Article

Models of Investor Forecasting Behavior — Experimental Evidence

Federico Bonetto [1], Vinod Cheriyan [2]* and Anton Kleywegt [2]

[1] School of Mathematics, Georgia Institute of Technology; federico.bonetto@math.gatech.edu
[2] School of Industrial and Systems Engineering, Georgia Institute of Technology; vinod.cheriyan@gatech.edu, anton@isye.gatech.edu
* Correspondence: vinod.cheriyan@gatech.edu

Academic Editor: name
Version October 14, 2017 submitted to J. Risk Financial Manag.; Typeset by LaTeX using class file mdpi.cls

Abstract: Investors' forecasting behavior affects their trading decisions and the resulting asset prices. It has been shown in the literature how different, apparently reasonable investor forecasting behaviors can lead to qualitatively different asset price trajectories. For example, it has been shown that the weights that investors place on observed asset prices when forecasting future prices, and the degree of their confidence or panic when prices changes are forecasted, determine whether price bubbles, price crashes, and unpredictable price cycles occur. In this paper, we report the results of behavioral experiments to study investors' forecast formation. We present the design of an experiment consisting of multiple investors that participate in a market for a virtual asset. We conducted three experimental sessions with different participants in each session. Unlike most other market experiments with known finite durations and predictable decreasing prices toward the end of the experiment, our experimental sessions had random durations chosen to emulate an infinite horizon market with discounting. Asset price bubbles and cycles occurred in all three sessions. We fit different models of forecast formation to the observed data. There is strong evidence that the investors forecast future prices by extrapolating past prices, even when they know the fundamental value of the asset exactly and the extrapolated forecasts differ significantly from the fundamental value. The rational expectations hypothesis seems inconsistent with the observed forecasts. There is some evidence for adjustment of forecasts towards the fundamental value as the extrapolated forecasts deviate from the fundamental value. The calibrated forecasting models of all participants were consistent with dynamical systems that exhibit price bubbles and cycles.

Keywords: Keywords: Behavioral OR, Forecasting, Finance

1. Introduction

Data and forecasting form the foundations of both long-term planning and operational control. In spite of the widespread use of computers and algorithms to assist with data processing and forecasting, human decision makers continue to affect forecasts in important ways, including the introduction of cognitive biases, strategic biases, and overconfidence into forecasts. For surveys of the impact of human decision makers on forecasts, see for example Armstrong [1] and Ramnath *et al.* [2]. In this paper we focus on human forecasting behavior in a setting in which forecasters know the rational expectations value of the variable to be forecasted, and they also observe data of the variable to be forecasted, where the individual data values are allowed to deviate from the rational expectations value. We conducted experiments in which the participants were asked to forecast future prices, in which the actual prices were endogenous to the experiment through the trading decisions of the participants, and in which the rational expectations value of future prices (fundamental value of the traded asset) was clear and known to all the participants. One question is whether such participants place more emphasis on the rational expectations value or on the observed data to make their forecasts.

If investors' forecasts were consistent with the rational expectations hypothesis, then the volatility in forecasts and prices would be caused by "external" surprises and not by the "internal" forecasting behavior itself, and the forecasts and resulting prices would be less volatile than the earnings produced by the assets. However, LeRoy and Porter [3], and Shiller [45], and later Shiller [6], Campbell and Shiller [78], West [910], and LeRoy [11], presented empirical evidence that stock prices are typically much more volatile than the earnings produced by the stocks. Thus empirical evidence indicates that the rational expectations hypothesis does not provide an accurate model of investor forecasts and asset prices. Hence we are interested in studying alternative models of forecasting behavior.

Models of boundedly rational forecasting behavior have shown that small changes in investors' forecasting behavior can qualitatively change price trajectories. For example, Cheriyan and Kleywegt [12] showed that a small change in the weight that investors' forecasts place on recent data relative to older data can change price trajectories from convergent trajectories to cyclical trajectories consisting of repeated bubbles and crashes, and the instability is affected by investors' overconfidence in the information contained in observed price data. Also, it was shown that if investors' forecasts exhibit behavior called panic, then the price cycles can become unpredictable. Thus, an understanding of forecasting behavior contributes to an understanding of the qualitative behavior of price trajectories.

Therefore, our interest in investor forecasting behavior lead us to the following questions: Are there simple models that can describe investors' forecasts accurately? How accurately does the rational expectations hypothesis describe investors' forecasts? What does observed forecasting behavior imply regarding the resulting behavior of price trajectories, that is, whether price trajectories converge, or predictable price cycles occur, or trajectories are unpredictable? Questions such as these have been asked by both researchers as well as policy makers who are interested in studying, and sometimes containing or controlling, asset price dynamics.

In this paper we report experiments conducted to study investor forecasting behavior. We designed a market for a virtual asset in which the experimental investors report their price forecasts to the experimenter, and enter their buy and sell orders in the market, repeatedly over time. The market clearing prices are determined by the investors' buy and sell orders, and are therefore endogenous, and the market clearing prices are reported to the investors in real time. Before each experimental session, the investors were reminded of the theory of the fundamental value of an asset, and they were told how the earnings of the asset would be determined and what the resulting fundamental value of the asset was. Nevertheless, the resulting price trajectories exhibited cycles. We calibrated and compared various models of investor forecast formation. The experimental results indicate that:

1. The rational expectations hypothesis does not provide an accurate model of investor forecasting behavior, even when the earnings rate and fundamental value of the asset are known with certainty by investors.
2. A simple one-parameter exponential smoothing model of investor forecasting behavior is remarkably accurate. Models with a larger number of parameters provide only a slightly better fit.
3. There is strong evidence of investor overconfidence in the information content of observed price data. In the experiment, the fundamental value of the asset was known to the investors and therefore the observed prices contained no additional information regarding the fundamental value of the asset. Nevertheless, investor forecasts relied mostly on observed price data.
4. Except for a few experimental investors, the price forecasts were increasing in the extrapolated prices. That is, most experimental investors did not exhibit evidence of panic, as defined in Cheriyan and Kleywegt [12], in their forecasts.

1.1. Definition of Key Terms

A typical experimental setting consists of a number of participants buying and selling units of a virtual asset with virtual currency in a market. Usually, the participants can buy and sell assets. In some experiments, the participants can also be forecasters who only generate forecasts and do not trade.

A ***virtual asset market*** implements a specified ***trading mechanism*** for trading units of a ***virtual asset*** and ***virtual currency***. All virtual asset markets have a finite duration. The fundamental value of the asset may be specified exogenously or it can be computed from the dividend payoffs. The asset may provide either periodic dividends, or only a terminal dividend. Also, it may or may not have a salvage value.

An ***experiment*** consists of one or more sessions. A ***session*** involves a group of participants participating in the asset market. Each session consists of one or more ***trading runs***. Two or more trading runs with common participants are called ***repeated markets*** by some authors. At the beginning of each trading run, the investors are given an ***initial endowment*** of virtual currency and units of the virtual asset. The investors use the trading mechanism to participate in trades. At the end of the trading run, the participants are rewarded, typically as a function of their final portfolio of virtual currency earned and virtual asset in possession. In some experiments, there is also a reward associated with forecast accuracy.

