Explaining Bubbles in Experimental Asset Markets

Eric O'N. Fisher*
Department of Economics
The Ohio State University
410 Arps Hall
1945 North High Street
Columbus, Ohio 43210
(614) 292-3906
fisher.244@osu.edu

April 21, 1998

Abstract

Using experimental data from asset markets without short selling, this paper shows that bubbles arise for two reasons. First, subjects take time to learn about the dividends, not trusting initially the experiment's instructions. Second, agents have heterogeneous prior beliefs. These data support a Bayesian model of asset prices against the alternative hypothesis that agents have common priors agreeing with the actual distribution of fundamentals. This evidence confirms a subtle version of rational expectations in which assets are priced according to agents' priors and the market's history.

*The author would like to thank Jim Peck, Marine Carrasco, Nelson Mark, Joel Watson, Molly Cooper, and seminar participants at The Ohio State University, University of California at San Diego, University of Arizona, and the Midwest Macroeconomics Meetings at St. Louis for useful comments on earlier drafts of this work. The author will make available at cost a diskette containing all the data and programs used in the text.

1. Introduction

This paper brings together two hitherto disparate literatures. The first stems from the study of experimental asset markets by Smith, Suchanek and Williams (1988), and the second grows from the theoretical work of Harrison and Kreps (1978) and Morris (1996) on the role of heterogeneity of prior beliefs in theory of asset pricing. Smith, Suchanek and Williams showed that bubbles were ubiquitous in the laboratory, and their findings have spawned a large literature, surveyed well by Sunder (1995). Sunder conjectures that these bubbles may be artifacts of learning; he makes the distinction between bubbles and so-called false equilibria, in which some traders make incorrect inferences about the state of nature. The main contribution of my paper is to show that learning does indeed matter, but typical experimental data are consistent with a fully articulated model of equilibrium.

This paper draws upon the prescient analysis of Harrison and Kreps (1978), who first explored a model in which agents had heterogeneous expectations about an asset's dividends. Harrison and Kreps made an important but subtle distinction between common knowledge and an agent's own forecasts. In their analysis, it is common knowledge that agents have different expectations. In a market with no short selling, the natural definition of the market fundamental is the valuation of the most bullish trader, and the startling insight of Harrison and Kreps was that this was a *lower bound* for the asset's price. Often an asset trades strictly above this price, formalizing Keynes's notion of a beauty contest. The intuition for their surprising finding is that purchasing an asset entails both a claim to its future stream of dividends *and* the option value of reselling it later to another more bullish person.[1]

Morris (1996) extended Harrison and Kreps by incorporating learning. In a penetrating philosophical analysis of the common prior assumption in economics, Morris (1995) argues that economists, who have largely accepted Savage's notion of subjective probability, ought not be averse to analyzing the implications of diverse prior beliefs in markets. In discussing the practical implications of his own work, he observes that diverse priors may explain the pricing of initial public offerings. Morris (1996, p. 1128) states tellingly, "[Consider] the thought experiment of imagining traders forming priors about the dividends of an as yet unobserved asset. It is extremely hard to conceive of any criteria—rational or otherwise—that might require traders to have the same prior. A number of different priors seem (*sic*) entirely reasonable."

But this is precisely the situation with which subjects in asset market experiments are faced. Often, they come into an experiment never having traded assets with random dividends. Although the experiment's instructions always emphasize the true distribution of dividends and thus attempt to create as much common knowledge as possible, there is no simple way of knowing how a subject interprets this information. Think of the difference between experiments in psychology and economics. In empirical psychology, some aspect of the design is often concealed from the subjects. Therefore, anyone having been in a psychology experiment may have a healthy skepticism about the true trading

[1] Tirole (1982) noted that bubbles might arise if agents had different priors, but he concentrates instead on situations in which agents receive different private signals about an asset's dividends. He shows that there can be no bubbles in a market with finitely many periods, but the analysis below shows that bubbles arise even in these markets if agents have different priors.

environment. Thus, in spite of economists' best efforts, it is plausible that subjects are reluctant initially to believe that the distribution of dividends is transparent.

This observation is not a criticism of asset market experiments. Indeed, only experimental data allow one to make a proper test for a bubble. Prices from field data present the econometrician with the thorny problem of how to disentangle an asset's fundamental from an assumed dividend process and agents' information sets. Thus an obvious advantage studying an asset in the laboratory is that its fundamental is a part of the experimental design. In essence, experimental markets allow one to disentangle an otherwise intangible fundamental from a real measure of "irrational exuberance."

The model developed and calibrated below goes well beyond the norm in the applied literature on tests for bubbles. First, both the model and its empirical implementation take heterogeneity seriously; the theory is based upon the assumption that agents have different prior beliefs, a postulate supported in data from many of the experiments. Second, a simple version of rational expectations is strongly rejected in these (and many other) experimental data. Thus the workhorse of much of the literature in applied finance just does not seem adequate to the task of carrying the data generated by these experiments. Indeed, this paper indicates that more attention should be paid to models where learning and informational heterogeneity are central.

The rest of this paper is structured as follows. The next section reviews the experimental designs, and the third section develops the model of asset prices. The fourth section discusses the calibration method, and the fifth describes the empirical findings. The sixth section presents brief conclusions and suggests directions for future research.

2. Experimental Design

These experiments were designed to test simple theories of exchange rate determination described in Fisher and Kelly (1997). Using an oral double auction, the subjects traded blue or red assets for cash. An asset was a claim to a finite stream of dividends, and every experiment lasted for fifteen periods. The dividends were independently and identically distributed on simple supports. In some designs, the assets were different, and in others, the distinction was completely spurious. All this information was made quite explicit in the instructions.

Indeed, these experiments made as much information as possible public at all times. Thus the instructions described the dividend process at great length, and the history of realized dividends for each asset was written at the top of the blackboard. At the beginning of each period, the maximal, minimal, and expected values of the remaining stream of dividends for each asset were announced and written down on the blackboard. Of course, this emphasis was intended to make the true process for an asset's dividend the focal point of the agents' priors. Finally, it was emphasized repeatedly that the assets had no redemption value, standard for experiments of this type.

