Vernon L. Smith, Gerry L. Suchanek, Arlington W. Williams

An Experimental Analysis of Stock Market Bubbles: Prices, Expectations and Market Efficiency*

1. Introduction

The recent unsettling gyrations of securities' prices in the world's financial markets pose a major challenge to financial theorists and practitioners, breathing new life into the debate between the behavioralists and the proponents of the efficient markets hypothesis. Germane to this debate are the results that researchers are obtaining from the application of two relatively new methodologies to the study of price fluctuations in asset markets: Chaos Theory and Experimental Asset Markets. Chaos Theory is a search for possibly complex mathematical patterns in actual securities' prices, and the preliminary results challenge the efficient markets hypothesis. Experimental Asset Market studies, on the other hand, investigate security trading behavior using real subjects motivated by cash rewards in a laboratory environment where critical information variables can be controlled by the experimenter. In this paper, we report results from our study of experimental asset markets which present yet another challenge to the received theory of efficient markets. We find that traders readily generate self-reinforcing market bubbles and crashes. It is not necessary to search for a cause hidden in the mechanics of trading, in the capital requirements of specialists and brokers, in the existence of index futures, or in program trading to explain the extraordinary recent volatility of asset prices. These are scapegoats, not fundamental causes.

For more than two decades financial economists have maintained that securities' prices correctly reflect all publicly available information; except for transitory or noisy fluctuations, prices react only to unanticipated events that affect the fundamental values of securities. In other words, investors are rational, in the narrow sense defined by this literature, and rapidly develop common expectations about securities' values even though initially they may have very diverse individual information sets or interpret available information differently. In the absence of new information, trading occurs to compensate for differences in risk attitudes. Prices that deviate from a security's commonly expected or fundamental value cannot be sustained, and therefore must be transitory. Moreover, such price fluctuations must be random and without trend.

This position is inconsistent with the recent market experience, and prominent efficient market theorists admit they are perplexed. Fischer Black, now a partner at Goldman, Sachs & Company, has suggested, however, that investors merely revised their expectations of price volatility upward, increasing the risk premia required by investors (*Wall Street Journal,* October 23, 1987). He does not purport to know why equity prices were so high relative to other types of assets prior to the recent crash, nor does he suggest why investors only recently revised their expectations about volatility. There is no obvious event that would have warranted such a large change, which has

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

prompted some to seek the culprit among institutional restrictions or trading innovations such as the New York Stock Exchange's Superdot automated order routing system. Still others suggest that institutional restrictions, such as daily price change limits, are needed to moderate price volatility and avoid the onset of panic trading.

By way of contrast, behavioralists simply insist that the market is not really efficient. They assert that psychological factors like greed, fear, and panic are important determinants of securities' prices (*Wall Street Journal,* October 23, 1987). Although this means that price movements are not random, it does not mean that traders can exploit price histories to predict future prices accurately as asserted by those who advocate technical analysis. The issue is not whether securities' prices exhibit patterns, but whether they exhibit readily recognizable patterns. It is precisely this question that a small group of financial economists have been investigating in recent years using the tools of Chaos Theory. Motivated by doubts that markets really are efficient, principally because of the growing literature on market anomalies and seasonalities, Professors BROCK, DECHERT, and SCHEINKMAN have pioneered new econometrics tests designed to identify the existence of patterns in prices. Prices that appear to be random or chaotic may in fact possess order where the order is obscured by noise, making the underlying pattern difficult to discern and virtually impossible to arbitrage. Their tests rely on methods that filter the noise and then test the filtered data for randomness. Their preliminary results assail the efficient market theory, make behavioralists' arguments more credible, and do not seem to support the claims of technicians.

The early experimental asset market literature focused on the rational expectations equilibrium hypothesis, posing the following question: If agents have different information about the value of an asset, will the asset's trading price adjust to a level that reveals this information?[1] If so, then agents who trade on their private information initially will adjust their expectations to accommodate the information of other agents as signalled by the asset's trading price.

The experimental asset market design used in these studies consists of a two- or three-period cycle that is replicated several times. Generally, there are three types of traders, distinguished according to the fixed payoff each trader type will receive at the end of a period for each unit of the asset held; these payoffs are different for each period within a cycle, but once stipulated, are maintained across cycles. Each trader knows his own payoffs initially, but not those of other traders. Agents are endowed with inventories of the asset and working capital at the beginning of the experiment, and these inventories are reinitialized at the beginning of each cycle to provide a pure replication of the trading environment. In each period, agents are permitted to trade the asset by submitting bids and offers which others may accept or amend by submitting improved bids or offers.

There are potential gains from exchange in these markets because agents' private (dividend) values are induced on the asset, and because these values are different for different groups of traders. Initially, each agent bids for the asset solely on the basis of his private information. As the value of the asset to other traders is signalled by their bidding activity, this information is gradually incorporated into the individual's bidding strategy. Contract prices usually converge to near the price predicted by rational expectations within five to eight replications of a two-period cycle. Convergence, although much slower, is also observed in the three-period cycle environment. Hence, the early experimental evidence supports the rational expectations hypothesis. It provides some support as well for the efficient markets hypothesis because the asset price adjusts to an equilibrium level that reveals agents' private value information.

This work has several disconcerting features, however. Agents tend not to trade for capital gains, possibly a consequence of the very short capital gains horizons. In addition, trade is induced artificially by assigning different fixed dividend values to different groups of agents. Therefore, an agent's uncertainty is confined to the dividend values assigned to other agents, and this uncertainty is quickly resolved by their bidding behavior which is keyed to privately known but fixed parameters. This might further mitigate any capital gains expectations. Finally, the rapidity of agents' learning behavior seems likely to be sensitive both to the number of periods in a cycle and to the process of pure replication across cycles. Unless the results are robust to more general trading environments, the conclusion that markets are efficient remains moot.

We study an experimental trading environment that corrects for these features to test the robustness of the results[2]. In our framework, there is a financial asset that pays the same random dividend at the end of each trading period to every agent holding the asset. That is, each trader faces the same distribution of possible dividends, and receives whatever dividend is drawn from this distribution at the end of each period. An experiment has 15 or 30 trading periods. Agents begin the experiment with randomly assigned endowments of working capital and shares of the asset drawn from a fixed set of alternatives; these endowments are not reinitialized at any time. Hence, in each experiment agents' asset and cash holdings are endogenously determined once the experiment begins. All agents are given complete information about the dividend distribution and, in some experiments, individual price forecasts are solicited at the end of each period.