The trading mechanism can be a continuous-time trading mechanism such as a double-auction, or a discrete-time trading mechanism such as a call market mechanism. For a market with a discrete-time trading mechanism, the trading run is divided into multiple ***(trading) periods***. During a period, traders enter buy and sell orders. At the end of the trading period, the price of the asset is determined by market clearing and the trades are executed based on the market clearing price. For a market with a continuous-time trading mechanism, the trading run consists of only one trading period. For example, in a double auction market, the traders may post buy and sell offers, or accept existing open offers resulting in trade; in this case the instantaneous price varies throughout the trading period.

Note that in our experiments, the object of interest is investor price forecast formation when both knowledge of the fundamental value of the asset as well as price data are available to the investor. As a result, each participant's forecast in each trading period in each session constitutes an observation. Thus, over the three experimental sessions that we conducted we collected 2760 observations.

2. Literature Review

Several researchers such as Smith *et al.* [13] have conducted experiments to study behaviors that may result in asset price bubbles. The subsequent literature in this area is vast. Here we focus on the body of work closely related to the topic of this paper. In particular, we focus on experimental work that study asset price expectation formation. We classify the experiments in this area based on the following properties: endogeneity of supply and demand, finiteness of the time horizon, market mechanism, reward mechanism, and purpose of the experimental work.

2.1. Exogenous and Endogenous Supply and Demand

The supply and demand of the asset are exogenously specified in experiments that study the effectiveness of market mechanisms in facilitating the convergence of prices to the equilibrium price. In such a setup, in each trading period some of the participants are given the role of buyers and the rest are given the role of sellers. Each buyer and seller is given an initial endowment of currency or the asset, respectively. Moreover, the value of the asset for each buyer and seller is specified privately to the buyer and seller. Thus, there are *a priori* supply and demand curves and a market clearing price in each trading run. Different market mechanisms, such as the double auction mechanism, are used to generate a price trajectory for the asset and this price trajectory is compared with the market clearing price. For example, Smith *et al.* [13], Ball and Holt [14] conducted such studies.

Most of the papers that our current work builds on use endogenous supply and demand [15–18]. In this setup, at the beginning of a trading run, every participant is given an initial endowment. Thereafter, in each period, each participant can enter buy and sell orders. Presumably, they make these decisions by taking into account the market clearing prices, their own forecasts of the future prices, as well as the revenue stream associated with owning the asset. Thus, in each period the supply and demand of the asset are endogenous to the experimental market. Our experimental setup also follows this endogenous supply and demand approach.

2.2. Finite Time Horizon and Learning

Many of the experiments reported in the literature have a pre-announced number of periods in each trading run. Typically, it is 9, 12, or 15 periods. For example, Haruvy *et al.* [17] studied price forecasts in trading runs consisting of 15 periods each. Lugovskyy *et al.* [18] also used trading runs consisting of 15-period-long trading runs to study the effect of tâtonnement mechanism on asset price bubbles.

In many experiments reported in the literature the asset is worthless at the end of the trading run. For example, Haruvy *et al.* [17] and Lugovskyy *et al.* [18] considered an asset that paid a fixed expected dividend d at the end of each period $t = 1, \ldots, 15$ and that had no salvage value. Assuming no discounting, the fundamental value of the asset at the beginning of period t was $\bar{p}_t = (16 - t)d$. Thus, the fundamental value decreased linearly throughout the life of the asset and the participants knew *ex ante* that the asset would be worthless at the end of the trading run. Consequently, prices fell towards the end of the trading run. Since the dividend process was pre-announced, the participants could anticipate this price trajectory. In repeated markets as in Haruvy *et al.* [17], one reasonable estimator of when the price process peaks is the peak time period in the previous trading run. Consequently, in each trading run the peak occurs earlier than in the previous trading run, and at the same time the size of the bubble decreases as the fundamental value is also larger in earlier periods. This observation leads to the conclusion that learning causes the bubbles to disappear.

One question that arises naturally is what happens when investors do not have the opportunity to form expectations using backward induction — e.g. reasoning such as "if it is certain that the price is 0 at the end of period 15, it may not be very high at the end of period 14". Hirota and Sunder [16] considered a setup which makes such backward induction difficult. In their setup, the asset paid a known terminal dividend only at maturity (at the end of 15 or 30 trading periods). In some trading runs, the duration was 15 periods and the asset matured at the end of 15 periods; in this setting, backward induction was possible. In other trading runs, the asset matured at the end of 30 periods; however the participants were told that the experiment would end at an unknown time most probably much earlier than the life of the asset. If the trading run lasted for 30 periods, they would receive the known dividend, otherwise they would receive a payout of the average price forecast for the asset when the trading run ends. In this setting, the apparently random end of the market removed the anchor that the terminal dividend provided for backward induction. Their results suggest that the latter setup is more conducive to formation of asset price bubbles.

Another approach to eliminate backward induction is to make the number of periods in the trading run random. In this kind of probabilistic stopping, each period the trading run ends with a pre-specified probability. It can be shown that for an asset paying constant dividends, the fundamental value under probabilistic stopping with probability p is equal to the fundamental value under an infinitely-long trading run with discount factor of $1 - p$ (see Lemma 1). Camerer and Weigelt [15] used trading runs with probabilistic stopping to study convergence of prices to fundamental value and to test if prices are rational forecasts of the fundamental value. Ball and Holt [14] used a variation of this in which each individual asset can expire with a pre-specified probability. However, in their experiment the total length of the trading run was pre-announced (10 periods), thus the probabilistic stopping only created an effective discounting but did not eliminate the end of horizon effect.

In the current work, we use probabilistic stopping; in each period the trading run can end with a pre-announced stopping probability similar to Camerer and Weigelt [15]. Thus the participants cannot use backward induction to value the assets. Moreover, we chose the total duration of the experimental sessions to be about three and a half hours, which is atypical in this literature. Our reasoning was that even with probabilistic stopping, towards the end of the time slot for the session, the participants may reason that even if the trading run does not end, the session will end soon; consequently they may want to apply backward-induction thinking and offer to sell their assets for lower prices. A long duration for the experimental sessions helps to reduce the effect of such thinking at least in the earlier trading periods.

2.3. Market Mechanism

Depending on the objectives of their studies, different researchers use different market mechanisms. For example, Ball and Holt [14], Camerer and Weigelt [15], Hirota and Sunder [16] used a continuous double auction mechanism. In this mechanism, participants announced their bids and asks;

at any time any participant could accept a bid or an ask, and the price associated was recorded as the current transaction price. In these studies, the average of the transaction prices was used as a proxy for the market clearing price for that period. On the other hand, researchers such as Haruvy *et al.* [17] used a call market mechanism. In this mechanism, in each trading period, each participant entered buy and/or sell orders. At the end of the period, all orders were aggregated into market supply and demand curves. An equilibrium price was determined that cleared the market. All feasible trades (bids above the equilibrium price, asks below the equilibrium price) were executed. Bids and asks at the equilibrium price might be only partially executed (in which case, some tie breaking or apportioning mechanism was used). Other authors have used other market mechanisms. For example, Lugovskyy *et al.* [18] used the tâtonnement mechanism. In this mechanism, in each period the market maker announced a suggested price and the participants placed their tentative buy and sell orders at that price. If the aggregate supply equaled the aggregate demand then the suggested price became the market clearing price and all trades were executed; otherwise, the market maker announced an updated suggested price and the participants again placed their buy and sell orders at the new suggested price. The process repeated until the market cleared or a particular exit criterion was met (e.g. if the number of iterations reached a maximum).