Since so much of the analysis depends upon agents' beliefs about the probabilistic natures of assets dividends, the Appendix reproduces the exact language used in describing the dividends in one design.[2] These experiments were run by hand, and bid and offer was written on a blackboard with the blue market on the left and the red market on the right. Since the administration of the experiments was so labor-intensive, time was at a premium, and a random number generator was used before beginning the session

[2] A complete set of instructions for Design 4 are given in Fisher and Kelly (1997), and I will make available instructions for any of the other designs.

to generate dividends. The dividends were written on a piece of paper that the experimenter brought into the classroom, and he publicly looked at it at the end of each period to announce what dividend had occurred in each market. Hence, the subjects did not see a tangible and obvious randomization device at the end of each six-minute period, but the dividends were drawn from precisely the distribution described to the subjects.[3]

Table 1 summarizes the designs for the 15 experiments. The agents were given heterogeneous initial mixes of money and assets, as is usual. For the purposes of this paper, the most important aspect of the design is the number of traders having previously participated in markets for assets with random dividends. *All* the subjects had had prior experience with oral double auctions, but only experienced traders had been in prior asset market experiments. The empirical analysis below shows that these subjects are more likely to have prior beliefs about the dividends that correspond with their actual distributions. At the beginning of each experiment, the experienced traders were publicly identified, breaking with some of the earlier experiments in this literature. Again, the intention was to have as much common knowledge as possible. Finally, it is worth emphasizing that these markets allow no short selling.[4] This aspect of the design allows one to build on the elegant theories of Harrison and Kreps (1978) and Morris (1996) to develop a model of asset prices that explains the data.

[3] Other authors typically run these kinds of experiments using computers, not by hand. In that case, the randomization device used to generate dividends is if anything more mysterious.

[4] King, Smith, Williams, and van Boening (1993) argue that bubbles in these experimental markets occur even when short selling is allowed and futures markets are introduced.

TABLE 1: Experimental Parameters

Experiment (Experienced Traders)	Trader Types (Number of Traders)	Endowments	Description of the Dividends
Design 1 11 November 1993 (0) 12 November 1993 (0) 11 January 1994 (0)[5]	I (2) II (2) III (2) IV (2) V (2) VI (2)	3 blue assets and \$2.25 2 blue assets and \$5.85 1 blue asset and \$9.45 3 red assets and \$2.25 2 red assets and \$5.85 1 red asset and \$9.45	Blue: uniform on {.6, .28, .08, 0} Red: uniform on {.6, .28, .08, 0} The assets' dividends are independently distributed.
Design 2 15 April 1994 (7)	I (2) II (2) III (1) IV (2) V (2) VI (1)	3 blue assets and \$2.25 2 blue assets and \$5.85 1 blue asset and \$9.45 3 red assets and \$7.65 2 red assets and \$9.45 1 red asset and \$11.25	Blue: uniform on {.6, .28, .08, 0} Red: uniform on {.3, .14, .04, 0} The assets' dividends are perfectly positively correlated.
Design 3 12 January 1994 (0) 28 January 1994 (0) 11 February 1994 (0)	Same as Design 1	Same as Design 2	Same as Design 2, but the assets' dividends are independently distributed.
Design 4 22 February 1994 (1) 1 March 1994 (5) 4 March 1994 (4)	Same as Design 1	3 blue assets and \$2.75 2 blue assets and \$6.50 1 blue asset and \$10.25 3 red assets and \$2.75 2 red assets and \$6.50 1 red asset and \$10.25	Blue: uniform on {.5, 0} Red: uniform on {.3, .2} The assets' dividends are independently distributed.
Design 5 6 May 1994 (7) 20 May 1994 (9)	Same as Design 1	Same as Design 4	Same as Design 4, but the traders are asked to forecast the exchange rate.
Design 6 2 March 1994 (0) 7 March 1994 (11) 9 March 1994 (11)	I (2) II (2) III (2) IV (2) V (2) VI (1)	Same as Design 1	Same as Design 1

3. The Model of Asset Prices

An asset is a claim to a stream of dividends that expires after T periods. Dividends are paid at the end of each period, and assets are traded at a price that includes a claim to the current dividend. The dividend is a binomial or multinomial random variable parameterized by $q = (q_1, ..., q_k)$, with $q_j \geq 0$ and $\sum_{j=1}^{k} q_j = 1$. Since the initial

[5] There was only one Type II trader and one Type VI Trader on 11 January 1994.

experimental designs were based upon those of Smith, Suchanek, and Williams, each asset had four possible dividends and $k = 4$. Some later experiments used a simpler design where $k = 2$. Thus q_j is the probability that a type *j* event occurs, in which case one receives d_j for each asset held.

Agent *i*'s priors are a distribution $p_i(x)$ defined on Δ^k, the simplex in $\Re^k$. For example, an agent who is certain that a binomial dividend is a fair coin toss has priors that are completely concentrated on (.5, .5), but one who is agnostic might have diffuse priors that are uniformly distributed on the simplex. Either agent has the same initial expected value for the dividend, but subsequent valuation of the agent with diffuse priors will depend upon the stochastic history that arises.

Let $0 \le t \le T$. Then a *t*-period history is a list of realized dividend $h^t = (d_1,...,d_t)$; the agents have observed only the null history h^0 in the first period. Let $n(h^t) = (n_1(h^t),...,n_k(h^t))$ be the sufficient statistic necessary to estimate q.[6] Then the agent's posterior after h^t is

$$j_i(q \mid h^t) = \frac{(q_1)^{n_1} \cdots (q_k)^{n_k} p_i(q)}{\int (x_1)^{n_1} \cdots (x_k)^{n_k} p_i(x) dx},$$

where this integral is taken over the relevant simplex and I have suppressed the dependence of n_j on h^t for notational convenience. Let $x = (x_1,...,x_k)$, and let $f : \Delta^k \to \Re$ have the rule:

[6] Note that $n(h^0) = (0,...,0)$.

$$f(x) = \begin{cases} d_j & if \;\; x = (0,\ldots,q_j,\ldots,0) \\ 0 & otherwise, \end{cases}$$

where again d_j is the value of the *j*-th dividend. Then agent *i*'s expectations about an asset's dividend after history h^t are summarized by:

$$m_i(h^t) = \int\int j_i(q \mid h^t)\, f(x) dq dx\,, \tag{1}$$

where both integrals are taken over the relevant simplices. The two sums in (1) have elegant interpretations. The inner integral is taken with respect to the agent's prior beliefs, and the outer one corresponds to the possible outcomes at the end of the *t*-th period. Since the dividend is a discrete random variable, (1) implies

$$m_i(h^t) = \sum_{j=1}^{k} d_j \frac{\int (q_1)^{n_1} \cdots (q_j)^{n_j+1} \cdots (q_k)^{n_k} p_i(q) dq}{\int (q_1)^{n_1} \cdots (q_k)^{n_k} p_i(q) dq},$$

a useful formula in calibrations of the model.

