) each to enter opening bids at 2.30, and for the trader with the largest share endowment (4 units) to enter an opening offer at 2.50. If the standing bid was accepted, it was backed by a new bid at the same price. If the standing offer was accepted the strategy was to immediately replace it with a new offer at 2.50. The insiders encountered unanticipated buying strength, and they were able to contain the surge in demand only by allowing some contract prices in excess of $E(\tilde{D}_t^T) + 0.10$ during the first three periods. As shown in Figure 3, this effort partially succeeded in that prices did not rise by very much in periods 4 and 5 before converging to near $E(\tilde{D}_t^T)$. Based on experiment 12xn we tentatively rejected the hypothesis that these markets are robustly sensitive to the "accident" of where they start. Strong endogenous expectations and behavioral uncertainty appear to determine the starting level as well as the subsequent course of prices, and these expectations are not easily neutralized even when 25% of the market is controlled by a confederate attempt to impose $E(\tilde{D}_t^T)$ expectations.

Since our first markets with experienced traders were yielding less than complete convergence to $E(\tilde{D}_t^T)$ share values, we continued to run paired experiments consisting of a 12 trader asset market followed by a market using a 9 trader subset of the first group (Figures 4–7). In two of these experiments (20x in Figure 5, and 23 in Figure 6) we imposed a computer enforced price ceiling at $E(\tilde{D}_t^T) + 0.10$ and a price floor at $E(\tilde{D}_t^T) - 0.10$ for the first two trading periods. These price controls would have the effect of forcing the market to trade within 10 cents of $E(\tilde{D}_t^T)$, which was the objective in experiment 12xn, but with the potentially important difference that it would be common knowledge that prices in this range would be the result of an externally imposed constraint. In

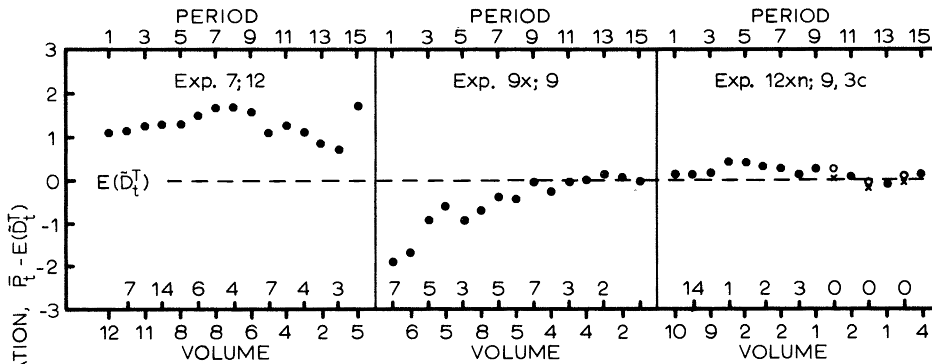

FIGURE 3

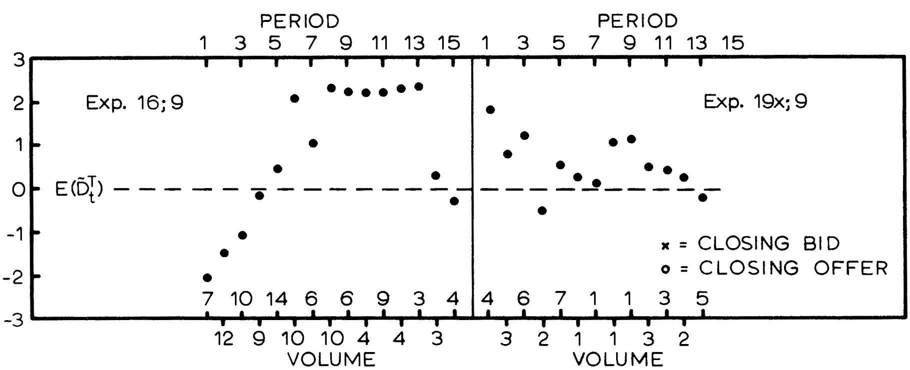

FIGURE 4

experiment 12xn we had sought surreptitiously to create the belief that such prices had occurred "naturally."

In experiment 16 (Figure 4) we observed our first full scale market bubble—a boom followed by a market crash. Replication of this experiment (19x) with experienced subjects failed to extinguish a boom-bust pattern of trading. In experiment 17 we observe a relatively smooth bell-shaped pattern of mean prices over time. A subset of these subjects returned to participate in experiment 20x which imposed a price ceiling and floor in periods 1 and 2 designed to see if such a constraint could induce $E(\tilde{D}_t^T)$ price expectations. This treatment worked as predicted, showing that an intrinsic value rational expectations price pattern could be approximated by combining experience (even bubble experience) with two initial periods of trading at controlled prices near $E(\tilde{D}_t^T)$. Would the same result be produced with inexperienced traders, and, if so, would it carry over into a subsequent market (without price controls) using a subset of these "conditioned" traders? From the chart of experiments 23 and 25x in Figure 6, we see that the answer is emphatically no. In experiment 23 the market traded near the ceiling price for the first two periods. Upon the removal of the price controls the market price increased by about one-third, with increased volume, then held

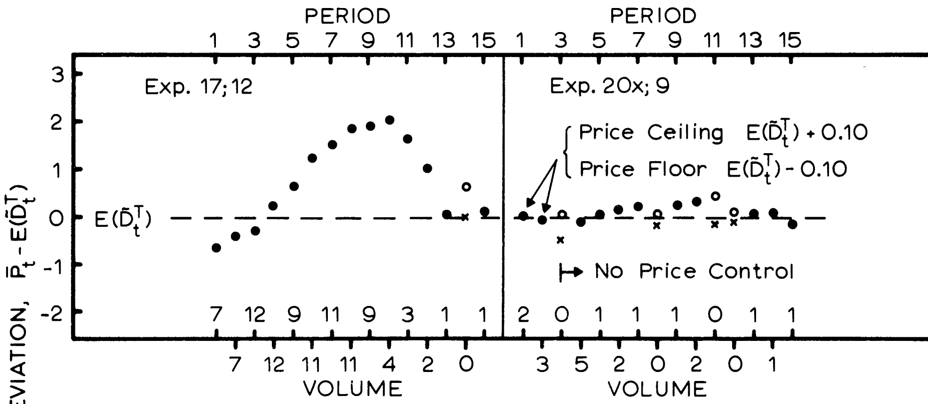

FIGURE 5

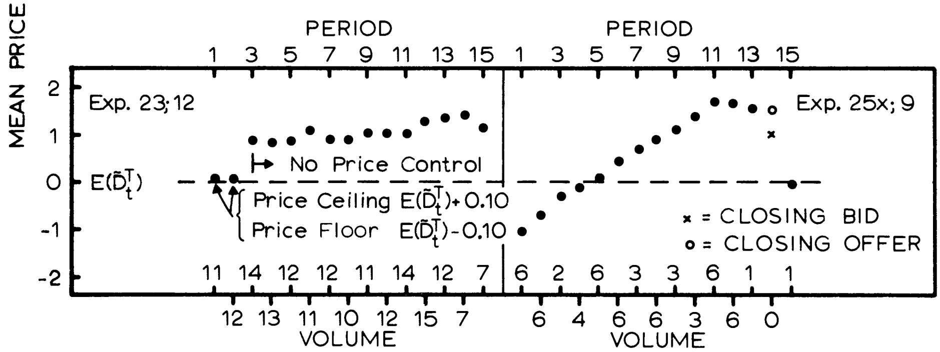

FIGURE 6

approximately steady until the last period. But the nine member subset of this group who participated in experiment 25x produced a substantial bull market measured relative to $E(\tilde{D}_t^T)$. This demonstrates the potential for endogenous expectations to dominate the objective underlying parameters of a market. Experiment 25x marked the beginning of our series of experiments with no end-of-horizon buyout at $T = 15$, and demonstrated that our use of such a buyout to enhance the volatility of expectations was unnecessary.

The two experiments in Figure 7 provide back-to-back market bubbles in which the first experiment (26) appears to have produced an expectation of a bubble in the replication (28x) causing the second bubble to rise faster, and break sooner than the first. (As subjects were arriving for 28x, one commented to an experimenter that he expected this market to "crash," which of course implied that he also expected it to first "boom.") This appears to be an excellent example of self-fulfilling expectations.