In our setup, we use a slight modification of the call market mechanism. In each period, each participant can place multiple buy and sell orders, effectively specifying their individual supply and demand curves. These are aggregated into market supply and demand curves. Their intersection determines the market clearing price. Trades are executed at the market clearing price — all buy orders above the market clearing price and all sell orders below the market clearing price are fully executed. Orders at the market clearing price are filled completely if the aggregate supply and demand match exactly; in general, the maximum number of orders at the market clearing price are filled, and they are filled in the temporal sequence in which the orders were entered (i.e. on a first-come first-served basis). Some orders can be partially filled.

We chose to use this market clearing mechanism primarily to make sure that each period results in a unique equilibrium price. The prices in prior periods are common knowledge to all participants, and we hypothesize that the participants use this information to form price forecasts.

2.4. Reward Mechanism

The reward mechanisms found in the literature are quite varied. Most experiments pay a combination of a fixed participation fee and a reward proportional to the participant's performance in the experiment. The participant's performance is evaluated based on factors like the total virtual wealth at the end of the experiment and the forecast accuracy (if the experiment collected forecast data). Some researchers used a fixed exchange rate (i.e. a pre-announced virtual currency to real currency exchange rate), a pre-announced payment schedule (e.g. for accuracy of predictions), or they divided a pot of money in proportion to the total virtual wealth at the end of the experiment (e.g. Hirota and Sunder [16]). In Hirota and Sunder [16] participants who were price predictors were paid based on the accuracy of their predicted prices; in Haruvy *et al.* [17] the participants also received payments for accurate predictions. Ball and Holt [14] argued against rewarding only the participant with the highest earnings as it may induce indifference or extreme risk seeking in people who are relatively behind in terms of their earnings.

In our experiment, rather than provide a fixed participation fee to each participant, three prizes were awarded at the end of each experimental session. Two of these were a function of their total wealth at the end of the experiment and one was for the best forecast accuracy. The details and rationale for these prizes are given in Section 4.4.

2.5. Purpose of the Experiments

Many of the experimental studies focus on convergence of asset prices to the theoretical fundamental value. For example, Ball and Holt [14] compared the prices resulting from trading with the fundamental value.

Some of the researchers went further to study price bubbles in repeated markets with the same set of participants. Haruvy *et al.* [17] found that the magnitude of price bubbles decreased as the same investors participated in successive trading runs.

Some of the experimental studies addressed belief or expectation formation among participants. For example, Haruvy *et al.* [17] explicitly captured participants' price expectations for future periods and concluded that prices converged to the fundamental value ahead of beliefs. Hirota and Sunder [16] studied the effect of first and higher-order beliefs by designing experiments with a publicly announced range of dividends that is bigger than the privately communicated actual range of dividends.

There are also other specialized objectives to some of these studies. Hirota and Sunder [16] studied the effect of investors' decision horizon on the presence of bubbles and they found that shorter decision horizons (compared with the maturity of the asset) can lead to larger bubbles. Lugovskyy *et al.* [18] studied the effect of the tâtonnement trading institution on price bubbles and concluded that the participants were able to learn about supply and demand during the tâtonnement process, thereby reducing price bubbles.

The goals of our experiments were two-fold: One goal was to collect investors' price forecasts so that various models of investor forecast formation could be calibrated and compared. Another goal was to determine which of the qualitative regimes of asset price trajectories identified in models (convergence of asset prices, cycling of asset prices, or unpredictable asset price trajectories) observed forecasting behavior corresponds to. We conducted these experiments in a setting with endogenous supply and demand, while minimizing the propensity of the participants to sell their assets for lower prices when they think that the experiment is near to completion.

3. Some Models of Price Forecast Formation

In this section we briefly review a class of models of asset price forecasts and the resulting market clearing prices. Time is indexed by $t = 1, 2, \ldots$. Each unit of asset pays a dividend d_t at time t. Let p_t denote the price of the asset at time t before the dividend d_t has been paid. Let $\hat{p}_{t+1}$ denote the investors' forecast at time t of the price at time $t+1$. Suppose that the investors in the market are indifferent between investing and not investing in the asset if they expect a rate of return of $\tilde{r}$, that does not depend on time. In each time period, investors can rebalance their portfolios without transaction cost, and they forecast only one period into the future. Then the indifference price at time t is

$$p_t \;=\; \mathbb{E}[d_t] + \frac{1}{1+\tilde{r}}\left[\hat{p}_{t+1}\right] \tag{1}$$

In the experimental sessions we conducted, the dividend process was revealed to everyone, enabling the participants to compute $\mathbb{E}[d_t]$ explicitly. Along with specification of a model to compute $\hat{p}_{t+1}$, and appropriate initial conditions, equation (1) defines a discrete time dynamical system model of the asset price process.

For the 'fundamentalist' investor, the price forecast $\hat{p}_{t+1}$ is simply the fundamental value $\bar{p}_{t+1}$

$$\bar{p}_{t+1} = \sum_{i=0}^{\infty} \frac{\mathbb{E}[d_{t+1+i}]}{(1+\tilde{r})^i}$$

assuming the infinite sum is convergent. (That is, the expected growth rate of the forecasted dividends is less than $\tilde{r}$.) However, in general, forecasts are also influenced by observed data, especially the most recent data. Different assumptions about how past data is used gives rise to different models for forecast formation.

3.1. Extrapolation-Correction Models of Forecast Formation

The model given here is an adaptation of the more generic model given in Cheriyan and Kleywegt [12]. We consider forecasts that depend on both observed data as well as the fundamental value, as explained in the next four subsections.

It is convenient to consider the prices and forecasts scaled by the fundamental value, as follows:

$$\begin{aligned} \pi_t &:= \frac{p_t}{\bar{p}_t} \\ \hat{\pi}_{t+1} &:= \frac{\hat{p}_{t+1}}{\bar{p}_t} \end{aligned}$$

Let y_t denote the growth rate of scaled prices in period t, that is,

$$y_t := \frac{\pi_t}{\pi_{t-1}}$$

3.1.1. Extrapolated Price

The extrapolation forecast $\hat{y}_t$ of the price growth rate is given by

$$\hat{y}_t = (1-\alpha)\hat{y}_{t-1} + \alpha y_{t-1}$$

where $\hat{y}_2$ is an appropriate initial value. Thus, the extrapolation forecast is given by the exponential smoothing forecast. Note that α corresponds to the weight the investor gives to the most recent price ratio. If the forecaster used only the extrapolation forecast, then the corresponding scaled forecast for the price at time $t+1$ would be

$$\hat{\pi}_{t+1} = \pi_{t-1}\hat{y}_t^2 \quad (2)$$

Note that at time t, $(p_1, \ldots, p_{t-1})$ are known to the investor. The price p_t is yet to be realized and will depend on the buy and sell orders of the investor.