Of course, an asset's price depends upon the history of realized dividends. The notation $p(h^t)$ makes this fact explicit. It is natural to impose that

$$p(h^T) = 0\,. \tag{2}$$

Thus an asset has no value at the end of the experiment, once all the dividends have been realized. This terminal condition is inherent in the experimental design.

Let $g : \Delta^k \to \Re$ have the rule

$$g(x) = \begin{cases} p(n_1,...,n_j+1,...,n_k) & if \;\; x = (0,...,q_j,...,0) \\ 0 & otherwise. \end{cases}$$

where $p(n_1,...,n_j+1,...,n_k)$ is the price that occurs when h^{t+1} is such that $n(h^{t+1}) = (n_1(h^t),...,n_j(h^t)+1,...,n_k(h^t))$. This price will prevail in period *t+1* if the *j*-th dividend is paid at the end of period *t*, given that the history h^t has been observed. Since there is no short selling, prices satisfy the recursion:

$$p(h^t) = \max_{i \in I} \left\{ \int\int j_i(q \mid h^t)[f(x)+g(x)]dqdx \right\}, \tag{3}$$

where the maximum is taken over the set of traders types *I*. Equation (3) is a generalization of the formula that Morris (1996) developed for an asset whose dividend was a simple binomial process. Of course, no discount factor enters because each trading period is no longer than six minutes!

The great advantage of a market with finitely many periods is that prices can be derived using backward induction from the terminal condition (2). Consider $\max_{i \in I}\{m_i(h^t)\}$, and let *i* be one of the agents placing the highest expected valuation on the dividend. Then *i* is a most *bullish agent after history* h^t. Using backward induction from (2), one can show that an asset's price is no less than the fundamental of the most bullish agent. Also, it strictly exceeds that valuation if agents' priors satisfy what Morris calls a switching condition. This situation occurs when priors are such that there is a subsequent information node at which there is another most bullish agent. In that case, buying an asset is the claim to a stream of dividends *and* the option value or reselling at a later history, so its price exceeds the natural definition of its fundamental. Thus Morris's arguments extend to the more general dividends analyzed here and to the case where assets have finite lives.

The model just developed applies *mutatis mutandis* to experiments in which several assets are traded simultaneously. Still, it is not possible to identify agents' priors from market prices, and it is appropriate to treat each market independently in the calibration. Of course, it is standard in Bayesian analysis to assume that agents' priors can be decomposed in simple ways; Smith (1988, chapter 6) gives some illuminating analysis and elegant examples. Also, treating agents' priors for the blue and red asset markets independently effectively doubles the amount of data available to test the theory. Thus one can use market prices for the blue and red assets to infer which priors about the relevant dividends that some (unidentified) agents must have held. The empirical implications of the theory will depend upon whether simple representations of these prior beliefs explain much of the observed data.

4. Calibration Techniques

Let p be the vector of an asset's actual prices. In these data, an asset's price is the average of the several transactions during a period in one market; if there is no transaction, then the price is the midpoint of the closing bid-ask spread.[7] Although there were often several transactions in a period, the information partition at the heart of the theory developed in Section 3 is only as fine as the frequency at which dividends are drawn. Thus nothing would be gained by analyzing every transaction price, and it is natural to give equal weights to transactions within a period.

The calibration technique searches for a list of priors $p = (p_i(x))_{i \in I}$ minimizing the distance between the data and the model's predictions. Let Π be a (computationally tractable) set of priors. Then a plausible list of priors is given by the solution to

$$\min_{p \in \Pi} (\hat{p}(p, h^T) - p)'(\hat{p}(p, h^T) - p)/T \,, \tag{4}$$

where $\hat{p}(p, h^T) = (\hat{p}_1(p, h^0), \ldots, \hat{p}_T(p, h^{T-1}))$ is the vector of the model's predicted prices. The notation makes explicit that these predictions depend upon all the agents' priors and the actual terminal history. A general solution to (4) is inherently difficult, since the space of distribution functions has infinitely many dimensions. Also, although the data consist of only *T* observations, it is obvious that an asset's price at any information node depends upon *all* the nodes that could be reached. For example, when there are four dividends, it is necessary to make 3060 calculations for each agent in order to compute (1) and then use the recursion (3) to predict an asset's price along the actual path h^T.

[7] On 9 March 1994, the subjects unanimously elected to close the markets early on occasion. Thus there were several periods with no transactions and no closing bid-ask spread. The theory has a further practical advantage: it prices assets at *any* information node. Thus one can still use the data from experiments in which there were fewer than fifteen prices, with the distance measure in (4) adjusted appropriately.

Thus it is natural to make some simplifying assumption about the agents and their prior beliefs. First, the index set of agents is limited to the three simplest cases. Thus:

Assumption 1: *I={1}*, *I={1,2}*, or *I={1,2,3}*.

The case with only one representative agent is an important benchmark. It does *not* represent the usual assumption of a competitive market with simple rational expectations. Instead, it captures the notion of learning about the true dividend process, a distinction that can be made only with experimental data. The case with two agents allows for both learning and heterogeneity of priors, the characteristic giving rise to the option value of reselling an asset. The case with three agents allows for increased diversity.

The simplest way to make the space of priors tractable is to restrict the domains upon which these functions are defined. An agent's priors are a (measurable) function $p_i : \Delta^k \to \Re$ such that $\int_{\Delta^k} p_i(x)dx = 1$, and the model has been developed at this level of generality. A natural empirical simplification is to choose a discrete domain $X \subset \Delta^k$ upon which the agents' prior beliefs are defined. Thus:

Assumption 2: If k=2, $X = \{(.01,.99),(.5,.5),(.99,.01)\}$, and if k=4,

$X = \{(.01,.01,.01,.97),(.01,.01,.97,.01),(.01,.97,.01,.01),(.97,.01,.01,.01),(.25,.25,.25,.25)\}$.