The intrinsic or fundamental value of the asset at any time is the expected dividend return from holding the asset until the end of the experiment. There are thus three conceptually distinct price adjustment paths: the process that describes changes in the asset's expected dividend value (which we control), the evolution of agents' price expectations as measured by their forecasts, and the asset's actual price adjustment path which we observe from trading. The rational expectations and efficient markets hypotheses require that these adjustment paths coincide, and the efficient markets hypothesis in particular seems to suggest that this intertemporal congruence occurs 'quickly'. Our principal observations are that no two of the three adjustment paths are ever congruent although all three are close occasionally, and actual prices often trace a bubble-crash path relative to intrinsic value. In effect, these observations imply that the rational expectations hypothesis is inadequate to explain trading behavior. Further, our characterizations of the price adjustment process implies that prices are predictable from current and historical data. More than historical price data must be used to accomplish this, however, contradicting the claims of technical analysts. These results do not mean that the rational expectations hypothesis is an inappropriate equilibrium concept. We suggest, in fact, that the opposite is true; what we observe is out-of-equilibrium trading behavior that reflects traders' uncertainty about the possible trading behavior of others. This generally prevents traders from developing truly common expectations over the course of an experiment, violating an important condition necessary for a rational expectations equilibrium to be attained. But across experiments they tend to converge to intrinsic value.

The remainder of the essay is organized as follows. We discuss the trading environment in greater detail in the next section. Our results are presented in section 3. Concluding remarks are given in section 4, where we discuss the significance of our results for conventional wisdom, for understanding the recent market experience, and for assessing some of the alleged causes of the market crash.

2. The Experimental Asset Market Environment

The experimental asset market mechanism is a computer-automated double auction that is an enhanced version of the PLATO system described by SMITH and WILLIAMS (1984). In all experiments, there are nine or twelve subjects (traders), each of whom is assigned to an interactive computer terminal. The computer is programmed to accept bids from traders who desire to buy and offers from traders who desire to sell; any trader can enter either a bid or an offer by switching their terminal between buying mode and selling mode, accomplished by depressing a single key. The computer passively enforces the rules of the market. Only the highest bid and the lowest offer are displayed to all traders in the market. To be displayed, new bids and offers must improve on the standing bid and offer, currently displayed on the terminals to all traders. That is, a new bid to buy must be entered at a higher price and a new offer to sell must be entered at a lower price. Therefore, price quotes evolve by reducing the standing bid-ask spread. Any bid or offer that violates this improvement rule is placed in a queue, ranked highest to lowest if a bid and lowest to highest if an offer. The queues are held in memory, and are not displayed. This corresponds to an electronic version of the exchange specialist's limit-order book where bids and offers away from the standing bid-ask spread are placed to await possible execution.

When a standing price quote is accepted by a trader, the computer automatically moves the highest bid and lowest offer in the queue into the standing bid-ask spread. To execute a sale (or purchase) at the standing bid (offer), a trader merely touches a touch-sensitive area highlighted on his terminal screen labeled ACCEPT BID (ACCEPT OFFER). A binding contract is formed when the acceptor touches a highlighted screen area labeled CONFIRM CONTRACT within five seconds. Since all bids and offers must specify a (limit) price, there is no such thing as an *at market order.* Trades always involve matched orders (a continuous-auction market). There is no assurance, however, that there will exist a standing bid or a standing offer, or that the spread will be narrow if both do exist.

Each contract constitutes a trade of a single share of the asset. All trades are between independent agents for their own accounts. When binding contracts are made, the appropriate information is recorded in the individual record sheets of the two parties to the contract. An agent's record sheet is maintained by the computer, and is displayed continuously on his terminal's screen. Each record sheet contains the following information: current asset inventory, current cash position, and the buying (selling) price of each transaction and the period in which the transaction occurred. For accounting purposes, inventories and capital gains (losses) are computed on a first-in-first-out basis. A copy of a representative trader's screen display is provided in Figure 1.

Each trader begins an experiment with a randomly assigned endowment of cash and shares of the asset. At the end of each trading period which lasts for 240 seconds, a dividend is declared. Each dividend declared is drawn independently from the same set of four uniformly distributed (or equally likely) alternatives. All traders are informed of this dividend structure before the experiment commences. Specifically, they are told what the four possible dividend declarations are, that each has a 25 percent change of being drawn at the end of each period, and the expected value of the dividend that will be drawn. Traders, however, do not know the actual dividend that

Figure 1: Screen display for asset market trader.

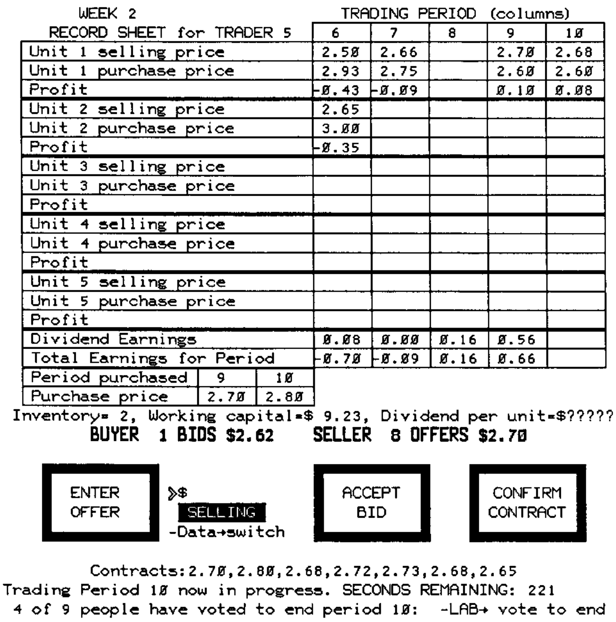

WEEK 2 — TRADING PERIOD (columns)

RECORD SHEET for TRADER 5	6	7	8	9	10
Unit 1 selling price	2.50	2.66		2.70	2.68
Unit 1 purchase price	2.93	2.75		2.60	2.60
Profit	-0.43	-0.09		0.10	0.08
Unit 2 selling price	2.65				
Unit 2 purchase price	3.00				
Profit	-0.35				
Unit 3 selling price					
Unit 3 purchase price					
Profit					
Unit 4 selling price					
Unit 4 purchase price					
Profit					
Unit 5 selling price					
Unit 5 purchase price					
Profit					
Dividend Earnings	0.08	0.00	0.16	0.56	
Total Earnings for Period	-0.70	-0.09	0.16	0.66	

Period purchased	9	10
Purchase price	2.70	2.80

will be declared at the end of any period until that period is concluded. Once a dividend is declared, agents are informed of its value, and dividend earnings are added to their working capital. Before the next trading period begins, agents are also advised of the expected holding value (i.e. the expected dividend earnings) for each unit of the asset, as well as the minimum and maximum possible dividend earnings from holding the asset for the remainder of the experiment. We emphasize that all agents are also told that the dividend structure and the dividends declared are the same for every participant. Therefore, the fundamental dividend value of a share is common information for all traders throughout the experiment since this value is simply the asset's expected holding value in any period.