3.1.2. Price Forecast

We consider a price forecast that is a function H of the extrapolation forecast and the fundamental value. Although the price forecast depends on both the extrapolation forecast and the fundamental value, it is convenient to use notation that shows only the scaled extrapolation forecast as an argument of H. Thus, the scaled price forecast $\hat{\pi}_{t+1}$ is given by

$$\hat{\pi}_{t+1} \quad = \quad H(\pi_{t-1}\hat{y}_t^2)$$

where $H : \mathbb{R}_+ \mapsto \mathbb{R}_+$ is a function that captures some behavioral characteristics of the investor. A fairly general class of H functions can be considered. The only property we require here is that

$$H(1) = 1$$

That is, if the extrapolation forecast is equal to the fundamental value, then the price forecast should just equal the fundamental value. As an example, $H(\theta) \equiv 1$ corresponds to rational expectations, while $H(\theta) \equiv \theta$ corresponds to extrapolation-only price forecast.

Generally, in a neighborhood of $\theta = 1$, H is nondecreasing, that is the price forecast increases with the extrapolation forecast. For larger values of θ, H may increase, perhaps at a decreasing rate. If, on the other hand, H increases and then decreases, we say that the investor exhibits ***panicking behavior***. This corresponds to the investor losing confidence in the extrapolation forecast and adjusting the price forecast closer to the fundamental value. Such functions H will be called non-mononotic.

Substituting the scaled price and forecast into (1), the market clearing equation becomes

$$\pi_t = \frac{\mathbb{E}[d_{t+1}]}{\bar{p}_t} + \frac{1}{1+\tilde{r}}\frac{\bar{p}_{t+1}}{\bar{p}_t}H(\pi_{t-1}\hat{y}_t^2)$$

Finally, in our experiments, $\mathbb{E}[d_{t+1}]$, $\bar{p}_t$ and $\bar{p}_{t+1}$ can be computed explicitly from the dividend process and the salvage value of the stock. These computations are given in Appendix A.

3.2. Linear Model of Forecast Formation

We also fit the data with a linear model given in Brock and Hommes [19]. According to this model, the investors price forecast for next period, $\hat{p}_{t+1}$ is given by

$$\hat{p}_{t+1} = \bar{p}_{t+1} + b_0 + b_1(p_{t-1} - \bar{p}_{t-1}) \tag{3}$$

where $\bar{p}_t$ is the fundamental value at time t, and b_0 and b_1 are constants. Thus, the forecasts are linear functions of the deviations of the past price from the fundamental value.

4. Design of the Behavioral Experiment

Our experiment consisted of three sessions of a virtual asset market with a discrete-time trading mechanism, as given in Process Flow 1.

Process Flow 1 Process flow for an asset market with a discrete-time trading mechanism

1. $t \leftarrow 0$
2. $t \leftarrow t + 1$ (New trading period begins)
3. Investors generate forecasts
4. Investors enter buy and sell orders
5. Market clears according to the specified market clearing algorithm. Price p_t is determined
6. Trades are executed based on the market clearing price p_t
7. Dividend d_t is paid for each unit of asset held
8. Check termination of trading run. If not ended, go to step 2; otherwise compute rewards and debrief.

4.1. Approaches to Reduce End of Horizon Effect

In most market experiments the asset becomes worthless at the end of the last period. Therefore, if the total number of periods in the trading run is known in advance, the asset price decreases to zero towards the end of the trading run. If there are repeated trading runs in an experimental session, the decreasing trend of prices during a trading run eventually reduces the occurrence of price bubbles in the later trading runs.

To alleviate this effect caused by a known time horizon, we used probabilistic stopping. At the end of each period the trading run can end with a stopping probability $p = 0.02$. Thus, the total number of periods in each trading run is a geometric random variable. This was explained to all participants, including the value of the stopping probability, in advance of each trading run.

Unlike traditional classroom experiments that last 1 or 2 hours, we set up our sessions to last for three and a half to four hours. Thus, at least during the initial part of the session, the participants did not consider the end of the session when making their trading decisions.

We also instituted a salvage value of 100 Experimental Currency Units (ECUs) for each unit of asset held when the trading run ended. Theoretically, this just adds a constant to the fundamental value of the asset. We decided to add the salvage value after an initial trial run of the experiment (the data from the trial run were not used) to reduce the fixation of some participants on the possibility that the asset may become worthless at any time.

4.2. Dividend Structure

In each period, a dividend was paid for each unit of asset held. The dividend was added to the cash-on-hand of each participant.

For the first two sessions, the dividend was fixed to 10 ECUs per unit of asset in each period. For the third session, the dividends were the same for all participants but were random for each period. In each period the market was in a state $X_t \in \{\text{low}, \text{high}\}$. The state transition followed a discrete time Markov chain with the transition matrix

$$P = \begin{pmatrix} 3/4 & 1/4 \\ 1/4 & 3/4 \end{pmatrix}$$

The initial state, X_1, was chosen at random with equal probability. The dividend D_t for each period had a distribution depending on X_t as given in Table 1.

Table 1. Conditional probability mass function for the dividends.

	$P(D_t = 5\|X_t)$	$P(D_t = 10\|X_t)$	$P(D_t = 15\|X_t)$	$\mathbb{E}[D_t\|X_t]$
$X_t =$ low	0.6	0.3	0.1	7.5
$X_t =$ high	0.1	0.3	0.6	12.5

4.3. Experimental Sessions

The experiment consisted of three sessions, each with a new group of participants. Each session lasted for about four hours in total and consisted of an overview, training trading run, actual trading runs, and post-session debriefing.

The overview included an explanation of the Process Flow 1, probabilistic stopping of trading runs including the memoryless property of the geometric distribution, the dividend process, the concept of fundamental value and its calculation, the process for placing buy and sell orders, and the market clearing process to determine the market clearing price and market clearing trades. (All participants were masters students in the Quantitative and Computational Finance program or Ph.D. students, and the explanation of fundamental value was a review of a concept already familiar to them.) The stopping probability for all trading runs was $1/50$ and this was announced to the participants.

The training trading run lasted four periods. The purpose of the training trading run was to familiarize the participants with the market simulation software, by entering forecasts, placing buy and sell orders, and answering questions that tested their understanding of the experiment. It was made clear that the number of periods was fixed only for the training trading run. The earnings in the training trading run were not taken into account for the determination of the eventual compensation.

After the training trading run, it was announced that a real trading run was about to start and that the total number of periods would be random as explained. The experiment started with a number of questions for each participant to test participants' understanding of the experiment. Each participant was endowed with an initial portfolio of 5000 Experimental Currency Units (ECUs) and 50 units of the virtual asset.

Thereafter, period 1 started. In each period each participant was shown a screen displaying the participant's current cash and asset balance, a chart of the previous prices and a table of previous prices, the dividend history, and previously executed buy and sell orders for the participant, input boxes to answer questions described below, an input box to enter the participant's price forecast, and input boxes to enter buy and sell orders. For the third session with Markov dividends, additional information was displayed as described in the next section.

In each period each participant had to enter answers to the following questions:

- The expected number of time periods remaining in the trading run.
- The expected total dividends paid by one unit of asset from the current period until the end of the trading run.
- The total of the expected dividend from the current period until the end of the trading run and the salvage value for one unit of asset held until the end of the trading run.
- The participant's forecast of the price per unit of asset.
- The current state of the market (only session 3 — Markov dividend case).

After they entered the feedback and forecast information, the order entry portion of the screen was enabled and they could enter multiple buy or sell orders. Participants could also not trade in a period, however, the software design enforced that there would be at least one order entered among all participants in every period. This was required to ensure that there is an equilibrium price in every period.

Once a participant made her entries, a wait screen was displayed. After all participants made their entries, the market was cleared based on the orders and the new equilibrium price was determined.