Experiment 10 (Figure 8) is noteworthy because of its use of professional and business people from the Tucson community, as subjects. This market belies any notion that our results are an artifact of student subjects, and that businessmen who "run the real world" would quickly learn to have rational expectations. This

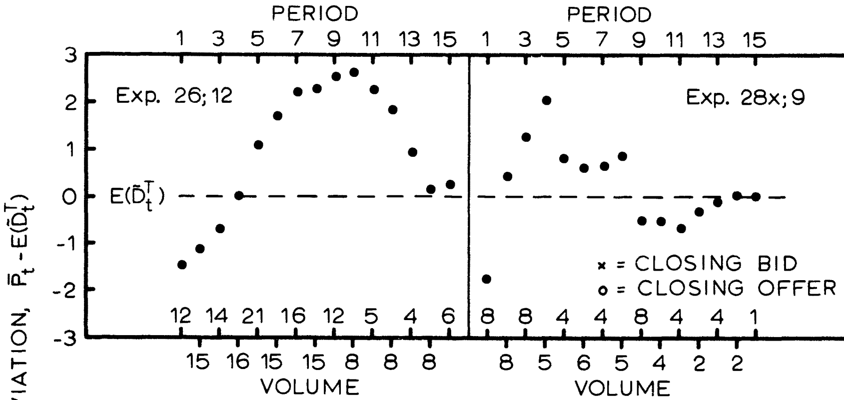

FIGURE 7

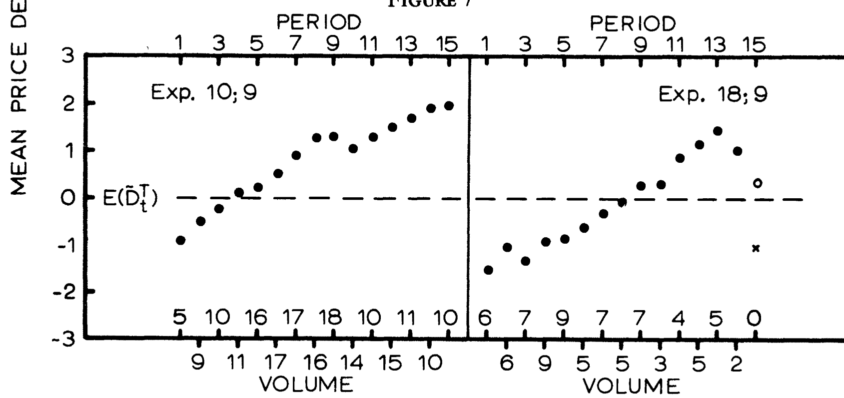

FIGURE 8

is the only experiment we conducted that closed on a mean price higher than in all previous trading periods. Of interest is the fact that because this experiment was conducted in the evening hours (9–11 pm CST), it had to be interrupted shortly (about 10 minutes) for the regular 10 pm CST PLATO shutdown for servicing. We informed subjects that they would be logged back in for period 10 at the same asset position each had at the end of period 9. In spite of our assurances that things would proceed as if there had been no interruption, the market steadied in anticipation of the interruption at the end of period 9, sold off in period 10 after restarting, then recovered to resume the steady-growth trend of periods 1–8. This result illustrates the sensitivity of an asset market to external sources of subjective uncertainty even when the experimenter uses instruction to attempt to neutralize their possible significance; it also corroborates the widely held belief that stock markets are vulnerable to "psychological" elements (factors other than "fundamentals" that create common expectations).

An empirical regularity in those markets that experience a price bubble is for the collapse in market prices to occur on a trading volume that is smaller than the average volume in the periods preceding the collapse. This is illustrated in

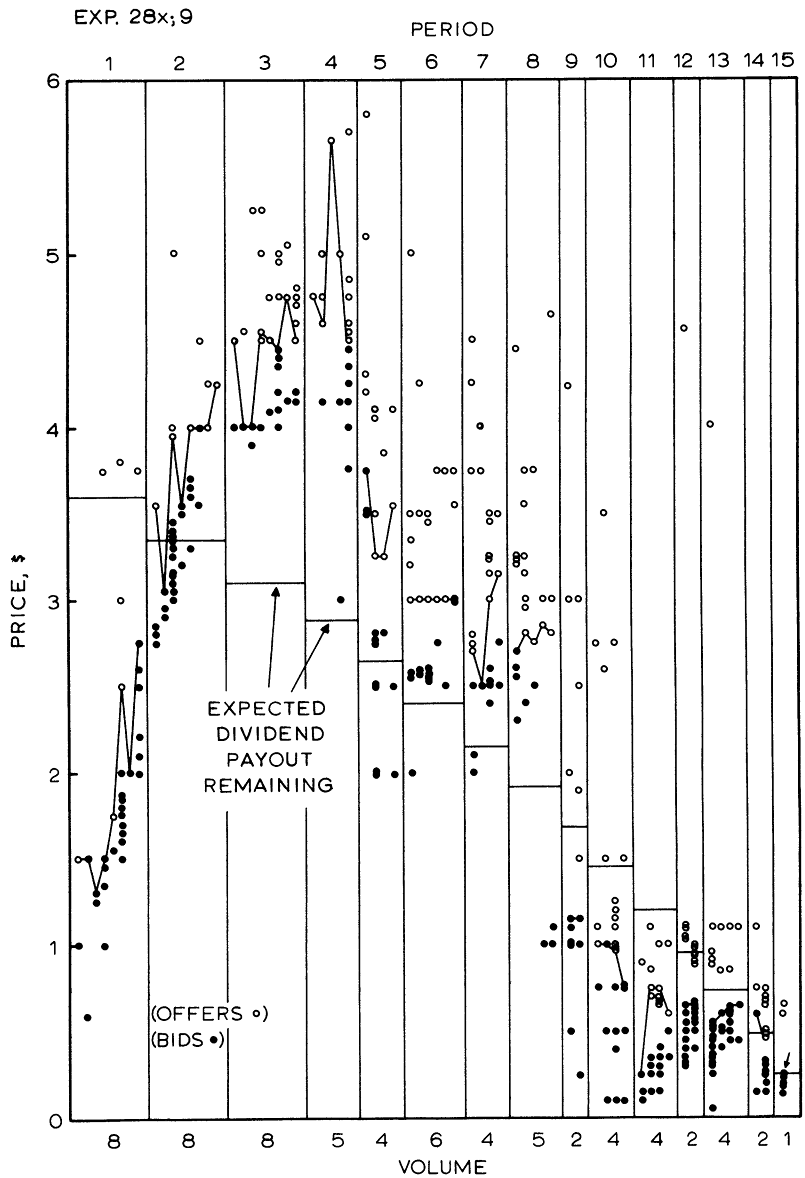

FIGURE 9

experiments 16 (periods 7 and 14), 17 (periods 11–15), 26 (periods 11–15), 28x (period 5), and 18 (periods 14–15). Even more telling is the tendency for volume to shrink in the period just *prior* to the collapse in prices.

Figure 9 provides a chart of all bids and offers and the resulting contract prices (joined by line segments) in sequence for experiment 28x, and illustrates the dynamics of price behavior both within and across trading periods in one market bubble. This market rose from an intraperiod low of \$1.30 in period 1 to an intraperiod high of \$5.65 in period 4. The significance of the changing pattern of bid-offer activity will be discussed in Section 5.2.

At this juncture in our research we posed the following questions. Are our results influenced artifactually by the confounding condition that when subjects participate in their first asset market experiment they simultaneously acquire training in the mechanics of asset trading *and* form expectations about the price behavior of such markets over time? In particular could it be that the price bubbles and market crashes with first-time asset traders are due to their inexperience, with similar bubble and crash phenomena repeated with second-time traders because of expectations created in the first market? To resolve these questions, in experiments 30xsf and 39xsf (Figure 10) subjects were experienced, but they did *not* acquire experience from a previous 15 period asset market. Their experience was obtained by participating in a sequence of single-period asset trading markets in which each trader's endowment was reinitialized at the beginning of each period, and no inventories of shares purchased in any earlier period could be carried over to any later period. Thus subjects were trained in an asset trading market in which no capital gains (or losses) were possible across trading periods. This treatment allowed subject experience in trading mechanics to be acquired while controlling for bubbles and crashes. The results charted in Figure 10 for experiments 30xsf and 39xsf show that bubbles and crashes can indeed occur in markets with experienced subjects who have not been inadvertently conditioned to expect bubbles and crashes in the process of acquiring experience. Although the trading patterns are quite different in 30xsf and 39xsf, each corresponds to one of the two major performance patterns identified in the earlier expermients.

Experiment 36xx was designed to see if the "superstar" traders in our previous experiments would yield intrinsic value market prices. These nine subjects had all participated in at least two previous asset markets (in addition to the basic supply and demand trainer). Also they had been screened for profit performance, so that eight of the subjects had been among the top earning subjects in all previous experiments. (One subject who was an exception to this screening rule earned the third highest profits in 36xx.) As indicated by the chart in Figure 10, experiment 36xx yielded a substantial (but very low volume) price bubble.