An experiment typically entails 15 trading periods, and terminates when the last announced period is concluded. All agents are told how many trading periods there will be before an experiment begins, and that the market will not be reopened for any reason after the announced number of trading periods have occurred. When an experiment terminates, each subject receives a cash payment equal to his closing cash balance. This balance is the sum of his initial cash endowment, all dividends earned by him during the experiment, and his accumulated capital gains and losses. A subject's total cash earnings can be greater than, less than, or equal to his initial cash endowment as a consequence of realized capital gains and losses in addition to realized dividend earnings. In one of our experiments, for example, subjects' individual earnings varied from $ 5 to $ 50; according to dividend fundamentals, each should have collected an average of $ 16. This means that some agents experienced significant capital gains whereas others experienced significant capital losses because aggregate capital gains must be zero; in our experimental markets, one trader's gain is another trader's loss.

In our formulation of the experimental asset market, the asset's fundamental value is easily computed in a general way. Let T be the number of trading periods the market will be open, and let t denote the current trading period. Then, at the end of period t, there are $(T-t)$ trading periods remaining in the experiment. Now let d and s^2 represent the mean and variance, respectively, of the dividend distribution; d is the one-period expected value of the dividend. Finally, let $D(t+1)$ represent the asset's expected dividend earnings if it is in inventory at the end of period $(t+1)$ and held to the end of the experiment. Similarly, let $S^2(t+1)$ denote the variance of the asset's uncertain dividend earnings over the remainder of the experiment. Since the distribution from which the dividend is drawn is the same for every period, it follows that the asset's expected dividend (or fundamental) value at the beginning of period $(t+1)$ is $D(t+1) = (T-t)d$ and the variance is $S^2(t+1) = (T-t)s^2$.

The subjects for our experiments comprized three distinct groups by age and background. The majority of our experiments were conducted using undergraduate economics students at the University of Arizona and at Indiana University. In general, these students were recruited from middle-level economics courses, and were in their second or third year at university. A second group of subjects consisted of graduate students in economics, all of whom had participated previously in several computer automated auction and market experiments. Consequently, when participating in our experiments, the subjects in this group already had considerable experience with experimental markets, and a reasonably extensive knowledge of the economic principles and theories underlying the functioning of these markets. The third subject group is distinctive because it consists of professionals and middle-level managers from the business community, some of whom were enrolled in the University of Arizona Executive Development Program. The individuals in this group have diverse backgrounds, but are alike in that all are well-educated with many possessing advanced or professional degrees, and all had had years of business management experience in various forms. None of the subjects, however, had extensive experience as professional money managers.

3. The Dynamics of Asset Market Trading

When we initiated this study, we wondered if subjects in our experiments would trade the asset beyond the first period. The consensus of the previous research, using replicated two- and three-period cycle market environments, was that trading behavior rapidly moves the asset's market price to the rational expectations price prediction, and that traders appear to be risk neutral. If traders are rational in the sense suggested by these studies and have homogenous risk preferences, then there should be no trading in our experimental markets; agents will simply hold their endowed inventories, collecting dividend income. If traders are rational, but heterogenous, then we should observe trading motivated only by differences in agents' risk preferences, and endowments, but these trades should be accomplished in the first trading period. Both assertions follow because all subjects know before an experiment begins that they face identical uncertain dividend payout schedules (assuming that this information is sufficient for traders to adopt common expectations about the value of the asset). This proviso is an important caveat. If trading is active in many periods, then there are two possible explanations. It may be simply because traders are not rational. Alternatively, it may be that traders are rational, but that the information they have been given is not sufficient for them to develop common expectations, which in turn precludes the rational expectations outcome.

We conjectured that agents are in fact rational, but that they are uncertain about how others will react in the marketplace. In general, this behavioral uncertainty varies from agent to agent, thereby engendering different individual expectations. Until the uncertainty is resolved, agents will trade at prices away from the asset's fundamental or dividend value, and the occurrence of contract prices that deviate from fundamental value may or may not be anticipated by traders. This, we knew, would be determined by the relationship between the adjustment dynamics that describe how actual and expected prices move through time, provided these processes could be characterized. Regardless, we conjectured that traders who were able to forecast future prices best would realize the greater total payoffs.

From some 60 experiments, costing approximately $10000 in cash payouts to subjects, we find that inexperienced subjects never trade consistently near fundamental value, and their trading activities usually generate prices that trace a boom-bust pattern (relative to fundamental value). This is illustrated by the experimental results given in Figure 2; trading prices start and stay well above fundamental value. Occasionally, the trading prices generated by even experienced traders appear to be almost chaotic, but with an underlying order characteristic of a boom-bust price pattern. This is illustrated by the experimental results given in Figure 3. With only one exception, maintaining the relative group integrity of a subject pool for a second market experiment does not eliminate the boom-bust price pattern, although the bubbles and crashes are driven by reduced volume. A second trading session with the same group of subjects also tends to mitigate both the duration and the volatility of the bubble. If the group's relative integrity is maintained for a third trading session, then the trading price is very close to the fundamental value throughout the experiment.

Figure 2

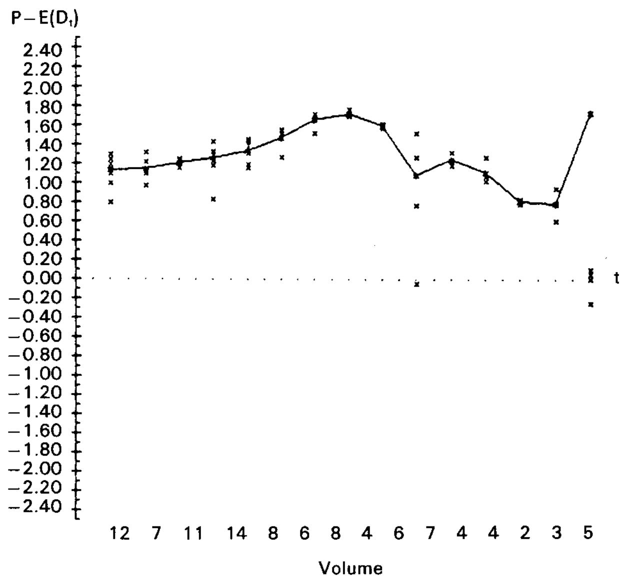

To interpret these results, it is important to stress that experience alone is not sufficient to eliminate the bubble. It is critical in general that subjects acquire the experience as a group, as demonstrated by the experiment given in Figure 4. In this experiment, the subjects selected were 'superstar' traders; all were in at least two asset market experiments previously, although in diverse groups, and all were successful traders as measured by their trading profits. Again