Figures 1(a) and 1(b) give a graphical representation of these priors as barycentric coordinates in the relevant simplexes. Because of the way that the dividends are defined in the empirical analysis, priors concentrated on {(.99, .01)} or {(.97, .01, .01, .01)} are bullish. Likewise those concentrated on {(.01, .99)} or {(.01, .01, .01, .97)} are bearish. Also, priors concentrated on {(.5, .5)} or {(.25, .25, .25, .25)} accord with simple rational expectations, a special case of the more general model that allows for learning. The points in these domains are in the interior of the simplex because (1) is not well defined if the priors are concentrated on the event that the asset never pays a positive dividend. Finally, note that these domains satisfy the "kernel-of-truth" assumption: they generically place positive weight on the true dividend process.

Figure 1(a): Support for the Priors with Two Dividends

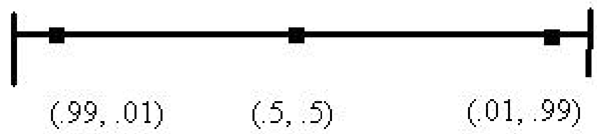

Figure 1(b): Support for the Priors with Four Dividends

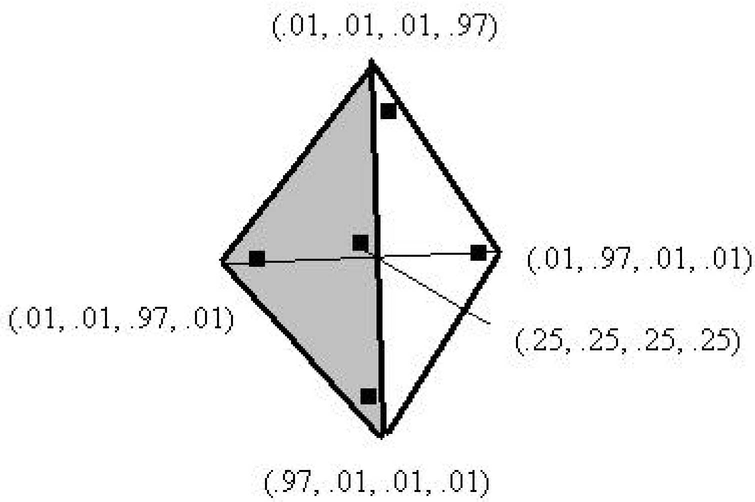

By imposing that each agent's priors are defined on a domain with finitely many elements, the problem has been reduced to a usual non-linear minimization. Indeed, it is necessary only to calibrate $\#I \times (\#X - 1)$ parameters, a computationally tractable task on a PC with a Pentium microprocessor. The empirical implementation uses the following procedure. First, Monte Carlo methods draw 100 random priors and save the analog of the minimum distance estimator in (4). Second, that prior is used as the starting point for a non-linear minimization routine, finding the best fit for a market with one agent.[8] Third, constraining the first agent's priors to be those found in the second step, the numerical minimization routine is run again to calibrate the priors for the second agent; now the starting point is that both agents have identical priors. Fourth, constraining two agents to have priors that are the outcome of third step, one searches for the third agent's priors; now the calibrated priors for the second agent are the initial condition for the third agent's priors. This procedure works well, but it depends upon the initial conditions.

This algorithm nests models in a natural way. In particular, the simple rational expectations model is a special case of the second step because imposing that a single agent's priors are concentrated on the truth generates the rational expectations prices, no matter what history is observed. Likewise, the third step is more general than learning because one can always impose that both agents' priors are identical to those generated in the second step, which were the best fit in a model with Bayesian updating. Hence the third step captures the importance of heterogeneity in prior beliefs. Finally, the fourth step allows for increased heterogeneity, and perhaps a higher likelihood that Morris's switching condition will obtain.

5. Empirical Results

The calibrated solutions to (4) are given in Table 2. The first column gives the predictions of the simple rational expectations model, the workhorse in applied economics and finance. The figures in this column serve as standards against which the alternatives are judged. The units are squared dollars per period.

---

[8] The "amoeba" method described in Press, Flannery, Teuklsky, and Vetterling (1988, pp. 292-93) is used. This method does not need derivatives, and it seems to behave well at the boundary of the parameter space.

Table 2: The Distance Measures

Experiment and Asset Market	Rational Expectations	One Agent	Two Agents	Three Agents
11 November 1993, blue	1.29	0.53	0.50	0.50
11 November 1993, red	1.23	0.47	0.47	0.47
12 November 1993, blue	0.81	0.56	0.49	0.49
12 November 1993, red	0.73	0.64	0.64	0.64
11 January 1994, blue	3.05	0.88	0.88	0.88
11 January 1994, red	3.06	1.11	1.11	1.11
15 April 1994, blue	0.92	0.61	0.61	0.61
15 April 1994, red	0.54	0.06	0.03	0.03
12 January 1994, blue	1.73	0.88	0.71	0.71
12 January 1994, red	0.37	0.36	0.16	0.16
28 January 1994, blue	2.36	1.41	0.86	0.86
28 January 1994, red	0.61	0.37	0.18	0.18
11 February 1994, blue	1.68	0.42	0.42	0.42
11 February 1994, red	0.36	0.11	0.11	0.11
22 February 1994, blue	0.95	0.74	0.74	0.74
22 February 1994, red	1.01	0.93	0.93	0.93
1 March 1994, blue	2.52	1.15	1.08	0.54
1 March 1994, red	1.60	1.57	1.57	1.57
4 March 1994, blue	0.97	0.75	0.75	0.75
4 March 1994, red	0.51	0.42	0.42	0.42
6 May 1994, blue	0.12	0.10	0.10	0.10
6 May 1994, red	0.06	0.02	0.02	0.02
20 May 1994, blue	0.35	0.29	0.27	0.27
20 May 1994, red	0.32	0.11	0.11	0.11
2 March 1994, blue	2.90	0.96	0.59	0.59
2 March 1994, red	3.01	1.11	1.03	1.03
7 March 1994, blue	0.43	0.16	0.12	0.11
7 March 1994, red	0.99	0.84	0.35	0.35
9 March 1994, blue	0.19	0.06	0.06	0.06
9 March 1994, red	0.17	0.13	0.13	0.13

The main message emerging from Table 2 is that learning matters in these markets. Compare the distances for the simple rational expectations model with those in the column labeled 'One Agent'. The model of Bayesian learning improves the predictive power of the model enormously in all but a handful of experiments. Graphs of the data show that the model predicts prices well in the first several periods; after about the fifth period, the simple rational expectations model and the model with learning often, but not always, coincide. An important conclusion is thus that the subjects have pessimistic priors; they do not trust fully that the dividend process described to them in the experiment's instructions is true.