The single period horizon experiments ($T = 1$) used to train subjects for experiments 30xsf and 39xsf produced no intraperiod price bubbles. Yet we often observe bubbles when $T = 15$. (One of us has observed bubbles when $T = 3$ using inexperienced subjects.) This suggests the extra theoretical hypothesis that bubble effects should be intensified if we double the horizon from 15 to 30 periods, since this would increase the scope for capital gains expectations to swamp intrinsic value. In experiment 42xf, we set $T = 30$ using a nine member subset of the subjects in 41f (Figure 11). Contrary to this view, experiment 42xf (charted in Figure 11) appears to converge quickly to intrinsic dividend asset value (but see Section 5 for qualifications) in spite of the trading group's experience with a sharp price bubble and collapse in 41f. We interpret this result as strengthening the interpretation that these markets are sensitive to group endogenous expecta-

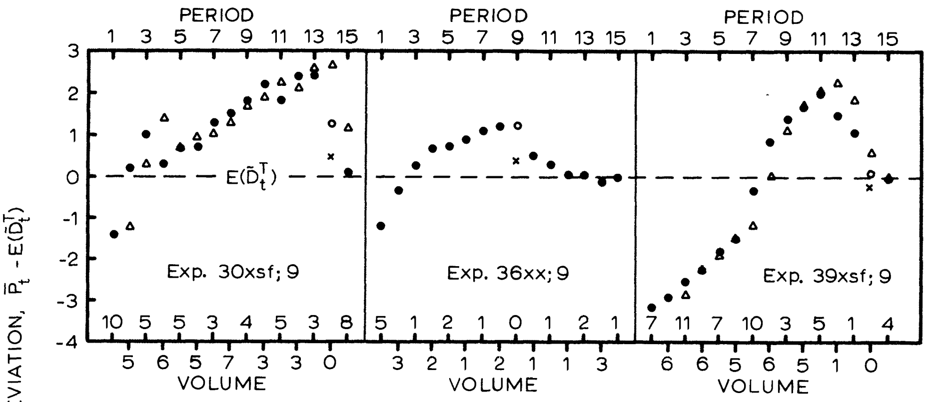

FIGURE 10

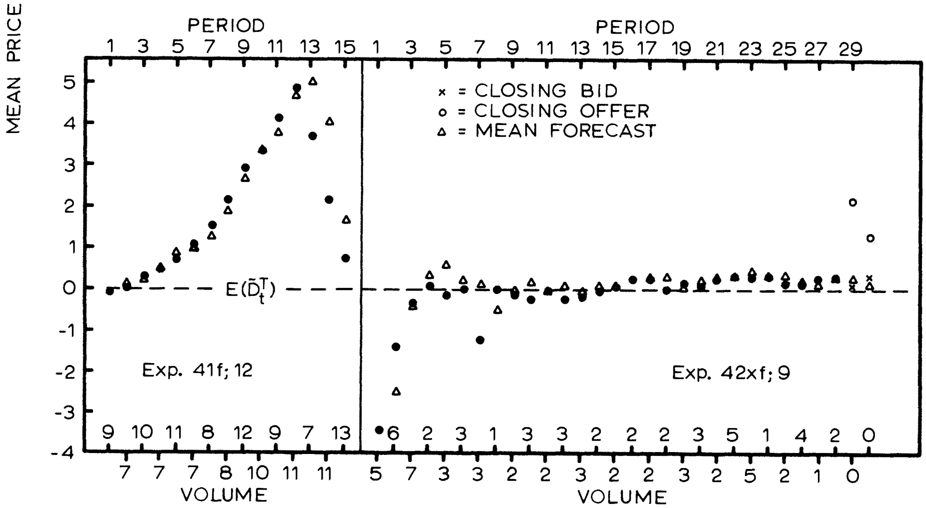

FIGURE 11

tional factors that are not reliably manipulated by such controllable treatment variables as experience, information, and horizon length.

Experiments 43xnf, 48xnf, and 49xnf represent an effort to mix experienced traders who had yielded intrinsic value equilibrium prices with traders who either had no experience or who had a bubble experience. This treatment was motivated by the conjecture that if the stock market was dominated only by professional traders, one might observe intrinsic value asset prices, but that the presence of uninformed novices who lose money, leave the market, and are replaced by new novices, prevents such equilibria from occurring. In 43xnf, six subjects were recruited from the "professionals" who had participated in experiment 42xf. The remaining three subjects in 43xnf were novices in the sense that one had no previous asset trading experience and two had experience only in a bubble

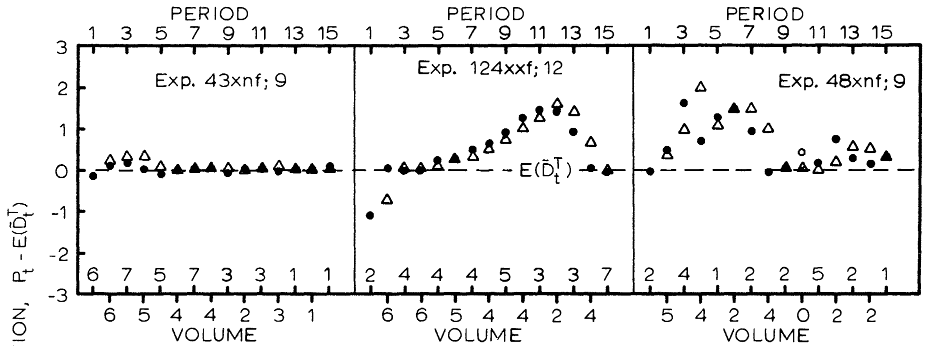

FIGURE 12

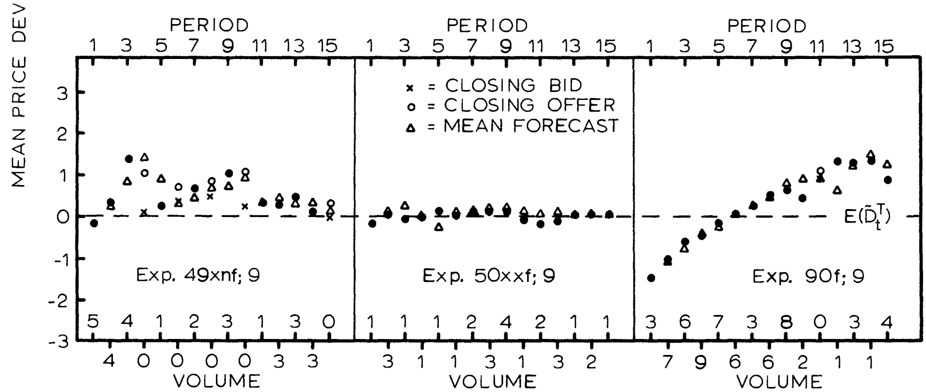

FIGURE 13

market (experiment 41f). From the chart of mean prices for 43xnf in Figure 12 it is seen that with two-thirds of the market made by "professionals," we observe convergence. But experiment 48xnf, with only three "professionals," four with bubble experience, and two inexperienced subjects, provided an erratic bubble for the first six periods. These results do not contradict a conjectured "professional effect" such that, if there are enough such experienced traders, they will dampen any bubble tendency. This conjecture is further supported by experiments 49xnf and 50xxf. In 49xnf we had a mixture of four "professionals" and five subjects with bubble, or no experience; these subjects created a small bubble (Figure 13). We then recruited these same subjects for a replication (one was replaced with a highly experienced subject), and the new market, experiment 50xxf, traded near intrinsic value. It appears that *replication with essentially the same subjects* eventually will create a "professional" market with common expectations in which bubble tendencies are extinguished and replaced by intrinsic value pricing. In this regard experiment 124xxf is of special significance in that the subjects were sophisticated graduate students with experience in at least two previous asset markets, but it was their first 15 period time horizon. Their previous experience was with the Forsythe et al. (1982) and Friedman et al. (1984) environments in which they converged to the rational expectations equilibrium. As shown in Figure 12, this market converged temporarily on dividend value then exhibited a bubble before crashing back to dividend value. These subjects clearly understood the dividend structure, but "played" the bubble.

Beginning with experiment 30xsf, at the end of each trading period each subject was asked to forecast the mean contract price in the next trading period. The forecasting exercise followed the procedures utilized in Williams (1987). The subject with the smallest cumulative absolute forecasting error over periods 2 to 15 earned an additional $1. Individual forecasts were private information, and to avoid providing an incentive to manipulate price in a "close" forecasting race no subject was informed how their own forecasts compared with those of others. While entering forecasts, subjects' screens displayed the entire history of their own forecasts, mean price, and absolute forecast errors. Williams (1987) provided evidence that $1 is sufficient incentive for serious forecasting, but not so large as to motivate strategic manipulation of the mean price in an effort to win the forecasting prize.

In Figures 10–13 are plotted the mean of these individual forecasts on the same scale with the mean contract prices realized in each period. These charts reveal several characteristics of the mean forecasts: (1) in many periods the mean forecast appears not to be a bad predictor of the mean price; (2) forecasts tend to be good when the mean price is approximately constant (as in 90f, periods 13–14), exhibits a small trend (as in 41f, periods 2–6), or follows intrinsic value (as in 42xf): (3) the forecasts lag behind larger changes or trends in price (as in 30xsf, periods 2–4, and 49xnf periods 2–10); (4) the forecasts invariably fail to predict turning points (as in 30xsf, period 4, 39xsf, period 12, and 41f, period 13). In short our subject's forecasting ability in these markets is similar to that of professional forecasters in the field.[4] Characteristics (3) and (4) are particularly interesting since experimental market prices, including price jumps and turning points, are determined entirely by the endogenous actions of the same individuals who are making the forecasts!