Figure 3

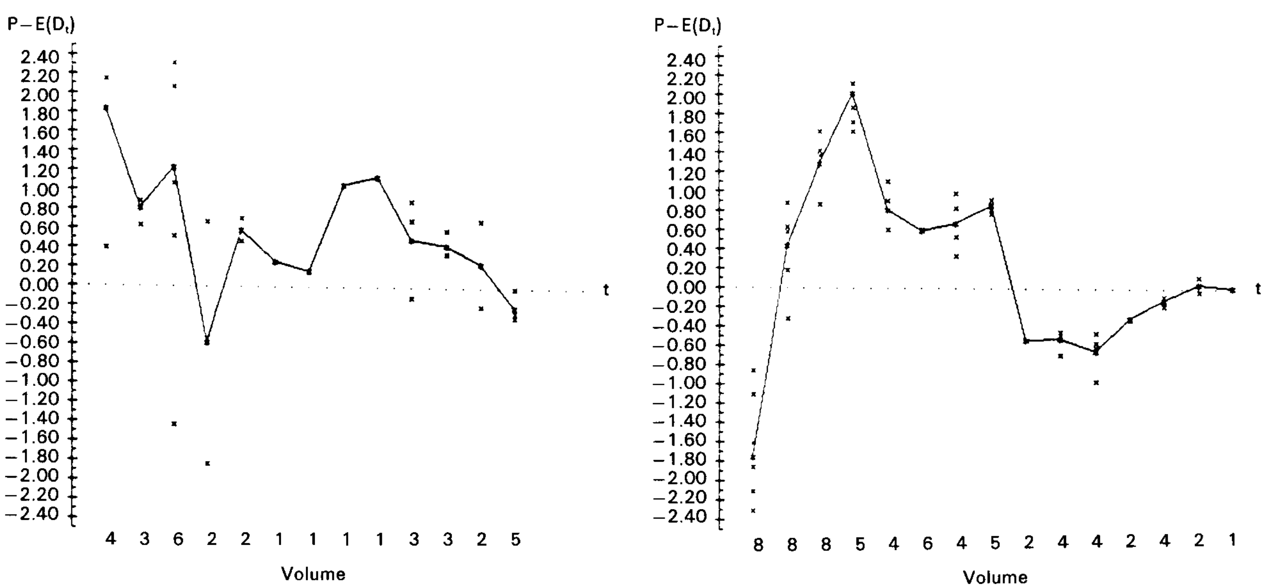

Figure 4

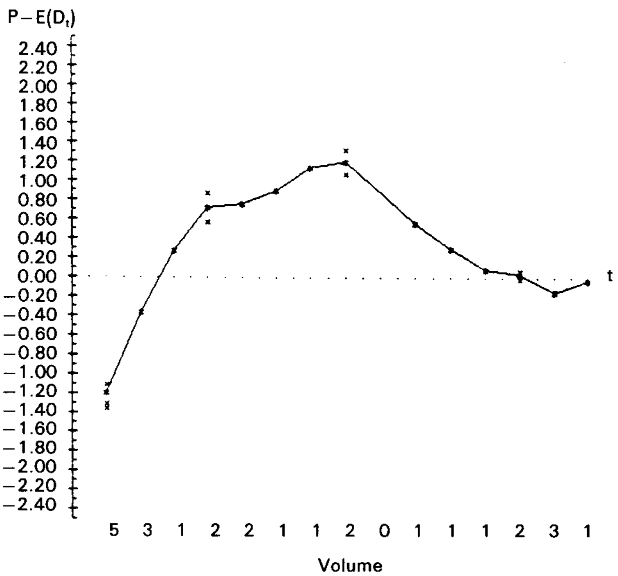

we see that the bubble's duration and volatility are reduced, and that volume is much less than with inexperienced traders. The bubble-crash phenomenon is not eliminated, however. This means that a group of traders will develop common and rational expectations only after the group resolves the uncertainty about how individual members will trade in the market in response to observed prices, and this requires that the group's composition be kept constant long enough for agents to assimilate the necessary information.

We conclude, therefore, that price bubbles and crashes can be expected to occur in any securities market where the set of traders is changing frequently, with some traders exiting the marketplace while others are entering to trade. Moreover, the greater the number of inexperienced traders, the more likely are transactions' prices to appear chaotic, but obscuring an order that approximates the boom-bust phases that are more apparent in markets with only experienced traders. Unlike the market bubbles and crashes that we observe in the laboratory, those that occur in actual securities' markets are incidental to much greater trading volume, and the crash gradients seem to be less steep initially and much steeper later. This is probably because, when market corrections occur in our laboratory markets, both bids and offers fall sharply to near fundamental value, which is known and assured. Since a security's fundamental value is not assured in the economy, a market correction is likely to occur on greater volume and the downturn is less likely to be as rapid initially as in laboratory markets. Furthermore margin calls force involuntary sales, and will tend to increase the rate at which prices fall off as the correction takes hold.

That traders in our laboratory asset markets develop what appear to be common and rational expectations when they have acquired sufficient experience does not mean necessarily that they learn to anticipate actual prices correctly, nor that actual or forecast prices coincide with the asset's fundamental value. For this to occur, the three adjustment paths must become congruent. This is going to be possible if the adjustment processes that describe how the actual price, the forecast price, and the fundamental value evolve over time in response to dividend and trading information differ only by noise terms that become inconsequential with learning. This, in fact, does not appear to be the case. Analysis of our experimental markets data indicates that each of the three adjustment processes is driven by distinct but interrelated forces. In particular, the fundamental asset value is determined strictly by the publicly known dividend structure; the actual mean price change is related to the change in fundamental value and to the magnitude of excess demand for the asset; agents' forecasts of the mean price are linearly adaptive, changing by a percentage (less than 100 percent) of their previous forecast errors where these errors depend on actual prices and hence excess demand.

We tested two expectations hypotheses[3]. First we examined how well traders' forecasts anticipate actual prices, testing the notion that agents' expectations are realized. We accomplished this by regressing actual prices on forecast prices. Let P(t) represent the mean trading price observed in period t, and let F(t, i) and F(t) denote, respectively, agent i's forecast and agents' pooled forecast of the mean trading price in period t. We estimated the equations,

$$P(t) = a_1 + b_1 F(t,i) + e(t,i), \quad (1)$$
$$F(t) - P(t) = a_2 + b_2[F(t-1) - P(t-1)] + e(t), \quad (2)$$
$$F(t) - P(t) = a_3 + b_3[P(t) - P(t-1)] + e(t). \quad (3)$$

All estimates are given in the Appendix. Our estimate of (1) indicates rejection of the joint null hypothesis, $(a_1,b_1) = (0,1)$, at any standard level of significance; this means that agents' forecasts are not realized in general, that their forecasts tend to deviate systematically from actual prices. It was apparent from our data, however, that some agents were better forecasters than others. Pursuing this observation, we regressed agents' profits on their absolute forecasting errors. We observed a marked tendency for greater trading profits to be associated with greater forecasting accuracy, suggesting that agents do trade on their forecasts.