A good example is the blue asset market on 11 November 1993. The simple rational expectations model is wrong on average by 1.29 squared dollars per period. The model with one representative agent who learns is wrong by only 0.53 squared dollars per period. The calibrated priors are given in Table 3. They show that the typical agent is almost sure at the outset that the dividend will pay only $.08, but he allows for some slight chance that the true dividend process might occur. The complete predictions of this version of the model are shown in Figure 2. The actual realizations of the dividend are shown in parenthesis along the bottom axis, and it is apparent that the agent learns the true process by about the fifth period. Still, the fact that his calibrated priors are so pessimistic allows for a much better fit during the initial periods, when the prediction of the simple rational expectations model are egregious. In this version of the model, Morris's switching criterion cannot be satisfied, since there is no second agent who might become more bullish after some future history.

Table 3: 11 November 1993,
Priors for the Blue Asset

Element of the Support	$p_1(x)$
(.01,.01,.01,.97)	0
(.01,.01,.97,.01)	0.9999
(.01,.97,.01,.01)	0
(.97,.01,.01,.01)	0
(.25,.25,.25,.25)	0.0001

Figure 2: 11 November 1993
Blue Market, One Trader Type

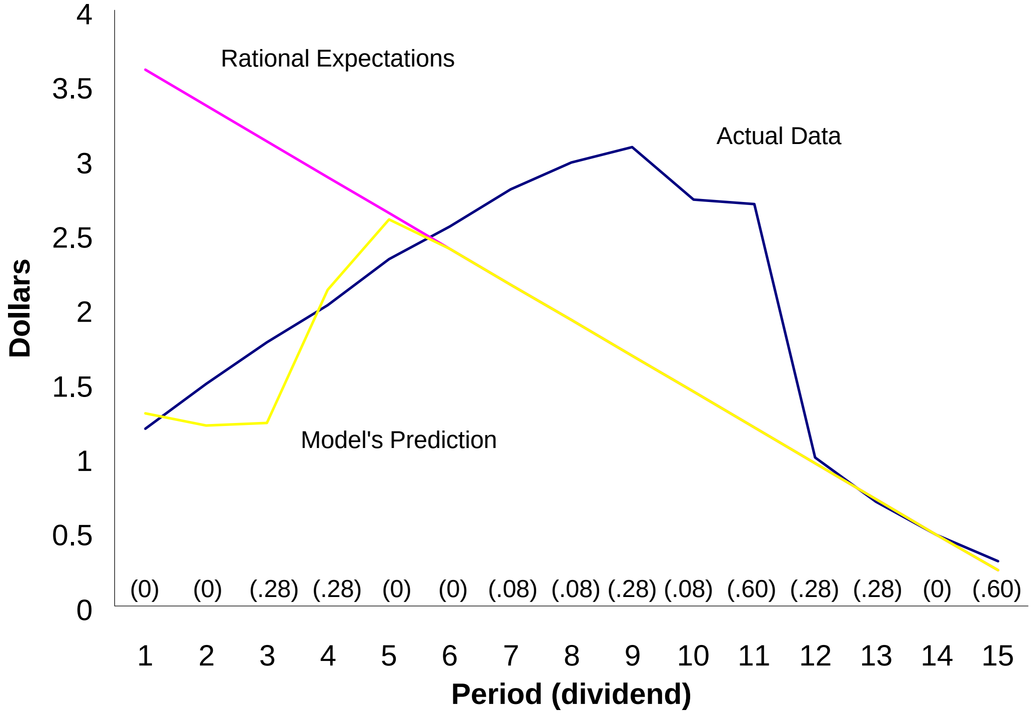

A second, and not unimportant, message in Table 2 is that the diversity of prior beliefs also matters. In about one quarter of the experiments, the distances in the column labeled 'Two Agents' are significant improvements over that in the model that only allows Bayesian learning. Of course, only a model with diverse priors can capture a second important aspect of these data: that prices often trade well above the remaining stream of dividends in the middle periods.

The red market on 28 January 1994 is a good example of how heterogeneous priors increase the model's predictive power. Table 4 gives the calibrated priors.

Table 4: 28 January 1994,
Priors for the Red Asset

Element of the Support	$p_1(x)$	$p_2(x)$
(.01,.01,.01,.97)	0.80	0.47
(.01,.01,.97,.01)	0.10	0.52
(.01,.97,.01,.01)	0	0
(.97,.01,.01,.01)	0.05	0.01
(.25,.25,.25,.25)	0.05	0

Both agents are bearish at the null history, but the first agent places a small probability on a bullish outcome and some on the true process. Figure 3 plots the data for this version of the model. It predicts the bubble remarkably well, and it is obvious that the two runs of good dividends, starting in the second and the eighth periods, raise the asset price.