5. PRICE FORECASTS AND PRICE DYNAMICS: HYPOTHESES AND EMPIRICAL RESULTS

In formulating some hypotheses implied by alternative models of forecasting behavior and price adjustments, we will distinguish between *rational expectations in the sense of Muth* (REM) and *rational expectations in the sense of Nash* (REN). In Muth's (1961) well known treatment the REM hypothesis is "... that expectations of firms (or, more generally, the subjective probability distribution of outcomes) tend to be distributed, for the same information set, about the

[4]Such characteristics of professional forecasters have a long history. For example, "... forecasters tend to rely heavily on the persistence of trends in spending, output, and the price level. To the extent that inertia prevails in the economy's movement, their predictions turn out to be roughly right... but... such forecasts suffer from missing business cycle turns and underestimating recessions and recoveries..." (Zarnowitz, 1986, pp. 17–18).

predictions of the theory (or the 'objective' probability distribution of outcomes)." However, it is perhaps less well known that many years earlier Nash (1950, p. 158) defined the concept less restrictively by stating that "... since our solution should consist of *rational* expectations of gain..., these expectations should be realizable." Thus REN implies only that expectations are sustained (or reinforced) by outcomes, while REM implies that expectations are sustained by outcomes that in turn support the predictions of some theory.[5]

5.1. *Are Subject Price Forecasts Accurate, Valuable, Adaptive?*

We begin our analysis with the experiments (designated with an "f" under the listing "design 4 experiments" in Table I and displayed in Figures 10–13) in which traders submitted forecasts of the mean price in the next trading period. The question of forecast accuracy is examined using OLS estimation of the equation

$$\bar{P}_{t,e} = \alpha_1 + \beta_1 F_{t,e,i} + \varepsilon_{t,e,i}, \tag{1}$$

where $\bar{P}_{t,e}$ is the mean price in trading period t ($t = 3, \ldots, 15$) of experiment e; $F_{t,e,i}$ is the forecast of the period t mean price in experiment e entered by trader i ($i = 1, \ldots, 9$); and $\varepsilon_{t,e,i}$ is the random error term. Forecasts are "accurate" if they are unbiased predictors of the mean price. The REN hypothesis implies the inability to reject the joint null hypothesis $(\alpha_1, \beta_1) = (0, 1)$. REN is the correct interpretation here since we are not asking whether prices correspond to some specific theoretical prediction but simply whether prices and forecasts are mutually supportive.

The OLS estimation of equation (1) yields

$$\bar{P}_{t,e} = \underset{(5.98)}{.208} + \underset{(-8.25)}{.844}\, F_{t,e,i}, \qquad R^2 = .823,\ N = 852. \tag{1'}$$

The numbers in parentheses are t ratios associated with the null hypotheses $\alpha_1 = 0$ and $\beta_1 = 1$. Both indicate rejection at any standard level of significance as does the test of the joint null hypothesis $(\alpha_1, \beta_1) = (0, 1)$ which yields $F_{(2,850)} = 38.9$. Clearly, there is a systematic tendency for forecasts to deviate from the observed mean price.[6]

We also estimated equation (1) for each of the ten forecasting experiments and for each individual subject in three forecasting experiments (39xsf, 41f, and

[5] When testing REM using field survey data, investigators assume implicitly that observed prices are randomly distributed about some theoretical equilibrium price. It should be emphasized that unless this assumption is satisfied, these investigations are testing REN. No distinction between REM and REN is possible without experimental control of dividends.

[6] Estimation of (1′) using the opening price as the dependent variable rather than the mean price does not alter this result. The coefficient estimates and test statistics are very similar to those using the mean price. The coefficient of correlation between the mean and opening price is $r = .97$. Forecast accuracy was also evaluated using the change in the observed mean price $(\bar{P}_t - \bar{P}_{t-1})$ as the dependent variable and the predicted change in the mean price $(F_t - \bar{P}_{t-1})$ as the independent variable. The results indicate that the null hypothesis $(\alpha, \beta) = (0, 1)$ must be rejected ($F_{2,779} = 93.8$).

TABLE II

REGRESSION ESTIMATES: PROFIT VERSUS ABSOLUTE FORECASTING ERROR

Experiment	N (Subjects)	Forecast Error Regression Coefficient
30xsf	9	−4.4
39xsf	9	−3.0
41f	12	−2.55**
42xf	9	−1.1*
43xnf	9	−0.058
48xnf	9	−0.83
49xnf	9	−1.4**
50xxf	9	−0.43*
90f	9	−1.1
124xxf	9	−0.74

*Significant Pr ≦ 0.05.
**Significant Pr ≦ 0.01.

124xxf). The results strongly parallel those shown in equation (1′). It is clear, however, that some subjects were much better forecasters than others. Furthermore, there was a tendency for the better forecasters to earn more money. We established this by regressing profit on absolute forecasting error across individual subjects in each of the ten forecasting experiments. Table II lists the coefficient estimate for the forecast error variable. The coefficient is negative in all ten regressions with four estimates being significant at the 95% level. Greater accuracy in forecasting is associated with greater profit. This is consistent with, but does not prove, the proposition that the better forecasters *acted* on their forecasts to earn higher profits.

Figure 14 summarizes the accuracy of subject price forecasts using the forecast error frequency polygon ($F_t - \overline{P}_t$ rounded to the nearest .05 node) generated by a pooling of all individual forecasts across trading periods 3–15. The sample distribution is not abnormal in appearance but is slightly skewed toward positive forecast errors with mean, median, and modal error of .049, −.01, and −.05, respectively. The vast majority of the forecasts are within one standard deviation of the mean. Can the forecast errors depicted in Figure 14 (pooled across time and subjects) be characterized as a sample of independent draws from a single random variable? We address this question by formally testing the null hypotheses: (i) serial independence of forecast errors, and (ii) no systematic relationship between forecast errors and changes in the forecasting objective.

Serial independence of forecast errors implies the inability to reject the null hypothesis $\beta_2 = 0$ in the equation

$$(2) \qquad (F_t - \overline{P}_t) = \alpha_2 + \beta_2(F_{t-1} - \overline{P}_{t-1}) + \varepsilon_t,$$

where indexing over experiments (e) and individuals (i) is implied. Our alternative hypothesis, based on evidence presented in Williams (1987) and inspection of the charts of the forecasting experiments, is that forecast errors are positively autocorrelated implying $\beta_2 > 0$. OLS estimation of equation (2) for $t = 3, \ldots, 15$

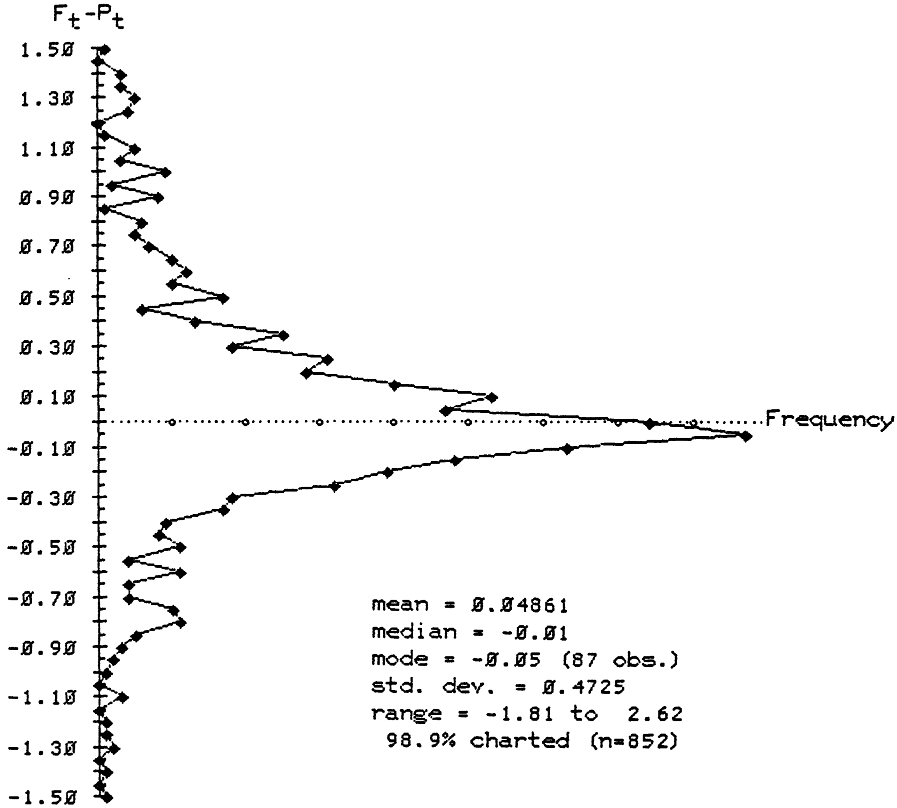

FIGURE 14.—Forecast error frequency polygon.

yields:

$$(2') \qquad (F_t - \bar{P}_t) = \underset{(3.03)}{.046} + \underset{(8.66)}{.282}\,(F_{t-1} - \bar{P}_{t-1}), \qquad R^2 = .089,\ N = 769.$$

The t ratio shown in parentheses under the slope coefficient estimate indicates that the null hypothesis of serial independence is rejected.