Regressions (2) and (3) lead us to reject the hypothesis that $b_2 = 0$ in favor of $b_2 > 0$, and to reject $b_3 = 0$ in favor of $b_3 < 0$. We conclude from this that forecast errors are positively autocorrelated, and that agents usually tend to underpredict the mean price as a bubble is developing $[F(t) < P(t)]$ and to overpredict the mean price during a crash $[F(t) > P(t)]$. In other words, agents' forecasts lag actual prices, and hence they fail to anticipate the turning points that signal the end of a boom or when prices bottom out. We interpret these results to mean that agents formulate their forecasts adaptively, viewing actual prices as important indicators of developing trends. We do not interpret this to mean, in turn, that agents ignore fundamental value information. To explore this further, we estimated the following equation:

$$F(t) - F(t-1) = a_4 + b_4 [P(t-1) - F(t-1)] + e(t). \quad (4)$$

According to the adaptive expectations model, agents should adjust their forecasts by a percentage of their most recent forecasting error and then adjust for any changes in the fundamental value (which means to subtract d, the expected one-period dividend, in our experimental framework). We found that $0 < b_4 < 1$, rejecting the hypotheses $b_4 = 0$ and $b_4 = 1$, but that $a_4 > -d$ consistently. Thus, agents' expectations are adaptive, but exhibit a persistent forecast bias that is consistent with agents being risk adverse in dividends. If we restrict attention to highly experienced traders, then we find that the average deviation from fundamental value and the variance of these deviations are significantly less than those associated with the pooled data, and especially with the data generated by inexperienced traders. Hence, trading experience is important for reducing market volatility, absent new information indicating fundamental values or their riskiness should be adjusted.

It is worth emphasizing that equation (4), a linear adaptive forecast error dynamic, provides a link between forecast prices and prices observed over time. By exploring this link, we found that the correlation between forecasts and observed prices is significant but less than 50 percent. This suggests that expectations are an important – but significantly imprecise – determinant of actual prices. Most importantly, however, is that this also means that empirical studies of asset price expectations which use observed prices for data (as they must since price expectations data are generally unavailable) are prone to considerable specification error. In short, the results reported in such studies must be viewed with skepticism.

At this point in our analysis, we are left with no explanation of how actual prices adjust over time even though we have considered the prevailing expectations theories. Expectations simply are not an adequate explanation of market price adjustments for we have shown that expectations in general are not realized systematically. Equally important is our observation that fundamental value is not generally a good predictor of market prices. Our data makes it clear that market prices are not random about the trend line given by fundamental value. Yet it seems appropriate to interpret fundamental value as the equilibrium market price since enough experience leads agents to trade the asset at prices near fundamental value. This means, of course, that rational expectations is a reasonable equilibrium concept for our experimental markets, and, more importantly, that we have been observing disequilibrium trading behavior characteristic of agents' learning process.

A disequilibrium adjustment dynamic that is entrenched in economic theory is based on the Walrasian excess-demand hypothesis. According to this hypothesis, the price of a commodity will increase (decrease) if demand is greater (less) than supply. There are two difficulties with this hypothesis for our analysis. It is a static equilibrium notion used to explain the price adjustment process in a Walrasian auctioneering context where, theoretically, trading occurs only in equilibrium and we do not have a direct measure of agents' aggregate demand for the asset being traded in our experimental markets. We conjectured, however, that agents' bids and offers do serve as surrogate measures of the demand for and supply of the asset. At a price below that which is market clearing, there are more willing buyers than willing sellers and we believed that this might be revealed in the context of our double-auction markets by the excess of the number of bids over the number of offers. This seemed especially promising since our data exhibit an important regularity: the number of bids usually exceeds the number of offers while a bubble is developing, the reverse is true in general during a crash, and the number of bids will thin relative to the number of offers just prior to a downturn in contract prices. We tested the idea by formulating a lagged Walrasian adjustment hypothesis based on excess bids.

We hypothesized that the change in the mean trading price from one period to the next has three components: the known change in fundamental value, $-d$; a risk adjustment, k; the excess demand for the asset implicit in the number of excess bids. Let $X(t) = B(t) - 0(t)$, where $B(t)$ is the number of bids posted in period t and $0(t)$ is the number of offers posted in period t; then $X(t)$ is the excess of bids over offers in period t. Our hypothesis is given formally by the equation, $P(t) - P(t-1) = (k-d) + bX(t-1)$ where $b > 0$. If traders are risk averse, then $k > 0$ as well, so that $(k-d) > -d$. We tested the hypothesis by estimating the equation,

$$P(t) - P(t-1) = a_5 + b_5X(t-1) + e(t). \quad (5)$$

Based on our estimate, we reject the hypothesis that $b_5 \leq 0$ at the 95 percent confidence level, enabling us to conclude that $b_5 > 0$. The change in the mean trading price from one period to the next is positively related to lagged excess bids. We are unable to reject the null hypothesis that $a_5 = -d$, so that agents' do not appear to be significantly risk averse as measured relative to the considerable noise in prices. We do reject the hypothesis that $a_5 < -d$, indicating that agents are not risk preferring on average. In fact, the estimate of a_5 is greater than $-d$, which provides weak evidence that traders are risk averse on average.

The salience of equation (5) for explaining observed prices is best illustrated by separating our experiments into three groups defined by: I) the market is stable and shares are traded at prices near fundamental value, II) the market first bubbles and then crashes, and III) the market booms throughout the experiment. For group III experiments, excess bids are generally positive in all periods [$X(t) > 0$ for all t], whereas for group II experiments, excess bids are positive [$X(t) > 0$] during the boom and negative [$X(t) < 0$] during the crash. More notable, however, is that excess bids tend to thin as the boom in prices slows, and excess bids is negative almost always in the period just prior to a decline in the mean trading price, signalling the beginning of a crash. Hence, excess bids appears to be a vital leading indicator of the direction of major price movements, being a critical signal that price bubbles are forming, that the bubble is about to burst, and that a market correction or crash is in progress. In view of this result, the important issue is whether excess bids is a general leading indicator of the direction of price changes, being equally applicable to stable markets.