Subsequent periods followed the same pattern until the computer determined that the trading run was over.

At the end of the trading run, the participants' portfolios were converted into virtual currency. Then, the winners were determined and were announced.

Finally, there was a short debrief session where we gathered feedback from the participants about the session.

The specific details of the experimental sessions are given in Table 2.

Table 2. Parameters for the Experimental Sessions

Session	Participants	Periods	Dividend Type	Mean Dividend (ECUs)	Salvage Value (ECUs)
1	16	37	Constant	10	100
2	17	64	Constant	10	100
3	15	72	Markovian	10	100

4.4. Incentive Mechanism

Rather than provide a fixed participation fee to each participant, three prizes were awarded at the end of each experimental session.

1. Two randomly selected players split $300 in proportion to the value of their final amount of virtual currency, after payment of the salvage value.
2. The participant with the largest final amount of virtual currency received a prize of $100.
3. The participant with the smallest mean square forecast error received a prize of $100.

The motivation for this incentive scheme is as follows. Other behavioral experiments have shown that people are more willing to perform a task when either it has no remuneration (i.e. it is considered as a help) or when the remuneration is significantly large. In our case, $18.75 (= $300/16) would not have been a lucrative enough participation fee for four hours' time. On the other hand, splitting the $300 among two randomly selected participants makes the incentive more lucrative. This is supported by the framing effect and pseudocertainty effect reported by Kahneman and Tversky (Tversky and Kahneman [20]) wherein when evaluating such conditional situations, people evaluate the options assuming that the selection process has already happened. Thus, in our case, the participants would be looking at an amount of the order of $150, which is lucrative. Also, random selection of two participants helped avoid participants dropping out of the experiment if they are doing poorly.

The prize for the largest portfolio incentivized playing the game strategically and thoughtfully and the prize for the smallest forecast error incentivized careful forecasting and reporting of forecasts.

Our incentive structure is somewhat unconventional. Our incentive structure is such that the risk faced by a participant who invests in the virtual "risky" asset is not the same as the risk that an investor would have faced if the market were real and the investor invested in an asset with the same dividends. Of course, the risk faced by a participant who invests in the virtual risky asset affects the price that the participant is willing to virtually pay for the asset during an experiment. Also, rewarding the participant with the largest virtual wealth at the end of a session disproportionately creates tournament incentives that may also affect the price that the participant is willing to virtually pay for the asset during an experiment. When making these observations, it is important to keep in mind that the purpose of the study was to model investor forecast formation and not price formation in a market.

4.5. Participant Demographics

The participants were students belonging to the Ph.D. (ISyE, Math) and Masters (Quantitative and Computational Finance — QCF) programs. We chose to include only Masters or Ph.D. students to ensure that they had sufficient background knowledge of basic probability. Basics of asset valuation, discounted cash flow, and the memoryless property of the geometric distribution were explained to all the participants during the overview session.

All subjects gave their informed consent for inclusion before they participated in the study. The study was conducted in accordance with the Declaration of Helsinki, and the protocol was approved by the Ethics Committee of TODO XXX (Project identification code).

5. Summary of Observations

Though the first session had two trading runs, the second trading run did not complete within the time allotted for that session. So we used only the completed trading run for data fitting. There was only one trading run per session for the remaining two sessions. Therefore, from here on we will use the terms session and trading run interchangeably.

5.1. Equilibrium Price

Since in each period, the trading run could end with a stopping probability of $1/50$, the number of time periods T was Geometric with parameter $p = 1/50$. The expected total returns for one unit of stock can be computed as

$$\bar{p}_t = \mathbb{E}\left[\sum_{t=1}^{T} d_t\right] + s = \mathbb{E}[T]\mathbb{E}[d_t] + s$$

where d_t is the dividend at period t. For sessions 1 and 2, the fundamental value was constant throughout and equaled 600 ECUs. For session 3, the computation of fundamental value is given in Appendix A. For session 3, the fundamental value at the beginning of the trading run was also 600 ECUs.

Figure 1 shows the realized equilibrium prices for the three sessions. Sessions 1 and 2 had prominent price cycles whereas session 3 had milder price cycles.

In all three sessions, the prices started from well below the fundamental value and started increasing thereafter. This phenomenon has been observed in many other market experiments, see for example Camerer and Weigelt [15], Haruvy *et al.* [17].