Figure 3: 28 January 1994
Red Market, Two Trader Types

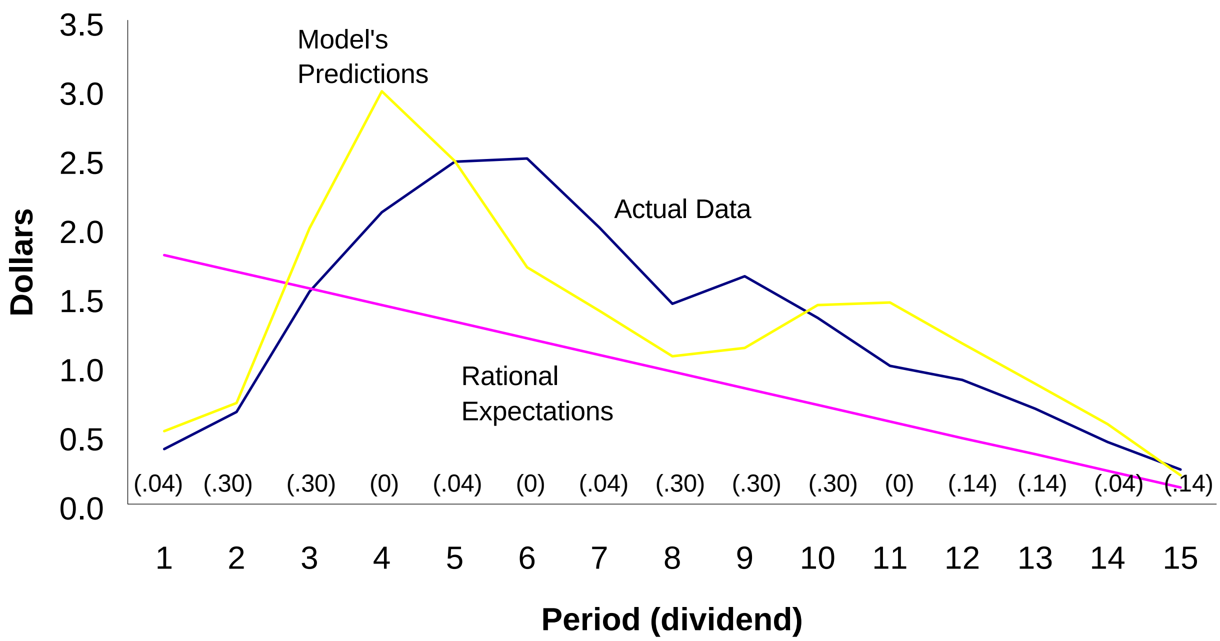

Tables 5(a) and 5(b) show how the posteriors evolved in this experiment. The first trader type learns the dividend process after four or five periods, but the second trader type never gets enough information to ascertain the true process. Also, the predicted priors satisfy Morris's switching condition. Thus the first agent holds the assets during the early part of the experiment, but the second agent becomes bullish after an early run of good dividends. The low dividends in the middle periods then turn him bearish, but the second run of good luck turns him bullish again, and he buys back the assets that he sold during the middle periods!

Table 5(a): 28 January 1994,
Red Asset, First Trader Type

$\hat{q}_1$ (\$.30)	$\hat{q}_2$ (\$.14)	$\hat{q}_3$ (\$.04)	$\hat{q}_4$ (\$.00)	History	Expected Dividend
0.07	0.02	0.12	0.79	∅	\$0.03
0.04	0.03	0.83	0.10	\$.04	\$0.05
0.27	0.17	0.38	0.19	\$.30	\$0.12
0.51	0.16	0.17	0.16	\$.30	\$0.18
0.27	0.24	0.24	0.25	\$.00	\$0.12
0.25	0.25	0.25	0.25	\$.04	\$0.12
0.25	0.25	0.25	0.25	\$.00	\$0.12
0.25	0.25	0.25	0.25	\$.04	\$0.12
0.25	0.25	0.25	0.25	\$.30	\$0.12
0.25	0.25	0.25	0.25	\$.30	\$0.12
0.25	0.25	0.25	0.25	\$.30	\$0.12
0.25	0.25	0.25	0.25	\$.00	\$0.12
0.25	0.25	0.25	0.25	\$.14	\$0.12
0.25	0.25	0.25	0.25	\$.14	\$0.12
0.25	0.25	0.25	0.25	\$.04	\$0.12

Table 5(b): 28 January 1994,
Red Asset, Second Trader Type

$\hat{q}_1$ (\$.30)	$\hat{q}_2$ (\$.14)	$\hat{q}_3$ (\$.04)	$\hat{q}_4$ (\$.00)	History	Expected Dividend
0.02	0.01	0.51	0.46	∅	\$0.03
0.01	0.01	0.96	0.02	\$.04	\$0.04
0.03	0.01	0.94	0.02	\$.30	\$0.05
0.65	0.01	0.32	0.01	\$.30	\$0.21
0.51	0.01	0.25	0.23	\$.00	\$0.16
0.03	0.01	0.94	0.02	\$.04	\$0.05
0.02	0.01	0.51	0.46	\$.00	\$0.03
0.01	0.01	0.96	0.02	\$.04	\$0.04
0.03	0.01	0.94	0.02	\$.30	\$0.05
0.65	0.01	0.32	0.01	\$.30	\$0.21
0.97	0.01	0.01	0.01	\$.30	\$0.29
0.96	0.01	0.01	0.01	\$.00	\$0.29
0.96	0.01	0.01	0.01	\$.14	\$0.29
0.96	0.01	0.01	0.01	\$.14	\$0.29
0.65	0.01	0.32	0.01	\$.04	\$0.21

It is not clear how many periods it would take for him to learn the true dividend process, but the model predicts that the asset's prices are well above the first agent's fundamental in the middle periods. The option value of reselling it to the second agent is significant indeed.

A third, and perhaps minor, message from Table 2 is that the simple rational expectations model occasionally predicts well. The blue market on 6 May 1994 is a good example; more than half the subjects had had experience in previous asset markets experiments. The calibrated priors for a representative agent are in Table 6.

Table 6: 6 May 1994,
Priors for the Blue Asset

Element of the Support	$p_1(x)$
(.01,.99)	0
(.5,.5)	0.84
(.99,.01)	0.16

The agent already places high prior probability on the true distribution of dividends, and he is slightly more bullish than is warranted, allowing the model to fit the data from the first period. Figure 4 shows that there was barely any bubble, but the model calibrates very plausible priors.