Before pursuing this issue, it should be noted that the trading behavior exhibited in the experiments in groups II and III is apparently a consequence of agents playing the market bubbles as they form in anticipation of capital gains returns. In the group III experiments, these expectations are generally not revised at any point, and the market price either continues to increase or stabilizes at a level well above fundamental value. These experiments provide good examples of markets where expectations are self-generating and, possibly, where agents on average are unwilling to sell assets at prices less than they paid when making their purchases. *Agents in these experiments are predominantly inexperienced traders.* In the group II experiments, agents appear to become less optimistic that prices will continue to increase once the price has risen well above fundamental value, and begin to sell to realize their gains. According to the Walrasian adjustment hypothesis, of course, this leads to declining prices.

Figures 5 and 6 show plots of three experiments from group I. The experiment in Figure 5 is of interest because it has 30 trading periods, and the early trades were consummated at prices well below fundamental value; contract prices quickly converged to near fundamental value, however, and then remained close to fundamental value for the remainder of the experiment. The experiments in Figure 6, on the other hand, demonstrate markets where shares of the asset trade very near to fundamental value from the opening trade to the end of the experiment. For all three experiments, the coefficient of price adjustment speed in equation (5) is positive (i.e. $b_5 > 0$), indicating that excess bids is a leading indicator of small changes in the mean trading price which tends to trace a pattern of fluctuation that might be termed mini-booms and mini-busts. In other words, there is a subtle trading dynamic that the excess bids variable picks up even in stable markets where prices closely approximate but fluctuate around fundamental value. Thus, excess bids is a good leading indicator of even minor price fluctuations which tend to feed agents' expectations of modest capital gains, and agents trade on these modest price movements to capture the small gains. *We find this to be a remarkable result.*

To conclude this section, we discuss briefly two features of our experimental stock market analysis that may be of concern to the reader. First, the reader may wonder just how robust the excess bids adjustment dynamic is, especially since our results are reported in terms of changes in the mean trading price. In partial answer to this question, we also estimated equation (5) using closing prices for each period instead of the mean price. This estimate is reported as equation (5′) in the Appendix, and shows that our conclusions are unchanged. This means that our result is tenacious, and is not just picking up intraperiod effects. The preliminary results from further research indicates that the excess bids variable is even a reasonable predictor of intraperiod price changes.

Figure 5

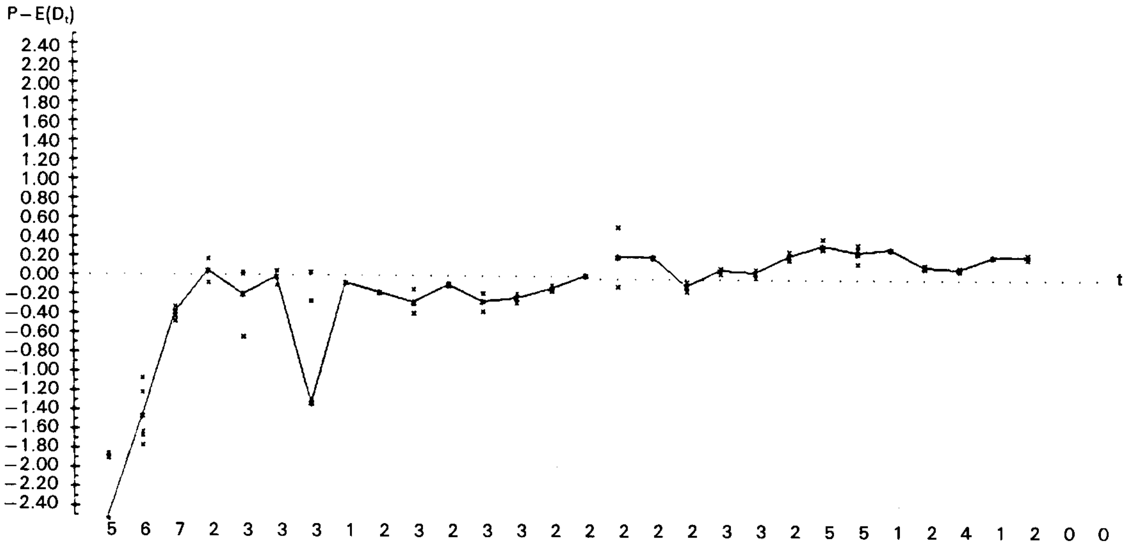

Figure 6

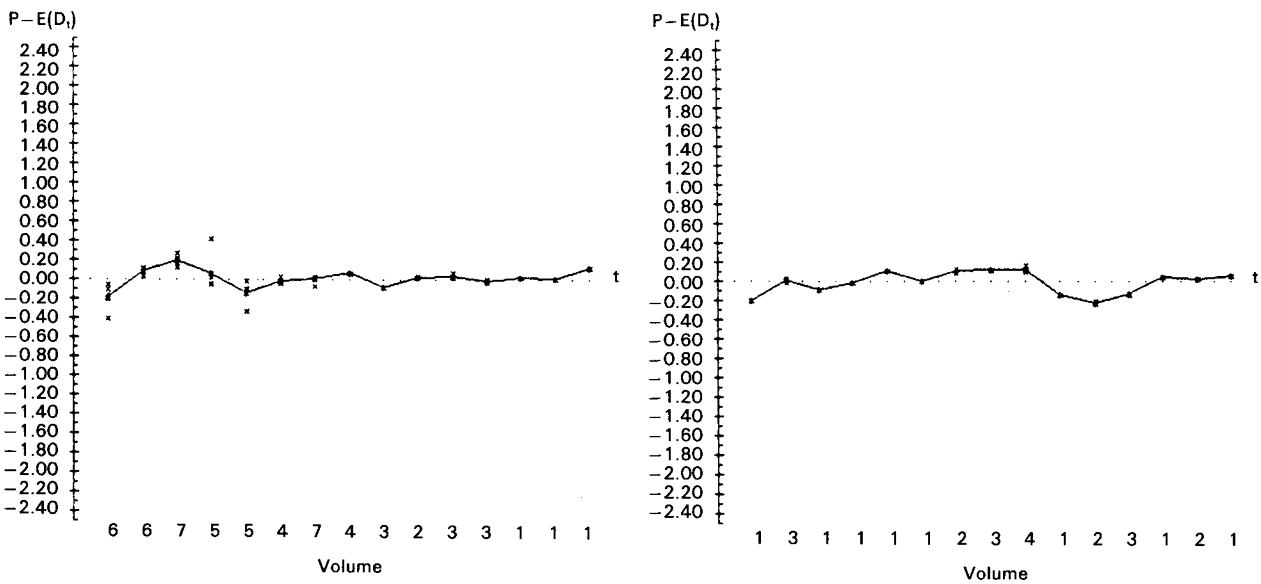

Finally, shortselling was not permitted in the experimental markets that we consider in this essay. A natural question is whether shortselling will stabilize prices near fundamental value earlier than is being observed, possibly eliminating the market bubbles altogether. Again, preliminary results from additional experiments wherein limited shortselling is permitted indicates that this is not the case, although shortselling does seem to reduce both the mean duration and the magnitude of price bubbles. This suggests that restricted shortselling is stabilizing, but will not eliminate price bubbles in general.