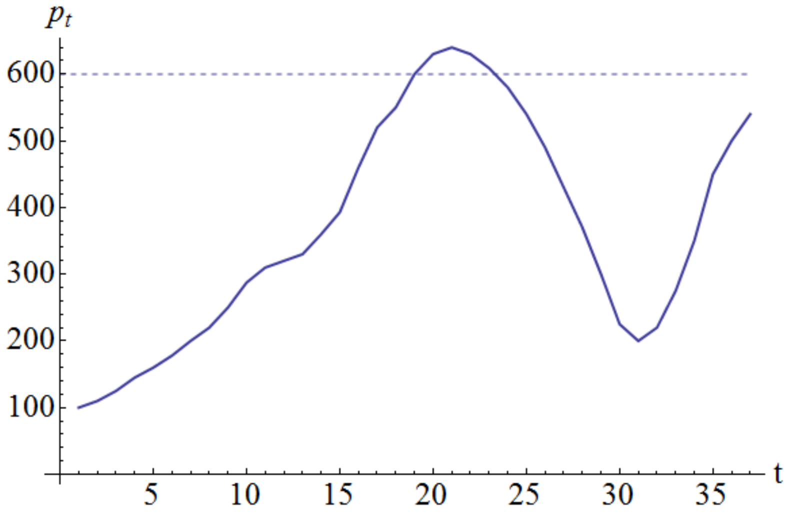

(a) Session 1

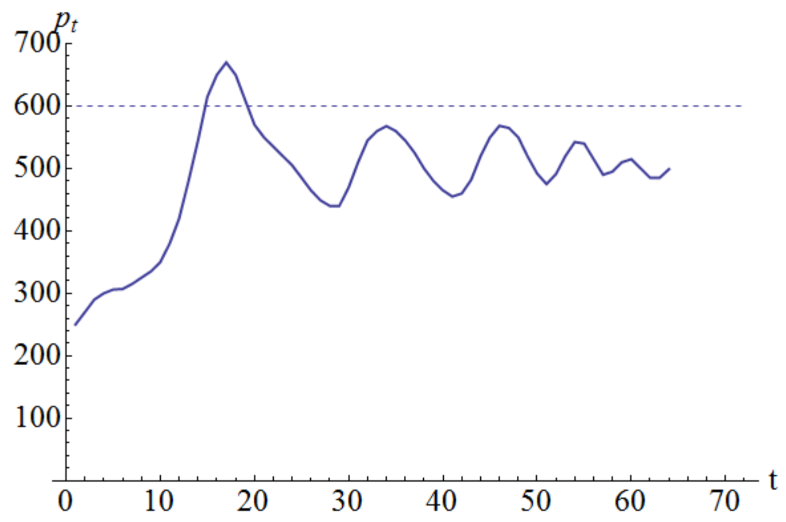

(b) Session 2

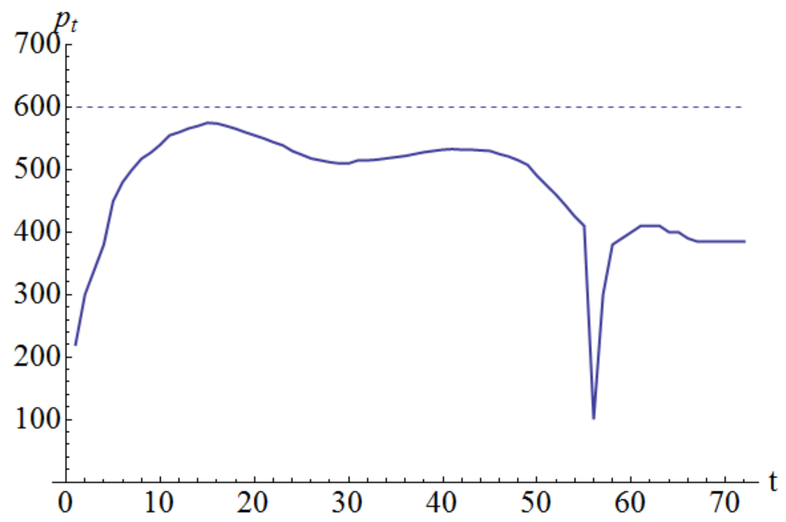

(c) Session 3

Figure 1. Equilibrium Prices From Sessions

In the third experimental session, there was a sudden downward spike in period 56. It was revealed to have been caused by one participant who panicked for some unknown reason and wanted to get rid of all his assets. So he offered to sell them for 0.01 ECU each. The market rebounded the very next period. (The data from this spike onwards were excluded from the calculations that follow.)

From the graphs of the price process, it appears that the participants had a perceived value that was lower than the fundamental value of 600 ECUs. In other words, they seemed to have made most buy-sell decisions as if the fundamental value was somehow smaller than the true fundamental value.

5.2. Price Forecasts

In each period, each participant was required to enter a price forecast for that period. Recall that the market clearing price in a period was determined only after the buy and sell orders were processed. Figure 2 shows a plot of the price forecasts for each session by each participant along with the market clearing price (solid line). It can be seen that the price forecasts usually followed the market clearing price.

The forecasts in the initial few periods are interesting. It seems that some of the participants started with a near-rational forecast and adjusted them down as the trading run progressed. On the other hand, some "skeptics" started with a very low forecast and adjusted them upwards as the trading run progressed.

5.3. Earning Forecasts

As part of the initial survey, each participant was asked the following question:

> Given the compensation rules and the number of participants, how much money do you think you will earn in this experiment?

Table 3 summarizes the responses. In every session, the average of the participants' expected earnings was at least twice the average prize money per participant for that session. Though there were some participants who expected to earn less than this average, a majority of them reported a higher than average earnings. Perhaps, this suggests an individual belief that they are better than the rest. This observation is reminiscent of the investor overconfidence mentioned in literature (e.g. Scheinkman and Xiong [21], Chuang and Lee [22]).

Table 3. Expected earnings for the three sessions.

Session	Participants	Number of Valid Responses	Average Expected Win	Average of Prize Money Per Participant
1	16	14	90.43	31.25
2	17	15	66.07	29.41
3	15	13	68.91	33.33

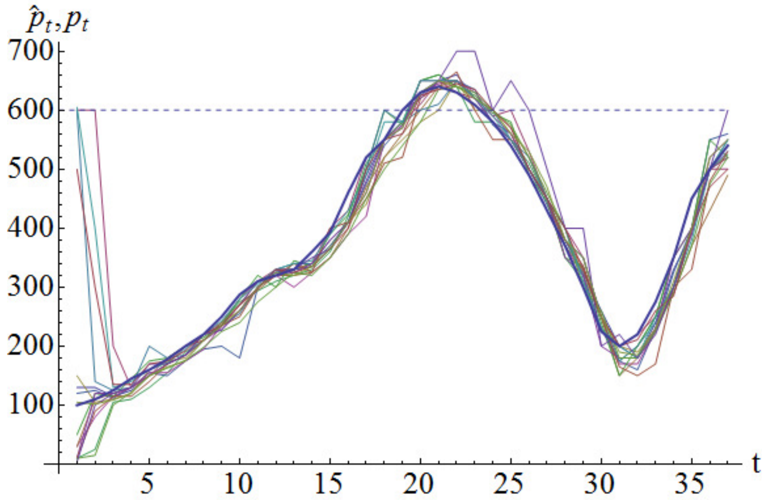

(a) Session 1

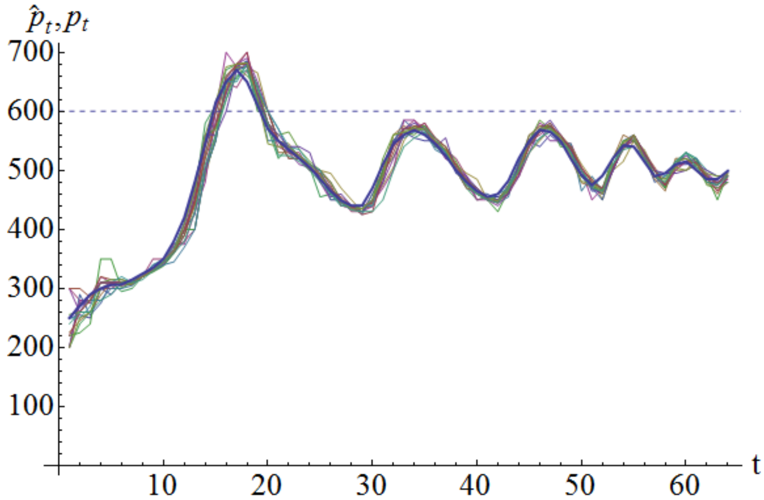

(b) Session 2

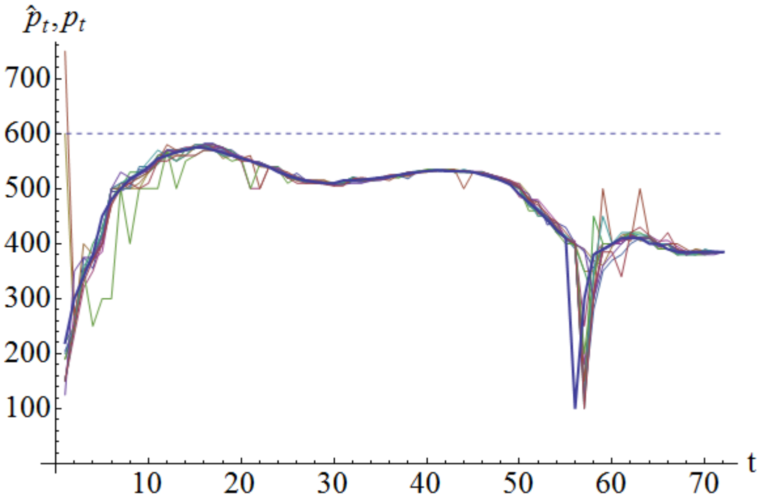

(c) Session 3

Figure 2. Price forecasts for each period.

6. Comparison of Models Fitted with Data

6.1. Data Used

The data used for model calibration are the equilibrium prices and the participants' forecasts from the three sessions

First, we cleaned the data by fixing obvious typographical errors. Typically, the error was omission of the decimal point.