Figure 4: 6 May 1994
Blue Market, One Trader Type

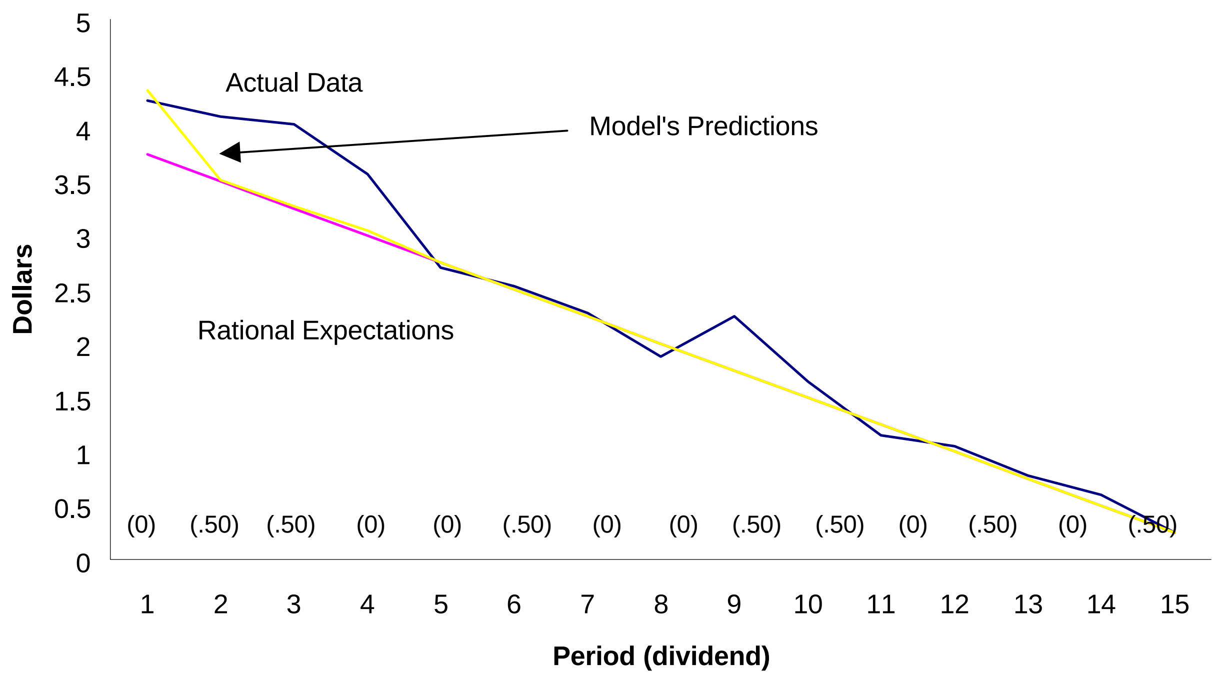

A fourth message from Table 2 is that only very rarely does one need a high degree of heterogeneity of priors to fit the data well. Only for the blue market on 1 March 1994 does the addition of a third agent make much of a difference. The calibrated priors are in Table 7. Figure 5 is again quite telling. The two strings of good runs of dividends seem to cause the bubbles, and it seems as though the addition of the third agent allows the model to capture the second small bubble occurring in periods 9 and 10.

Table 7: 1 March 1994,
Priors for the Blue Asset

Element of the Support	$p_1(x)$	$p_2(x)$	$p_3(x)$
(.01,.99)	.8123	.985	.817
(.5,.5)	.0007	.015	0
(.99,.01)	.187	0	0.183

Figure 5: 1 March 1994
Blue Market, Three Trader Types

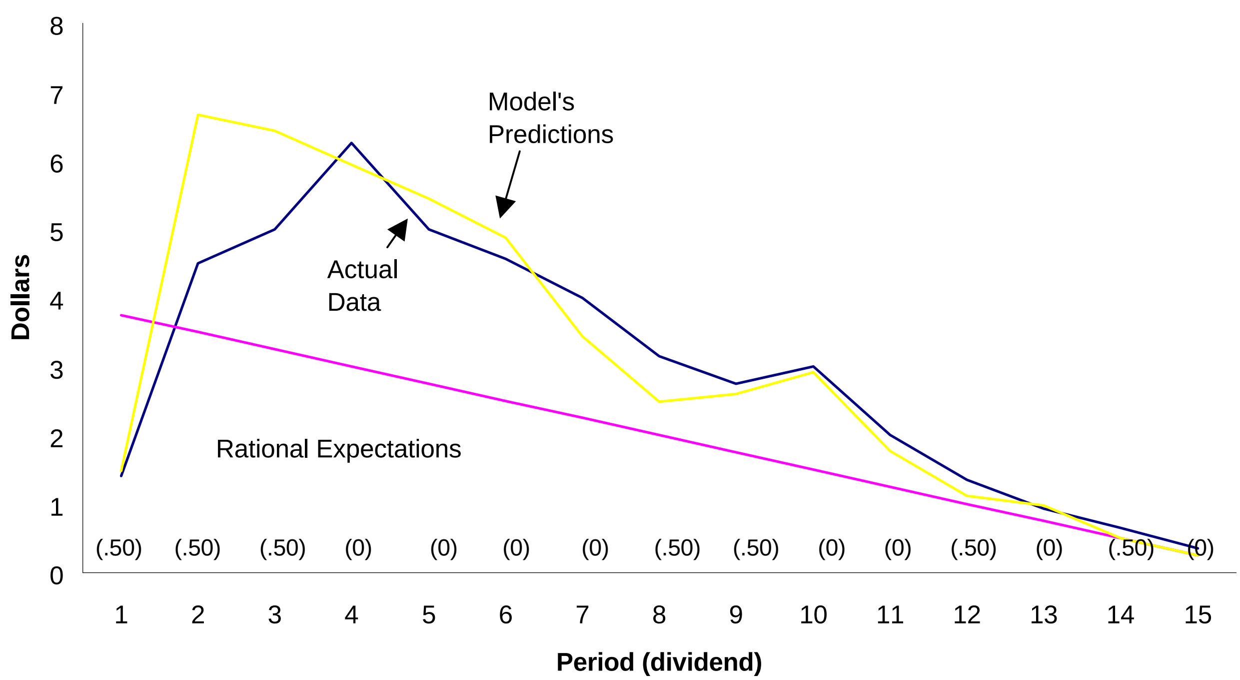

Now consider the predictive power of the model across experiments. The data fall naturally into two groups. The first, consisting of the experiments from Design 1, Design 3, 22 February 1994, and 2 March 1994, is likely to have agents with diverse beliefs about the asset dividends. The second, comprising the experiments from Design 2, Design 5, 1 March 1994, 4 March 1994, 7 March 1994, and 9 March 1994, has experienced traders, whose beliefs about the dividends are more likely to be homogeneous. The simple rational expectations predictions are the benchmark against which any more general model must compete.