Given that forecast errors tend to persist over time, we now ask how forecast errors are linked to changes in the forecasting objective. Figure 14 shows that the distribution of forecast errors is fairly symmetric with a slight tendency for subjects to over-predict the mean price. However, inspection of the charts of the forecasting experiments with bubbles clearly indicates a tendency for the mean forecast to under-predict the mean price during booms ($F_t < \bar{P}_t$) and over-predict the mean price ($F_t > \bar{P}_t$) during crashes. Thus, forecast errors appear to be inversely related to changes in the forecasting objective. More formally, for the equation

$$(3) \qquad (F_t - \bar{P}_t) = \alpha_3 + \beta_3(\bar{P}_t - \bar{P}_{t-1}) + \varepsilon_t$$

TABLE III

STATISTICS FOR THE DISTRIBUTION OF $x_{i,t} = F_{i,t} - E(D_t^T)$

Group[a]	N	Mean[b]	Median	Range	Variance	F_{12} ratio[c]	F_{23} ratio[c]	F_{13} ratio[c]
x^1: 30xsf, 39xsf, 41f, 90f, 124xxf	672	0.93	0.86	−3.36 +5.70	2.5953	5.04		
x^2: 42xfa, 43xnf, 48xnf, 49xnf	504	0.32	0.22	−3.66 +2.87	0.5155		7.40	37.27
x^3: 42xfb, 50xxf	251	0.14	0.14	−2.63 +0.76	0.0696			

[a] x^1, experiments in which subjects had not previously participated in a 15 period asset market. kx^2, experiments in which some (all in 42xfa) subjects had participated in a previous 15 period asset market. (Experiment 42xfa represents periods 1–15 of 42xf.) x^3, experiments in which all subjects had been in the same two previous 15 period asset markets. (Experiment 42xfb represents periods 16–30 of 42xf.)
[b] Each of these means is significantly different from each other using either a t test or a rank sum test.
[c] F_{ij}, the F statistic for groups i, j.

this implies rejection of the null hypothesis $\beta_3 = 0$ in favor of the alternative hypothesis $\beta_3 < 0$. Estimation of equation (3) yields

$$(3') \qquad (F_t - \bar{P}_t) = \underset{(-7.30)}{-.077} \; \underset{(-33.4)}{-.824}\,(\bar{P}_t - \bar{P}_{t-1}), \qquad R^2 = .589,\ N = 781.$$

The null hypothesis $\beta_3 = 0$ is easily rejected in favor of the one-tailed alternative $\beta_3 < 0$, and we see a pronounced tendency to under-predict in expansions and over-predict in contractions. This lagged updating of forecasts relative to movements in the mean price combined with positively autocorrelated forecast errors suggests that forecasts were being formed adaptively.

Expectations are considered adaptive if $0 < \beta_4 < 1$ and $\alpha_4 = -E(\tilde{d}_t)$ (if all agents are risk neutral) in the equation

$$(4) \qquad (F_t - F_{t-1}) = \alpha_4 + \beta_4(\bar{P}_{t-1} - F_{t-1}) + \varepsilon_t.$$

The adaptive expectations model states that the current forecast updates the previous forecast by subtracting the expected single period dividend and adding a fraction of the previous period's forecast error. OLS estimation of equation 4 for $t = 3, \ldots, 15$ yields

$$(4') \qquad (F_t - F_{t-1}) = \underset{(12.11)}{-.117} + \underset{(38.16)}{.815}\,(\bar{P}_{t-1} - F_{t-1}), \qquad R^2 = .632,\ N = 850.$$

The t ratios shown in parentheses under the coefficient estimates indicate rejection of the null hypotheses $\alpha_4 = -.24$ and $\beta_4 = 0$, respectively. The null hypothesis $\beta_4 = 1$ is also rejected ($t = -8.66$). These results indicate that forecasts are adaptive in the sense that $0 < \beta_4 < 1$; however, there is a persistent forecasting bias with $\alpha_4 > -E(\tilde{d}_t)$. As will be shown below, this bias is consistent with agents being risk averse in dividends.[7]

[7] If agents utilize the most recent price information to formulate their forecasts, this implies that the right-hand side of equation (4) should utilize the closing price in period $t-1$ rather than the mean price. This change yields qualitatively similar results to the estimates reported in (4′): $\hat{\alpha}_4 = -.136$, $\hat{\beta}_4 = .628$, $R^2 = .491$ with strong rejection of the null hypotheses $\alpha_4 = -.24$, $\beta_4 = 0$, and $\beta_4 = 1$.

5.2. *Do Subjects' Forecasts Converge to REM with Experience?*

Table IV provides statistics for the distribution of $x_{i,t} = F_{i,t} - E(D_t^T)$ for three different poolings of experiments according to subject experience. All subjects in the x^1 grouping were inexperienced. Those in x^3 had all participated in two previous 15 period markets, while only some of the subjects in x^2 had been in such previous asset market experiments. These groups show clearly that as the experience level increases across experiments, both the mean deviation and the variance of forecasts relative to the dividend value decline significantly. Consequently, with increasing experience our subjects tend to acquire common intrinsic value expectations as behavioral uncertainty decreases.

5.3. *Price Adjustment Dynamics*

The empirical analysis of Section 5.1 distinguishes forecast (expected) prices, F_t, and observed prices $\bar{P}_t$. According to equation (4′), the linear adaptive forecast error dynamic characterizes the adjustment of forecasts over time, and provides a link between forecast prices and observed prices. In this section we consider a (mean) price adjustment hypothesis, H, for characterizing the intertemporal behavior of observed prices. This hypothesis includes the risk neutral and risk adjusted REM hypotheses as special equilibrium cases.

H: Walrasian adaptive, expected capital gains adjusted, REN,

$$\bar{P}_t - \bar{P}_{t-1} = -E(\tilde{d}) + K + \beta(B_{t-1} - O_{t-1}), \qquad \beta > 0. \tag{5}$$

The mean price change from one period to the next is separable into (at most) three components: a term expressing the decline in expected dividend value, $-E(\tilde{d})$; an adjustment term for risk, K; and a measure of the revealed excess demand for shares arising from capital gains expectations. We postulate that excess demand is positively correlated with excess bids (number of bids entered minus number of offers) in those markets which spontaneously self-generate an expectation of capital gains (losses).[8] This hypothesis was formulated after

[8] The first two terms are easily derived assuming constant absolute risk aversion (CARA). Applying the Arrow-Pratt measure of the risk premium (or charge if the agent is risk preferring), the value of a share is given by

$$V_t = E(\tilde{D}_t^T) - \sigma^2(\tilde{D}_t^T)U''[E(\tilde{D}_t^T)]/2U'[E(\tilde{D}_t^T)].$$

Since $\tilde{D}_t^T = \Sigma_{\tau=t}^T \tilde{d}_\tau$, and $\tilde{d}_\tau = \tilde{d}$ (independent dividend realizations), it follows that $E(\tilde{D}_t^T) = \Sigma_{\tau=t}^T E(\tilde{d}) = (T-t+1)E(\tilde{d})$, and $\sigma^2(\tilde{D}_t^T) = \Sigma_{\tau=t}^T \sigma^2(\tilde{d}) = (T-t+1)\sigma^2(\tilde{d})$. Hence

$$V_t = (T-t+1)E(\tilde{d}) + (T-t+1)\sigma^2(\tilde{d})\frac{U''[(T-t+1)E(\tilde{d})]}{2U'[(T-t+1)E(\tilde{d})]}.$$

If we have CARA, then $U''(m)/U'(m) = -a$, and

$$V_t = (T-t+1)[E(\tilde{d}) - a\sigma^2(\tilde{d})/2].$$

Now assume that market prices average the effect of individual risk attitudes in such a way that the mean price equation has the same form as the typical individual's valuation of a share. Then if capital gains expectations are nil, we have $\bar{P}_t = (T-t+1)(E(\tilde{d}) - K)$ and $\bar{P}_t - \bar{P}_{t-1} = -E(\tilde{d}) + K$, as in (5).