In the initial few periods of each of the three sessions, the equilibrium price increased from a value much lower than the fundamental value. A possible explanation is that the participants were trying to learn how the rest of the participants would behave. However, the models that we wanted to fit with the data were not intended to capture such initial learning or adaptation. Consequently, we used a subset of the data when the initial effect had passed. Moreover, we wanted to choose a subset of the data in a way that is endogenous to the data itself and not dependent on the model that we fit. To achieve this, we used the following approach: In the equilibrium price chart, let p_0 be the price in the first price trough. We discarded the data from the initial periods during which the price was smaller than p_0. According to this approach, the number of dropped periods for the three sessions was between 7 and 11. Finally, to use the same number for all three sessions, we simply discarded the data from the first 10 periods in each session.

We also checked this approach using the calibrated models with the quadratic H function given in Section 6.3.5. We calibrated a sequence of models by successively dropping more initial periods from the data set. We observed that the fitted parameter values stabilized by the time we dropped the first 10 periods' data.

In addition, for session 3, we dropped the data from period 56 onwards. This was done to avoid the effects of the spike that occurred in period 56.

6.2. Parameter Estimation Method

Parameter estimation was done by least squares fitting. We use the following notation.

$$\begin{aligned} \sigma &= 1,2,3 \qquad \text{session index} \\ n_\sigma &= \text{number of periods in session } \sigma \\ N_\sigma &= \text{number of participants in session } \sigma \\ X_{\sigma,j} &= \text{equilibrium price in period } j \\ Y_{\sigma,j,u} &= \text{price forecast reported by participant } u \text{ in session } \sigma \text{ in period } j \end{aligned}$$

($n_3 = 55$, as discussed above.)

Consider a family of models of forecast formation represented as

$$Y_{\sigma,j,u} = \psi_j(X_{\sigma,1}, X_{\sigma,2}, \ldots, X_{\sigma,j-1}; \gamma) + \varepsilon_j$$

where γ denotes the vector of parameters of the model and ε_j are i.i.d. $N(0, \sigma^2)$ random variables.

The parameter fit for participant u in session σ is given by

$$\hat{\gamma}(\sigma, u) \quad := \quad \arg\min_{\gamma} \sum_{j=11}^{n_\sigma} (Y_{\sigma,u,j} - \psi_j(X_{\sigma,1}, X_{\sigma,2}, \ldots, X_{\sigma,j-1}; \gamma))^2$$

The parameter fit using all data in all sessions is given by

$$\hat{\gamma} \quad := \quad \arg\min_{\gamma} \sum_{\sigma=1}^{3} \sum_{u=1}^{N_\sigma} \sum_{j=11}^{n_\sigma} (Y_{\sigma,j,u} - \psi_j(X_{\sigma,1}, X_{\sigma,2}, \ldots, X_{\sigma,j-1}; \gamma))^2$$

6.2.1. Leave-One-Out-Cross-Validation (LOOCV) Approach

To compare the different models that were fitted, we use the Leave-One-Out-Cross-Validation (LOOCV) approach. According to this approach, the observations in the data set are partitioned into subsets. The model is fitted using all but one subset, and then the fitted model is used to calculate forecasts for the omitted subset. The forecast error is recorded. This step is repeated with each subset of the data set left out in turn. Then the root mean square error (RMSE) of all the recorded forecast errors is calculated. The LOOCV RMSE of different models are compared. We perform the LOOCV validation with two types of subsets: (1) leave-one-period-out LOOCV, in which each period's data for each participant is a subset, and (2) leave-one-session-out LOOCV, in which each session's data is a subset.

Leave-One-Period-Out LOOCV

In this case, separate parameters are fit for each individual participant. The resulting RMSE gives a measure of predictability when a separate model is fitted for each participant. The leave-one-period-out LOOCV RMSE for a participant is computed as follows.

Let $\hat{\gamma}_{-i}(\sigma, u)$ denote the vector of parameters fitted for the participant u in session σ after dropping the data for period i. That is,

$$\hat{\gamma}_{-i}(\sigma, u) \quad := \quad \arg\min_{\gamma} \sum_{j=11, j\neq i}^{n_\sigma} (Y_{\sigma,u,j} - \psi_j(X_{\sigma,1}, X_{\sigma,2}, \ldots, X_{\sigma,j-1}; \gamma))^2$$

Then, the leave-one-period-out LOOCV RMSE of the model ψ for the participant is given by

$$RMSE^P_\psi(\sigma, u) \quad := \quad \sqrt{\frac{\sum_{i=11}^{n_\sigma} (Y_{\sigma,u,i} - \psi_i(X_{\sigma,1}, X_{\sigma,2}, \ldots, X_{\sigma,i-1}; \hat{\gamma}_{-i}(\sigma, u)))^2}{n_\sigma - 10}}$$

The leave-one-period-out LOOCV coefficient of variation is the leave-one-period-out LOOCV RMSE scaled by the fundamental value, given by

$$CV^P_\psi \quad := \quad \frac{RMSE^P_\psi(\sigma, u)}{600} \times 100\%$$

Leave-One-Session-Out LOOCV

In this case, one set of parameters are fit for all participants and all periods in all sessions but one (thus in two sessions). The fitted model is then used to calculate forecasts for all particpants and all periods in the omitted session. The resulting RMSE gives a measure of predictability if a common model is fitted for all sessions and participants.

The leave-one-session-out LOOCV RMSE is computed as follows. Let $\hat{\gamma}_{-s}$ denote the vector of parameters fitted after dropping all observations in session s. That is,

$$\hat{\gamma}_{-s} := \arg\min_{\gamma} \sum_{\sigma=1,\sigma\neq s}^{3} \sum_{u=1}^{N_\sigma} \sum_{j=11}^{n_\sigma} (Y_{\sigma,j,u} - \psi_j(X_{\sigma,1}, X_{\sigma,2}, \ldots, X_{\sigma,j-1}; \gamma))^2$$

Then, the leave-one-session-out LOOCV RMSE is given by

$$RMSE_\psi^S := \sqrt{\frac{\sum_{s=1}^{3} RMSE_s^2}{3}}$$

$$RMSE_s := \sqrt{\frac{\sum_{u=1}^{N_s} \sum_{j=11}^{n_s} (Y_{s,j,u} - \psi_j(X_{s,1}, X_{s,2}, \ldots, X_{s,j-1}; \hat{\gamma}_{-s}))^2}{N_s(n_s - 10)}}$$

The leave-one-session-out LOOCV coefficient of variation is the leave-one-session-out LOOCV RMSE scaled by the fundamental value, given by

$$CV_\psi^S := \frac{RMSE_\psi^S}{600} \times 100\%$$

6.3. Comparison of Fitted Models

We fit various models with the data. The individual models are covered in subsequent sections; the key parameters of these models are given in Table 4.

The model $BASE$ is the pure rational expectations model — according to the rational expectations model, the participant's forecast equals the fundamental value (also recall that the participants were told the fundamental value in each period). Note that model BASE has no parameters to be fitted. Next we consider two one-parameter models. The model F is a "modified rational expectations" model where the hypothesis is that the participants believe that the fundamental value is some number other than the fundamental value communicated to them, and they forecast their belief of the fundamental value. Model ES assumes that the investors use only simple exponential smoothing and therefore has only one parameter. Model BH has 2 parameters, and is based on Brock and Hommes [23]. Models $ESH2$ and $ESHE$ belong to the family of extrapolation-correction models proposed in Cheriyan and Kleywegt [12], with 3 and 4 parameters respectively.