4. Concluding Remarks

In this section, we consider what we have learned from our experimental asset markets for evaluating the conventional wisdom about rational expectations and efficient markets, and for assessing the recent market experience and the alleged causes of the radical price declines that occurred recently in the world's securities' markets.

4.1 The Conventional Wisdom

The conventional wisdom holds that investors are rational where this is generally interpreted to mean that they will usually trade assets at prices near fundamental values, quickly developing common expectations given the same available information. Analysis of our experimental market data demonstrates quite clearly that agents are rational in the sense of the conventional wisdom only after they have acquired considerable experience with an unchanging environmental market structure. Endowing all agents with identical and full objective information about fundamental values will not induce rational expectations trading behavior. Thus, the rational expectations hypothesis is a valid equilibrium concept around which the profession can organize its theoretical studies, but it does not provide a good description of trading behavior in securities' markets where the trader sets and information structures are changing constantly. Such markets are continually in *disequilibrium,* and there is no well-developed received theory of disequilibrium trading behavior. Presumably, the efficient markets hypothesis voids the need for a disequilibrium theory, a notion that we evaluate next.

There are three forms of the efficient markets hypothesis, the strong, semi-strong and weak forms, and all are generally well-known (for a review, see FAMA, 1970). It is generally accepted that markets are not strong form efficient, and that markets are at least weak form efficient; the debate continues as to whether markets are semi-strong efficient. For our purposes, it suffices to remember that markets are efficient in the weak sense if price histories cannot be used to predict future prices, and markets are efficient in the semi-strong sense if no publicly available information can be used to predict future price movements (i. e. all public information must already be discounted in current prices). The semi-strong form is often interpreted to mean that, absent new unanticipated information, market prices will be a random walk about fundamental values. As we mentioned in the introduction, Chaos Theory is challenging this interpretation. Our results also challenge the random walk hypothesis because the trading prices that we observe in general in our laboratory markets exhibit pronounced regularities. What we learn about the efficient markets hypothesis is more subtle, however, and potentially of much greater importance.

Recall that the early experimental asset market research of FORSYTHE, PALFREY and PLOTT (1982) and others found that private information is quickly capitalized in the asset's market price. Although they were addressing the rational expectations hypothesis, this result provides some support for the strong form of the efficient markets hypothesis which the profession rejects since insider trading is known to yield extraordinary returns[4]. By way of contrast, no agent has any information not possessed by any other agent in our experimental markets. Consequently, we cannot say anything directly about strong form efficiency. We can use our results to evaluate semi-strong and weak form efficiency, however.

Based on our analysis, we endorse weak form efficiency. Simply put, an asset's price history does not provide enough information for agents to anticipate upper or lower turning points. On the other hand, we suspect that markets probably are not semi-strong efficient, although a relatively weak argument supporting semi-strong form efficiency can be made using our results. The difference is one of interpretation. We find that the number of excess bids is a reliable predictor of an asset's market price, and accurately presages both upper and lower turning points by at least one period. In our laboratory markets, the number of excess bids is public information since it is computed by subtracting the number of offers in a period from the number of bids in a period, and all bids and offers are known to all agents. Hence, information that would enable traders to predict the most likely direction of future price movements is available. In view of the definition, this suggests that the market is not semi-strong efficient.

It is worth noting that agents generally do not assimilate the information content of the number of excess bids. Agents systematically change the direction of their price forecasts only after observed prices have already changed direction. This reinforces our conclusion that our laboratory asset markets are not semi-strong efficient. Since our experimental market environment is simpler and the subjects much better informed than in natural market environments like the New York Stock Exchange, we conclude that actual securities' markets probably are not semi-strong efficient either, even though information about the number of excess bids is generally not available in such markets (except to the market specialist through the contents of his limit order book)[5]. And if markets are not semi-strong efficient, then they certainly are not strong efficient. Alternatively, one can argue that our experimental markets are *approximately* semi-strong efficient since agents react quickly to reversals in the direction of price movements. This skirts the intent of the semi-strong efficiency concept, however, and is not a very convincing counterargument.

4.2 The Recent Market Experience

When the Dow Jones Industrial Average plummeted 508 points on bloody Monday (October 19, 1987), climaxing the substantial decline occurring the previous week, investors were incredulous, angry and worried. According to received theory, the 22.5 percent decline was unwarranted. Efficient market theorists were perplexed, and many investors and politicians were clamoring for regulatory reform to prevent a recurrence. Billions of dollars had been lost, and fingers were pointed accusingly at the capitalization requirements of specialists and brokerage firms, at the failure to have more stringent margin requirements, at the growth of program trading, at the growth of index futures and options trading, and at the mechanics of trading itself, in particular at the increasing reliance on 'programmed' block trading strategies. Our studies of asset markets in the laboratory suggest that the price decline was to be expected, and that it had little or nothing to do with any of the reasons advanced above. It was simply a market 'correction', but one of panic proportions. These interpretations of our laboratory studies are corroborated by the results of survey evidence reported by Shiller (1987) that he conducted at the time of the October crash.

The price collapse was not unique to the US securities' markets, but was suffered by securities' markets worldwide. Yet these markets are quite diverse in their trading rules, capital requirements, and the use of electronic trading mechanisms. This suggests that it is futile to look for a cause of the price declines hidden in the mechanics of trading or in the institutional market structures, especially a cause that can be regulated easily to prevent a recurrence. A scramble to regulate *suspected* causes could actually do more harm than good, interfering with important normal market functions. Price bubbles and panics go back centuries in history before any of these new suspected causes were operative.

When assessing the recent record market decline, it is important to remember that the boom preceding it was also punctuated by unprecedented record price increases. This market pattern is essentially the same pattern we have been observing in our laboratory markets, and our laboratory markets are simple stylized market structures, lacking all of the factors inherent in actual securities' markets that many claim contributed to the recent market collapse. It is apparent to us at least that the similarity is fundamental.

People are more than happy to ride a market bubble as it develops, which according to our results often provides even greater impetus to the increasing prices. As market prices continue to climb, greater numbers of inexperienced investors tend to move into the market, and our results suggest that this will fuel further price increases and generate greater volatility. Experienced investors, however, will not continue to hold or to buy securities once prices have moved well above fundamenal values, realizing that such prices cannot be sustained. At some point, traders will begin selling so as to realize their capital gains, initiating a general price decline that will gather momentum as other investors revise their expectations downward and begin to sell as well. According to our results, this decline – sharp though it may be – represents a *market correction* where prices are dropping to near fundamental values. Granted, the decline may continue until prices are below fundamental values. If this occurs, then the situation is temporary, and a recovery will follow as our laboratory market experiments demonstrate. And so it goes, just as it has in a recurring pattern throughout the long history of stock markets.