TABLE IV

ESTIMATION OF EQUATION $\bar{P}_t^e - \bar{P}_{t-1}^e = \alpha^e + \beta^e(B_{t-1}^e - O_{t-1}^e) + \beta_t^e$

Experiment	$E(\tilde{d}^e)$	$\hat{\alpha}^e$	$\hat{\beta}^e$	R^2	dw[c]
Group I, Stable Price Markets					
5	0.08	−0.22 (−1.04)	0.014 (0.88)	0.06	1.6
7	0.08	−0.10 (−0.21)	−0.013 (−1.1)	0.09	1.6
42xf	0.24	−0.18 (+0.62)	0.025[b] (2.8)	0.23	2.2
43xnf	0.24	−0.22 (+0.56)	0.0044 (1.7)	0.20	1.9
50xxf	0.24	−0.23 (+0.29)	0.006 (1.2)	0.11	2.0
Group II, Growing Price Markets					
9x	0.12	0.027[a] (+2.6)	−0.00038 (−0.07)	0.0004	2.3
10	0.12	0.20[a] (+5.4)	−0.010 (−2.1)	0.26	1.7
90f	0.24	−0.053[a] (+2.2)	0.0063 (0.50)	0.02	1.6
Group III, Bubble-Crash Markets					
6x	0.12	−0.16 (−0.74)	0.014[b] (3.4)	0.49	2.1
16	0.12	0.058 (+0.76)	0.038[b] (2.2)	0.28	3.0
17	0.08	−0.23 (−1.6)	0.035[b] (4.5)	0.63	2.9
18	0.12	−0.17 (−0.40)	0.029[b] (1.8)	0.21	1.4
25x	0.24	−0.47[a] (−2.2)	0.033[b] (4.3)	0.60	1.6
26	0.16	−0.082 (+0.66)	0.039[b] (3.2)	0.46	1.3
28x	0.24	−0.12 (+0.74)	0.063[b] (3.9)	0.56	2.5
30xsf	0.24	−0.32 (−0.57)	0.073[b] (4.5)	0.63	2.5
36xx	0.24	−0.20 (+0.35)	0.012 (0.71)	0.04	0.5
39xsf	0.24	−0.28 (−0.32)	0.044[b] (4.0)	0.57	2.6
41f	0.16	−0.58 (−1.5)	0.049[b] (2.4)	0.32	0.7
48xnf	0.24	−0.31 (−0.45)	0.031 (1.6)	0.17	2.6
49xnf	0.24	−0.37 (−0.74)	0.030 (1.3)	0.12	2.3
124xxf	0.24	−0.31 (−0.50)	0.025[b] (2.0)	0.25	1.3

[a] Intercept is significantly different from $-E(\tilde{d}^e)$ (two-tailed test, $p < 0.05$).
[b] Walrasian coefficient of adjustment speed is significantly positive (one-tailed test, $p < 0.05$).
[c] dw: Durbin-Watson statistic.

examination of the data from two experiments suggested that there might be a tendency for the number of bids (as a measure of demand intensity) to thin relative to the number of offers in the period (or periods) prior to a crash in contract prices. Our interest in this potential regularity in the data was heightened when we realized that it might be expressed as a lagged Walrasian adjustment hypothesis in which excess bids are a surrogate for excess (capital gains) demand. We conjectured that excess bids might be correlated with excess demand because at a price below that which is market clearing there are more willing buyers than willing sellers, and this might be revealed in the context of the double auction institution by the simple numerical excess of bids over offers.

Concerning the interpretation of H, three remarks are appropriate.

1. If we interpret the capital gains component of the price change in H as a price change that is literally *expected* by the traders, then why do rational traders not act on that expectation in period $t-1$ and drive this component of the price change to zero? The answer is contained in a different interpretation of H in which it is proposed that traders *do not expect* the price that occurs in period t; they expect the price they forecast, which is adaptively error prone. In particular, traders fail to predict large price changes and turning points. By this interpretation what the excess bids variable does is to predict trader excess demand in the next period; i.e. excess bids measures *potential* excess demand, with that excess demand impinging on subsequent price realizations. These price realizations *then* produce or reinforce new price expectations. This scenario is consistent with the behavior of price forecasts, and with the excess bids hypothesis. The behavioral mechanism postulated by this interpretation is as follows. In the market's bull phase, if bidding activity is strong in t, with many bids not being accepted, this signals a strong willingness-to-pay and presages an *incentive* induced increase in bid levels in the next period; i.e. traders experience rejection (nonacceptance) of their bids and are motivated to bid higher. Similarly, a thinning of bids, even though at a higher contract price level in t, with few bids failing to be accepted, presages an incentive induced decline in bid levels in the subsequent period. A symmetrical argument would also apply to offers. Traders do not expect this change in prices either because they fail to be aware of excess bid activity, or fail to anticipate the incentive response to high levels of bid rejection. Figure 9 for experiment 28x, illustrates changes in the bid-offer activity over the course of a bubble which are consistent with this interpretation. Notice that the large volume of excess bids in periods 1 and 2 are followed by jumps in bid levels in periods 2 and 3. In periods 3 and 4, when excess bids become negative, there follows a reduction in bid levels in periods 4 and 5, and so on. We have no insight concerning the deeper homegrown source of the endogenous expectations that give rise to positive or negative excess bids. It is not evident that the ultimate "cause" or source of such expectations can be formulated in terms of a traditional dynamic.[9]

[9] What we have in mind has been articulated by Coleman (1979, p. 280) in his discussion of "a panic of the sort that sometimes occurs in a crowded theater. The most puzzling question here is not why panics occur, but why their occurrence is so uncertain. In one situation, a panic will occur,... In

2. We regard H as directly representing rational expectation in the sense of Nash because capital gains (losses) expectations, if they persist, must be sustained by the subsequent observation of rising (falling) prices. Thus in Figure 9, an expectation of rising prices is sustained by the price outcomes in periods 1–4. The price decline in period 5, although presaged by the relative thinning of bids in periods 3 and 4 but not anticipated (forecast) by the traders, induces an adaptive reversal of trader price expectations which are then sustained by falling observed prices. However, in an equilibrium sense, H represents rational expectations in the sense of Muth, since it implies convergence to risk-neutral or risk-adjusted intrinsic dividend value if, and when, excess bids stabilize at zero. This interpretation of H is not inconsistent with the view articulated by Lucas (1986), which is supported by the examples he cites and the experiments reported by Williams (1987), in which adaptive expectations may be part of a transient (learning) process that culminates in a rational expectations equilibrium.

3. The risk-neutral and risk-adjusted REM hypotheses are special cases of H. Thus, risk neutral REM (H1) yields

$$\overline{P}_t - \overline{P}_{t-1} = -E(\tilde{d}), \tag{5.1}$$

risk adjusted REM (H2) yields

$$\overline{P}_t - \overline{P}_{t-1} = -E(\tilde{d}) + K. \tag{5.2}$$

We propose to test H, and its special cases H1 and H2, by estimating for each experiment the equation

$$\overline{P}_t - \overline{P}_{t-1} = \alpha_5 + \beta_5(B_{t-1} - O_{t-1}) + \varepsilon_t. \tag{5.3}$$

If we reject the null hypothesis $\beta_5 \leqq 0$ this supports H.[10] If we are unable to reject $\beta_5 = 0$, but we reject $\alpha_5 = -E(\tilde{d})$, with $\hat{\alpha}_5 > -E(\tilde{d})$, this supports a risk averse interpretation of H2, while if we reject $\alpha_5 = -E(\tilde{d})$ with $\hat{\alpha}_5 < -E(\tilde{d})$, this supports the risk-preferring interpretation of H2. Finally, if we are unable to reject the null hypotheses $\beta_5 = 0$ and $\alpha_5 = -E(\tilde{d})$ this supports H1. We were somewhat skeptical, a priori, that α would be statistically very close to $-E(\tilde{d})$.

other apparently similar situations, a panic fails to take place. Why? Another observation is that training, such as fire drills, is effective for panics... initiated by fire." Our market bubbles project this same kind of uncertainty. Two groups seem to have similar experiences (e.g. 28x and 42xf), but one path is "near" intrinsic value, the other yields a bubble. Yet it appears that if most of the members of any group return repeatedly (42xf, 43xnf and 50xxf), this will produce "near" intrinsic value pricing, much like the effect of fire drills on the propensity to panic in the face of fire.

[10] The research (Walrasian) hypothesis is that $\beta_5 > 0$, so a one-tailed test of the null alternative, $\beta_5 \leqq 0$, is appropriate. We should note, however, that the statistical meaning of coefficient tests will be compromised for this particular sample of experiments in an imprecise way by the fact that casual examination of two experiments (25x, 26) was a key factor leading to the formulation of H. This is not a significant problem for experimental methodology since one can always run new experiments. In the field, one cannot rerun the world, and all tests suggested by the data are questionable, if not irrelevant. A conservative way to report on the present sample is to exclude the two experiments where data were examined in advance.

From previous double auction experiments with induced supply and demand arrays we have observed that different subject groups confronted with the same market parameters vary considerably in terms of the number of price quotations entered by each side. Such group asymmetry need not reduce the *relative* effect of excess bids on price changes but price changes might disappear with nonzero excess bids yielding a contribution to the intercept that is unrelated to the expected dividend.