The simplest model with learning and heterogeneity is described the column entitled "Two Agents" in Table 2. The *model improvement* is the decrease in the distances between the column called "Rational Expectations" and the more general model's prediction. The median model improvement is 0.31. Table 8 gives the contingency cells for a test based upon the null hypothesis that the model improvement is equal across the two groups.[9]

Table 8: Equal Model Improvement

	Group I	Group II
Above Median	11	5
Below Median	5	9

The test statistic is $\chi^2(1) = 4.82$, whose p-value is greater than .95. Thus we reject the null hypothesis and may conclude that the model with learning and heterogeneous priors is better at explaining the experiments where there are fewer experienced traders.

6. Conclusion

The typical bubble in experimental data has three parts. First, the asset trades below the expected value of its future stream of dividends. Second, there is an intermediate stage when it trades above this valuation. Third, the asset price crashes down to its fundamental value, often about two-thirds of the way through the experiment.

[9] See Sprent, (1989, p. 86).

The theory developed by Harrison and Kreps (1978) and Morris (1996) shows how to explain this canonical pattern in the data. During the early stages of the experiment, inexperienced subjects simply do not believe that the dividend's assets are as described in the instructions. Inexperienced subjects almost uniformly have pessimistic priors. Still, some may attach a slight probability to a more bullish event, and then a run of a couple good dividends can generate the second stage of a bubble. It is surprising that it takes only two different agents to generate realistic bubbles, but Morris's switching condition allows for an asset's price to be well above its fundamental for many periods. The third stage, when agents learn the true dividend process, often occurs earlier in these calibrations than in the data themselves. Still, there are experiments in which the calibrations are remarkably accurate for the entire run of the experiment.

Perhaps the model predicts too much learning. Of course, few people can calculate Bayes's rule while participating actively in an oral double auction, and it may be plausible that subjects use simpler rules of thumb for updating their beliefs. Timmerman (1997) explores other models of learning in asset markets and argues that non-stationary dividend processes can have long-lasting effects on asset markets.

The model fits especially well when an asset has a run of good luck. The calibrations indicate that small prior probabilities on extreme events can generate significant bubbles in that case. These calibrations based on data from experiments may have important implications for how asset markets behave in the field. Further empirical applications of models where informational heterogeneity matters and learning is explicit seem a promising area of future research.

References

Fisher, E. O'N. and F. S. Kelly. "Experimental Foreign Exchange Markets." Unpublished manuscript, The Ohio State University, 1997.

Harrison, J. M. and D. M. Kreps. "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations." *Quarterly Journal of Economics 92* (1978), 323-336.

King, R., V. Smith, A. Williams, and M. van Boening. "The Robustness of Bubbles and Crashes in Experimental Spot Markets" in I. Prigogine, R. H. Day, and P. Chen, eds., *Nonlinear Dynamics and Evolutionary Economics*, Oxford: Oxford University Press, 1993.

Morris, S. "The Common Prior Assumption in Economic Theory." *Economics and Philosophy* 11 (1995), 227-253.

__________. "Speculative Investor Behavior and Learning." *Quarterly Journal of Economics* 111 (1996), 1111-1134

Press, W. H., B. P. Flannery. S. A. Teuklsky, and W. T. Vetterling. *Numerical Recipes; The Art of Scientific Computing*. Cambridge: Cambridge University Press, 1988.

Smith, J. Q. *Decision Analysis: A Bayesian Approach*. London: Chapman and Hall, 1988.

Smith, V., G. L. Suchanek, and A. W. Williams. "Bubbles Crashes and Endogenous Expectations in Experimental Asset Markets." *Econometrica* 56 (1988), 1085-1118.

Sprent, P. *Applied Nonparametric Statistical Methods*. London: Chapman and Hall, 1989.

Sunder, S. "Experimental Asset Markets: A Survey" in *The Handbook of Experimental Economics* (J. H. Kagel and A. E. Roth, eds.). Princeton: Princeton University Press, 1995.

Timmerman, A. "Asset Prices on Learning Paths." Unpublished manuscript, University of California at San Diego, 1997.

Tirole, J. "On the Possibility of Speculation Under Rational Expectations." *Econometrica* 50 (1982), 1163-1182.

Appendix: Excerpts from the Instructions for Design 4

Some Information on Probability

The dividends are drawn randomly each period. The possible values of the dividend per unit and the associated probabilities of occurrence are given below:

Blue Dividends

dividend	\$0.00	\$0.50
probability	1/2	1/2

Red Dividends

dividend	\$0.20	\$0.30
probability	1/2	1/2

Thus, the average blue dividend is \$0.25. Also, the average red dividend is \$0.25. The blue dividends and the red dividends are independent of one another. This means that if the blue asset pays a high dividend in period 1, it tells you nothing about what the red dividend will be in the same period.

Think of the blue dividend as being generated by a flip of a coin. If heads comes up, the blue dividend is \$0.00. If tails comes up, then the blue dividend is \$0.50. Similarly, if I flip another coin and heads comes up, then the red dividend is \$0.20. Or if tails comes up with this other coin, then the red dividend is \$0.30.

Notice that the blue asset earns the same on average as the red one does. However, the blue asset is much riskier than the red one. Holding a blue asset may earn you \$0.50 per period, but you may be stung and earn nothing in any period. On the other hand, holding a red asset earns you on average \$0.25 per period, and you will always get at least \$0.20 per period.

It is important to recognize that past dividend draws do not affect the next dividend draw. The dividend draws are independent of one another. Over many draws, the average or "expected" dividend is \$0.25 per blue unit and \$0.25 per red unit. Therefore, if you were to hold your inventory endowment of one blue unit over the course of the 15 trading periods you would expect to earn dividend payments in the amount of 15 x \$0.25 = \$3.75. If you were to hold your inventory endowment of one red unit over the course of 15 periods, then you would expect to earn dividends in the amount of 15 x \$0.25 = \$3.75. Again, the blue asset is riskier than the red one. You can earn a lot of money during fifteen periods by holding a blue asset, but you can also have a long string of bad luck. On the other hand, the red asset is safer. You will never get burned because you can count one earning at least \$.20 in each period, and you will earn \$.30 about half the time. Holding both assets for all fifteen periods yields an expectation of \$0.50 at the end of each trading period. Of course, you will be able to change your end-of-period inventory holdings through buying and selling assets of both colors over the course of each trading period.