To be sure, there probably are technical factors that will exacerbate bubbles and crashes. Halting trading in 91 stocks in New York, for example, while their options continue to trade in Chicago cannot help to moderate a crash, nor can it help to have futures trading suspended in Chicago while stock trading is suddenly resumed in New York (*Wall Street Journal,* October 26, 1987). The problem here is not institutional, but simply a lack of coordination among technically interdependent markets. Increased coordination of these markets requires more electronic trading to reduce uncertainty created by lags in prices and trades, not less as claimed by some. Finally, trading suspensions and limit price movements are probably counterproductive; once a bubble has developed, a market correction will most likely occur, and trading suspensions and limit price movements simply delay the inevitable. Recognition of this is implicit in the practice of increasing the limits on maximum daily price movements on the futures' exchanges as volatility increases. Preliminary evidence from laboratory investigations of limit price changes being conducted at the Economics Science Laboratory, University of Arizona, and at Indiana University provides additional support for this conclusion.

Particular mischief may be done in boom and bust markets by *at market* orders. At market orders provide for immediate execution at whatever price is available. In the short run in a crash market, for example, a flood of at market sell orders will temporarily yield incompatible demand and supply flows. There will be no clearing price if the mismatch cannot be accommodated from the specialist's limit order book, or from his own account. Consequently, trading either must be suspended until orders are better balanced (as in New York) or until the next day (as in Chicago) under a limit price change rule. If all orders had to specify a limit price, such indeterminate price episodes would not exist, and investors would be forced to think in terms of prices at which they personally are prepared to buy and sell. Increasing margin and capital requirements probably would be helpful in this regard. However, these changes would not be a panacea, preventing all sharp price movements, as our laboratory market experiments show. In our experimental markets, all bids and offers are price orders, and traders cannot borrow to invest.

Appendix: Estimates

Source: SMITH, SUCHANEK, WILLIAMS (1988): 'Bubbles, Crashes and Endogenous Expectations in Experimental Spot Asset Markets', *Econometrica,* forthcoming.

$$P(t) = 0.208 + 0.884\,F(t,i);\ R^2 = 0.823,\ N = 852. \quad (1)$$
$$\qquad (5.98)\quad (-8.25)$$

$$F(t) - P(t) = 0.046 + 0.282[F(t-1) - P(t-1)];\ R^2 = 0.089,\ N = 769. \quad (2)$$
$$\qquad (3.03)\quad (8.66)$$

$$F(t) - P(t) = -0.077 - 0.824[P(t) - P(t-1)];\ R^2 = 0.589,\ N = 781. \quad (3)$$
$$\qquad (-7.3)\quad (-33.4)$$

$$F(t) - F(t-1) = -0.117 + 0.815[P(t-1) - F(t-1)];\ R^2 = 0.632,\ N = 850. \quad (4)$$
$$\qquad (12.11)\quad (38.16)$$

$$P(t) - P(t-1) = -0.23 + 0.027\,X(t-1);\ R^2 = 0.24,\ N = 182. \quad (5)$$
$$\qquad (0.29)\quad (7.55)$$

$$P^c(t) - P^c(t-1) = -0.219 + 0.02\,X(t-1);\ R^2 = 0.126. \quad (5')$$
$$\qquad (0.55)\quad (5.09)$$

Footnotes

[1] See FORSYTHE, PALFREY and PLOTT (1982), FRIEDMAN, HARRISON and SALMON (1984), and PLOTT and SUNDER (1982).

[2] For a complete and technical development of the regression results presented in this essay, see SMITH, SUCHANEK and WILLIAMS (1988).

[3] NASH (1950) first defined rational expectations to mean that agents' expectations are realized. MUTH (1961) strengthened the definition to require that agents' expectations conform on average to the objective distribution. In our markets, this is given by D(t).

[4] This statement is not entirely fair. In the experimental markets of FORSYTHE, et al., there is no dividend infor-

mation that is private to a single trader; there are at least three traders of each type in every experiment. This dispersion of information and the level of competition that it induces may be sufficient to classify the information as *almost public*.

[5] This may well explain the specific nature of the market specialist's ability to earn extraordinary returns from the information contained in his limit order book. The orders maintained in the book should enable the specialist to estimate accurately the number of excess bids, which in turn (according to our results) would enable him to anticipate major price changes better than other investors who do not have this information.

References

ARROW, K. J., and L. HURWICZ (1962): 'Competitive Stability Under Weak Gross Substitutability: Nonlinear Price Adjustment and Adaptive Expectations', *International Economic Review* 3, 233–255.

BROCK, W. A. (1986): 'Distinguishing Random and Deterministic Systems', *Journal of Economic Theory,* October.

BROCK, W. A., and W. D. DECHERT (1986): 'Theorems on Distinguishing Deterministic and Random Systems', *Dynamic Econometric Modelling,* Proceedings of the Third Austin Symposium in Economics, W. BARNETT, E. BERNDT and H. WHITE, eds. Cambridge University Press, Cambridge.

BROCK, W. A., W. D. DECHERT and J. A. SCHEINKMAN (1987): *A Test for Independence Based on the Correlation Dimension,* Mimeo, University of Wisconsin-Madison.

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

FORSYTHE, R., T. PALFREY and C. PLOTT (1982): 'Asset Valuation in an Experimental Market', *Econometrica* 50, 537–567.

FRIEDMAN, D., G. HARRISON and J. SALMON (1984): 'The Informational Efficiency of Experimental Asset Markets', *Journal of Political Economy* 92, 349–408.

LUCAS, R. (1986): 'Adaptive Behavior and Economic Theory', *Journal of Business,* Supplement (October).

MUTH, J. F. (1961): 'Rational Expectations and the Theory of Price Movements', *Econometrica* 29, 315–355.

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

PLOTT, C. R., and S. SUNDER (1982): 'Efficiency of Experimental Security Markets with Insider Trading', *Journal of Political Economy* 90, 663–698.

SHILLER, ROBERT J. (1987): *Investor Behavior in the October 1987 Stock Market Crash: Survey Evidence,* Working Paper No. 2446, National Bureau of Economic Research, November.

SMITH, V. L., G. L. SUCHANEK and A. W. WILLIAMS (1988): 'Bubbles, Crashes and Endogenous Expectations in Experimental Spot Asset Markets', *Econometrica,* forthcoming.

WILLIAMS, A. W., and V. L. SMITH (1984): 'Cyclical Double-Auction Markets With and Without Speculators', *Journal of Business* 57, 1–33.