Pooling across all 12 experiments with $E(\tilde{d}) = .24$, OLS estimation of equation (5.3) for $t = 2, \ldots, 15$ yields

$$(5.3') \qquad (\bar{P}_t - \bar{P}_{t-1}) = \underset{(0.29)}{-.230} + \underset{(7.55)}{.027}\,(B_{t-1} - O_{t-1}), \qquad R^2 = .240,\ N = 182.$$

The t-values indicate that the null hypothesis $\alpha_4 = -.24$ cannot be rejected but the null hypothesis $\beta_5 = 0$ can be rejected at the 95% confidence level. This result generally supports the risk-neutral interpretation of H but since α_5 is estimated to be greater than $-E(\tilde{d}) = -.24$ there is weak evidence that subjects tend, on average, to display risk aversion.

Our application of H to *interperiod* mean price adjustments might be compromised if excess bids are positively correlated with *intraperiod* price movements. That is, $(\bar{P}_t - \bar{P}_{t-1})$ may be at least partially generated by price changes within period $t-1$ that are related to $(B_{t-1} - O_{t-1})$. To test for this *intraperiod* effect of excess bids on prices we estimate, for the same sample of twelve experiments used in equation (5′),

$$(6) \qquad (P^C_{t-1} - P^O_{t-1}) = \alpha_6 + \beta_6(B_{t-1} - O_{t-1}) + \varepsilon_{t-1}$$

where P^C_{t-1} and P^O_{t-1} are the closing and opening price in period $t-1$. Rejection of the null hypothesis $\beta_6 = 0$ in favor of the one-tailed alternative $\beta_6 > 0$ would indicate a significant *intraperiod* excess bids effect which tends to confound our interpretation of equation (5′). The null hypothesis $\alpha_6 = 0$ implies that intraperiod price trends will tend to be absent when the number of bids is equal to the number of offers. OLS estimation of (6) for $t = 2, \ldots, 15$ yields

$$(6') \qquad (P^C_{t-1} - P^O_{t-1}) = \underset{(1.19)}{.045} + \underset{(2.71)}{.010}\,(B_{t-1} - O_{t-1}), \qquad R^2 = .039.$$

The null hypothesis $\alpha_6 = 0$ cannot be rejected. However, the null hypothesis $\beta_6 = 0$ can be rejected confirming the existence of a small intraperiod excess bids effect. This result suggests that a more rigorous test of the ability of excess bids to predict *interperiod* price movements is provided by the regression specification

$$(7) \qquad (P^C_t - P^C_{t-1}) = \alpha_7 + \beta_7(B_{t-1} - O_{t-1}) + \varepsilon_t$$

since changes in the closing price from period $t-1$ to period t cannot be due to price adjustments within period $t-1$. OLS estimation of equation (7) yields

$$(7') \qquad (P^C_t - P^C_{t-1}) = \underset{(0.55)}{-.219} + \underset{(5.09)}{.020}\,(B_{t-1} - O_{t-1}), \qquad R^2 = .126,$$

which is quite consistent with equation (5′) although the predictive power of the model and the significance of the excess bids coefficient are somewhat diminished. Our conclusion that the data are generally supportive of a weakly risk-averse version of H is unchanged.[11]

The regression results for equation (5.3) are listed in Table IV for each of the 22 (of 26) experiments in which there were no interventionist treatment conditions (price controls, computer crashes, or the use of confederates). Table IV groups the experiments into a "stable markets" (group I) class, "growing markets" (group II), and "price bubble—crash markets" (group III). This classification was made on the basis of the charts, Figures 2–13, in advance of the regression estimates. Group I consists of those markets in which prices appeared to follow dividend value, or were constant or followed approximately parallel with dividend value over most of the horizon. Group III consisted of markets that produced price bubbles that collapsed sometime *before* the final period. Group II consisted of the experiments not in I and III, and are called "growing price" markets. Group II includes experiment 9x which grew asymptotically from below to dividend value, and thus seems to suggest risk-averse REM.

From Table IV it is seen that the support for H in Group III markets is exceptionally strong. The adjustment speed coefficient, $\hat{\beta}_5$, is positive in every experiment in this group. Furthermore, we reject the null hypothesis, $\beta_5 \leqq 0$, in 11 of the 14 experiments in Group III.[12] As we interpret it the strong support for H in Group III is because (a) lagged excess bids is indeed a consistent and, in eleven cases, a strong predictor of price changes, and (b) this group is rich in price jumps as well as turning points, thereby allowing any potential predictive power of excess bids to swamp the noise in price adjustments. We think that excess bids is a good leading indicator of stock price changes because in the period prior to a jump in contract price traders fail to be aware of the greater relative intensity (number) of bids being entered (and not all accepted) or fail to anticipate that this portends an increase in bid prices the next period. Hence, their forecasts under-predict price increases, but excess bids is relatively accurate. Similarly, just prior to a downturn in prices in a bull market, traders fail to be aware that bids are relatively thinner even though contract prices are still increasing (their forecasts at the end of the period will now over-predict realizations), and that this portends lower prices in the future. But these characteristics are only tendencies obscured by much noise, with our traders having to rely on their perceptions unreinforced by data analysis. *We* were not aware of it until after studying our data.

[11] We also estimated equations (5), (6), and (7) for a pooling of the six experiments (with no experimenter intervention) where $E(\tilde{d}_t) = .12$. The results are qualitatively similar to those shown in (5′), (6′), and (7′).

[12] It is natural to conjecture that capital gains expectations, and therefore price adjustments, are heavily influenced by end effects, but if we add $T - t + 1$ as a presumed "independent" variable in the regressions of Table IV we get no important improvement. The coefficient of the added variable, $T - t + 1$, is significant in only two experiments (90f and 39xsf). The dynamics associated with the horizon time remaining is already adequately taken into account by excess bids.

In groups I and II there is also some support for H, namely, 5 of the 8 experiments in these two groups yield $\hat{\beta}_5 > 0$, and $\beta \leqq 0$ is rejected in favor of $\beta > 0$ in one case, 42xf. But this case is of particular interest because the chart for this experiment (Figure 11) appears to the eye to show early and strong convergence to REM. But $\hat{\beta}_5$ is significantly positive suggesting that the small changes in mean price constitute a pattern of fluctuation (mini-booms and busts) which were, on average, anticipated by lagged excess bids. One is reminded of what commodity traders call "scalping"—trading on small price movements of only a few cents. Similarly, experiments 43xnf and 50xxf exhibit positive (if not significant) coefficients of adjustment speed, although the charts of mean prices for these experiments (Figures 12 and 13) suggest that the market trades very near to the REM price from beginning to end. These cases show that in a macromarket sense one might have "close" support for rational expectations, but within the interval of "close" there may be a subtle trading dynamics fed by expectations of modest capital gains. In experiments 42xf, 43xnf, and 50 xxf, we estimate $\hat{\alpha}^e > -E(\tilde{d}^e)$, but in none of these experiments can we reject the null hypothesis that $\alpha^e = -E(\tilde{d}^e)$. Experiment 9x is the only market providing strong support for a risk averse adjusted REM. Pooling the four experiments (9x, 42xf, 43xnf, and 50xxf) which exhibit the strongest support for REM we estimated (5.3) as follows:

$$\overline{P}_t^e - \overline{P}_{t-1}^e = \underset{(0.99)}{-0.15} + \underset{(3.1)}{0.013}\left(B_{t-1}^e - 0_{t-1}^e\right), \qquad R^2 + 0.12,\ n = 71.$$

Thus across all "REM experiments" the Walrasian coefficient of adjustment speed is significantly positive. Although the intercept shows risk aversion $\hat{\alpha} = -0.15 > -E(\tilde{d})$ $(= -0.21)$ across the four experiments, the difference is not significant.

Only 4 of 22 experiments yield $\hat{\alpha}_5$ estimates significantly different from $-E(\tilde{d}^e)$ (experiments 9x, 10, 25x, and 90), and across all 22 experiments 11 show $\hat{\alpha}^e > -E(\tilde{d}^e)$ and 11 show the reverse, suggesting no consistent tendency toward either risk aversion or risk preferring. We think this is because any adjustment for risk is small, relative to price variability due to capital gains expectations.

6. SUMMARY: WHAT HAVE WE LEARNED?

Our conclusions will be summarized under four headings:

6.1. *General Conclusions*

1. Inducing different private dividend values on different traders, as has characterized previous asset market experimental designs, is not a necessary condition for the observance of trade. Exchanges, sometimes in large volume,

occur when identical probabilistic dividends are to be paid on share holdings at the end of each period, and this fact is common knowledge. Consequently, it appears that there is sufficient homegrown diversity in agent price expectations and perhaps risk attitudes to induce subjective gains from exchange. This is not inconsistent with the rational bubble literature which assumes that traders have common priors. Our subject traders tend, with experience, to acquire common, intrinsic dividend value, rational expectations.

2. Previous studies that reinitialize and replicate a two or three period dividend environment all report convergence toward REM prices in successive replications. All of our experiments with experienced traders, and most of those with inexperienced traders converge to "near" the REM price prior to the last trading period. Thus our results, and those of the more structured markets reported in previous asset market studies all support the view that expectations *are adaptive*, and the adaptation over time is to *REM equilibrium outcomes* when asset value "fundamentals" remain unchanged over the horizon of trading.

3. Of the 22 experiments that did not involve experimenter intervention or inadvertent disruptions, the modal outcome (14 experiments of which 9 used experienced subjects) was a market characterized by a price bubble measured relative to dividend value.

4. Four experiments, all using experienced subjects, provide the strongest support for the REM model of asset pricing.

5. Regardless of the pattern of price movements the volume of exchange tends to be less for experienced than for inexperienced subjects. Although the divergence in agent price expectations tend to persist with experience, this divergence is attenuated, and markets become thinner.

6. None of the above conclusions are inconsistent with the Fama (1970) criterion for REM (no arbitrage profits), or with the Tirole (1982) model (agents have common priors). Here is what we learn from these experiments, their immediate predecessors, most of experimental economics, and from the examples cited in Fama (1970) and Lucas (1986). Real people in any environment usually do not come off the stops with common expectations; they usually do not solve problems of maximization over time by ex ante reasoning and backward induction, nor is this irrational when there is insufficient reason to believe that expectations are common.[13] What we learn from the particular experiments reported here is that a common dividend, and common knowledge thereof is insufficient to induce initial common expectations. As we interpret it this is due to agent uncertainty about the behavior of others. With experience, and its lessons in trial-and-error learning, expectations tend ultimately to converge and yield an REM equilibrium.

[13]Other experimental evidence supports our interpretation, namely that it is the failure of the assumption of common expectations, not backward induction incompetence by subject agents that explains bubbles. Thus, Cox and Oaxaca (1986) find that subject behavior is strongly consistent with the predictions of a job search model, requiring maximization over time using backward induction, but their subjects are making decisions in a game against nature which requires them to form expectations only about their own future behavior. Behavioral uncertainty is thus minimized.

6.2. *Forecasting Behavior*

1. In every (ten) forecasting experiment agent forecasts fail to be unbiased predictors of the mean contract price in period t.

2. The forecasts fail to predict abrupt increases and decreases in price and consistently fail to predict both upper and lower turning points. Both the mean forecast and the individual forecasts show a tendency to over-predict the mean price. However, in the bubble experiments the forecasts under-predict in the boom phase and over-predict in the crash.

3. Individual agents vary in their prediction accuracy with some agents being better forecasters than others. Furthermore, the better forecasters tend to earn more money.

4. The forecasts are highly adaptive: i.e. the change in forecasts from one period to the next is significantly and positively related to the forecasting error in the previous period. Also, the forecasting errors are autocorrelated.

5. Both the mean deviation and the variance of individual forecasts relative to dividend value decline significantly with increasing subject experience across experiments. With experience subjects tend to converge to common dividend value expectations as behavioral uncertainty decreases.

6.3. *Empirical Characteristics of Market Bubbles*

1. Experienced subjects frequently produce a market bubble, but the likelihood is smaller than for inexperienced subjects. When the same group returns for a third market, the bubble disappears (except that we do observe "scalping" on small price fluctuations).

2. In every market bubble experiment (Group III, Table IV), the mean price in the first period was below $E(\tilde{D}_t^T)$. This suggests the possibility that risk aversion plays a role in market bubbles by depressing prices at first, with the subsequent recovery (after such preferences are satisfied) helping to create or confirm expectations of capital gains.

3. The crash in market prices following a boom, whether with experienced or inexperienced subjects, occurs on a trading volume that is smaller than the volume during the bubble phase.

4. The collapse of price bubbles tends to be presaged by a thinning of bid relative to offer activity, as measured by excess bids (number of bids minus number of offers), in the period or periods immediately before the collapse. Similarly, a subsequent recovery or stabilizing of prices tends to be presaged by an increase in excess bids. Thus the change in mean price in all bubble-crash experiments is positively related with lagged excess bids, and the null hypothesis that this adjustment speed coefficient is nonpositive is rejected in 11 of 14 cases.

5. The tendency for experienced traders to produce price bubbles is not eliminated if we first "train" subjects in a sequence of single-period asset markets that controls for interperiod capital gains by initializing the asset holdings prior to trading in each period. This result is contrary to the conjecture that bubbles

with experienced subjects are caused by expectations of a bubble created in the markets in which the subjects acquired their experience.

6.4. *Characteristics of Markets That Most Strongly Support REM*

1. All four markets providing the strongest support for REM yield intercepts in the regression equation (5.3) that exceed the risk neutral prediction, $-E(\tilde{d})$. This supports the risk-averse adjusted version of REM. In only one of the four cases can we reject the null hypothesis that the intercept is $-E(\tilde{d})$. Also this null hypothesis cannot be rejected if we pool all four experiments in estimating equation (5.3). We conclude that there is weak support for the risk averse model of REM.

2. The three experiments (42xf, 43xnf, and 50xxf) that appear to converge to $E(\tilde{D}_t^T)$ in the first 1–3 periods, and to follow closely the path of $E(\tilde{D}_t^T)$ thereafter, all yield a positive coefficient of adjustment speed in equation (5.3). In one experiment (42xf) we reject the null hypothesis that this coefficient is nonpositive; we also reject the null hypothesis when the three experiments are pooled to estimate equation (5.3). Thus the Walrasian adaptive capital gains adjustment hypothesis receives support even in those experiments which appear to provide the strongest support for REM. We conclude that even these experiments are not an exception to the general conclusion that the REM model of asset pricing is supported only as an *equilibrium concept* underlying an adaptive capital gains price adjustment process.

College of Business Administration, University of Arizona, Tucson, AZ 85721, U.S.A.

and

Department of Economics, Indiana University, Bloomington, IN, U.S.A.

Manuscript received February, 1986; final revision received September, 1987.

REFERENCES

Coleman, James S. (1979): "Future Directions for Work in Public Choice," in *Collective Decision Making*, ed. by C. S. Russell. Baltimore: The Johns Hopkins University Press.

Cox, James, and Ronald Oaxaca (1986): "Laboratory Research with Job Search Models," Western Economic Association Meetings, July.

Fama, Eugene F. (1970): "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance*, 25, 383–417.

Forsythe, Robert, Thomas Palfrey, and Charles Plott (1982): "Asset Valuation in an Experimental Market," *Econometrica*, 50, 537–567.

Friedman, Daniel, Glenn Harrison, and Jon Salmon (1984): "The Informational Efficiency of Experimental Asset Markets," *Journal of Political Economy*, 92, 349–408.

Lucas, Robert (1986): "Adaptive Behavior and Economic Theory," *Journal of Business*, 59, S401–S426.

Miller, Ross M., Charles R. Plott, and Vernon L. Smith (1977): Intertemporal Competitive Equilibrium: An Empirical Study of Speculation," *Quarterly Journal of Economics*, 91, 599–624.

Muth, John F. (1961): "Rational Expectations and the Theory of Price Movements," *Econometrica*, 29, 315–355.

NASH, JOHN F. (1950): "The Bargaining Problem," *Econometrica*, 18, 315–335.

PLOTT, CHARLES R., AND SHYAM SUNDER (1982): "Efficiency of Experimental Security Markets with Insider Trading," *Journal of Political Economy*, 90, 663–698.

PLOTT, CHARLES R., AND GUL AGHA (1982): "Intertemporal Speculation with a Random Demand in an Experimental Market," in *Aspiration Levels in Bargaining and Economic Decision Making*, ed. by R. Tietz. Berlin: Springer-Verlag.

SMITH, VERNON L. (1982): "Microeconomic Systems as an Experimental Science," *American Economic Review*, 72, 923–955.

SMITH, VERNON L., AND ARLINGTON W. WILLIAMS (1983): "An Experimental Comparison of Alternative Rules for Competitive Market Exchange," in *Auctions, Bidding and Contracting: Uses and Theory*, ed. by R. Engelbrecht-Wiggans, M. Shubik, and R. Stark. New York: New York University Press.

TIROLE, JEAN (1982): "On the Possibility of Speculation Under Rational Expectations," *Econometrica*, 50, 1163–1181.

WALKER, JAMES M., VERNON L. SMITH, AND JAMES C. COX (1986): "Bidding Behavior in First Prize Sealed Bid Auctions: Use of Computer-Nash Competitors," Indiana University, November, 1986.

WILLIAMS, ARLINGTON W. (1979): "Intertemporal Competitive Equilibrium: On Further Experimental Results," in *Research in Experimental Economics*, Vol. 1, ed. by Vernon L. Smith, Greenwich, Conn.: JAI Press.

——— (1987): "The Formation of Price Forecasts in Experimental Markets," *Journal of Money, Credit and Banking*, 19, 1–18.

WILLIAMS, ARLINGTON W., AND VERNON L. SMITH (1984): "Cyclical Double-Auction Markets With and Without Speculators," *Journal of Business*, 57, 1–33.

ZARNOWITZ, VICTOR (1986): "The Record and Improvability of Economic Forecasting," NBER Working Paper No. 2099, December, 1986.