Table 4. Details of the models fit to the data

Model	Generalized Mean (m)	Correction Function (H)	Parameters to fit
$BASE$	N/A	N/A	None
F	N/A	N/A	$\bar{f}$
ES	Exponential Smoothing	$H(\theta) = \theta$	α
BH	N/A	N/A	b_0, b_1
$ESH2$	Exponential Smoothing	H_2	$\alpha, \bar{f}, \rho$
$ESHE$	Exponential Smoothing	H_9	$\alpha, \bar{f}, \rho, \eta$

508 6.3.1. Rational Base Case

For comparison, we take as base case the rational expectations forecast of a fully informed participant (BASE). That is

$$\psi_j(\mathbf{X}) = \bar{p}_t = 600 \qquad \text{(BASE)}$$

509 Thus, there are no parameters to fit for this model. For the Markov case in session 3, we used the
510 reference fundamental value of 600 to fit the model.

511 Table 5 gives the leave-one-period-out LOOCV RMSE for each session.

Table 5. Forecast errors of the rational base case $(BASE)$ for the three sessions.

	Session 1		Session 2		Session 3	
Part.	$RMSE^P_{BASE}$	$CV^P_{BASE}(\%)$	$RMSE^P_{BASE}$	$CV^P_{BASE}(\%)$	$RMSE^P_{BASE}$	$CV^P_{BASE}(\%)$
1	217.66	36.28	105.27	17.55	82.29	13.72
2	211.61	35.27	107.25	17.88	86.70	14.45
3	220.79	36.80	105.18	17.53	83.20	13.87
4	221.02	36.84	106.15	17.69	84.93	14.15
5	221.92	36.99	106.30	17.72	82.99	13.83
6	230.33	38.39	110.22	18.37	81.28	13.55
7	231.04	38.51	107.19	17.86	84.32	14.05
8	218.86	36.48	107.87	17.98	84.37	14.06
9	220.13	36.69	110.19	18.37	84.01	14.00
10	225.42	37.57	108.19	18.03	82.76	13.79
11	233.16	38.86	103.88	17.31	82.94	13.82
12	230.39	38.40	101.05	16.84	88.98	14.83
13	218.78	36.46	100.63	16.77	84.03	14.00
14	218.34	36.39	102.67	17.11	83.79	13.97
15	218.26	36.38	106.43	17.74	84.16	14.03
16	219.51	36.59	104.59	17.43		
17			111.33	18.55		

512 Also, the leave-one-session-out LOOCV RMSE is

$$\begin{aligned} RMSE^S_{BASE} &= 150.34 \\ CV^S_{BASE} &= 25.06\% \end{aligned}$$

513 6.3.2. Modified Rational Expectations

Model F is similar to BASE, except that a parameter representing "perceived fundamental value" is now fitted with data.

$$\psi_j(\mathbf{X}) = \bar{f} \qquad (F)$$

514 where $\bar{f}$ is the parameter to be fitted.

515 Table 6 gives the leave-one-period-out LOOCV RMSE for each session.

Table 6. Forecast errors of the modified rational expectations model (F) for the three sessions.

	Session 1		Session 2		Session 3	
Part.	$RMSE_F^P$	$CV_F^P(\%)$	$RMSE_F^P$	$CV_F^P(\%)$	$RMSE_F^P$	$CV_F^P(\%)$
1	145.38	24.23	60.61	10.10	35.17	5.86
2	142.01	23.67	62.79	10.46	37.51	6.25
3	138.03	23.01	62.73	10.45	35.63	5.94
4	147.77	24.63	67.34	11.22	37.29	6.22
5	150.40	25.07	66.81	11.14	36.33	6.06
6	153.93	25.65	68.90	11.48	34.45	5.74
7	148.48	24.75	65.64	10.94	35.31	5.89
8	152.48	25.41	63.15	10.53	36.32	6.05
9	147.76	24.63	64.75	10.79	37.08	6.18
10	150.34	25.06	64.49	10.75	35.80	5.97
11	149.16	24.86	63.13	10.52	35.70	5.95
12	146.90	24.48	61.53	10.26	35.78	5.96
13	149.37	24.89	62.06	10.34	36.22	6.04
14	164.94	27.49	61.36	10.23	36.82	6.14
15	143.05	23.84	66.05	11.01	33.69	5.62
16	148.35	24.72	64.53	10.76		
17			60.75	10.12		

516 The leave-one-session-out LOOCV RMSE is given by

$$\begin{aligned} RMSE_F^S &= 109.04 \\ CV_F^S &= 18.17\% \end{aligned}$$

517 6.3.3. Simple Exponential Smoothing

518 Model ES uses simple exponential smoothing with initial price ratio of 1.

$$\begin{aligned} \psi_j(\mathbf{X}) &= X_{j-1} y_j && (ES) \\ y_j &= (1-\alpha) y_{j-1} + \alpha \frac{X_{j-1}}{X_{j-2}} && j = 3, \ldots, n \\ y_2 &:= 1 \end{aligned}$$

519 Table 7 gives the leave-one-period-out LOOCV RMSE for each session.

Table 7. Forecast errors of the simple exponential smoothing model (ES) for the three sessions.

	Session 1		Session 2		Session 3	
Part.	$RMSE^{P}_{ES}$	$CV^{P}_{ES}(\%)$	$RMSE^{P}_{ES}$	$CV^{P}_{ES}(\%)$	$RMSE^{P}_{ES}$	$CV^{P}_{ES}(\%)$
1	29.35	4.89	15.95	2.66	5.88	0.98
2	40.88	6.81	19.56	3.26	10.59	1.76
3	27.40	4.57	11.61	1.94	4.96	0.83
4	24.49	4.08	10.80	1.80	4.47	0.75
5	22.60	3.77	13.17	2.19	2.55	0.43
6	30.87	5.14	18.16	3.03	5.59	0.93
7	25.79	4.30	11.53	1.92	4.31	0.72
8	19.71	3.28	23.18	3.86	2.19	0.36
9	28.72	4.79	25.37	4.23	2.41	0.40
10	30.54	5.09	18.30	3.05	2.45	0.41
11	37.00	6.17	15.42	2.57	9.71	1.62
12	23.11	3.85	16.74	2.79	15.59	2.60
13	23.95	3.99	14.36	2.39	5.02	0.84
14	38.43	6.40	22.18	3.70	3.39	0.57
15	26.32	4.39	11.71	1.95	5.50	0.92
16	24.34	4.06	9.91	1.65		
17			20.86	3.48		

520 The leave-one-session-out LOOCV RMSE is given by

$$\begin{aligned} RMSE^{S}_{ES} &= 22.87 \\ CV^{S}_{ES} &= 3.81\% \end{aligned}$$

521 6.3.4. Brock and Hommes Model (BH)

We adapt the model of expectation formation used in Brock and Hommes [23] given in (3):

$$\hat{p}_{t+1} = \bar{p}_{t+1} + b_0 + b_1(p_{t-1} - \bar{p}_{t-1})$$

Note that the forecast in period $t+1$ is a linear function of the price in period $t-1$. However, in our experiment, at the beginning of period t, the participants entered the price forecast for that period. Therefore, we modify the BH model to

$$\hat{p}_t = \bar{p}_t + b_0 + b_1(p_{t-1} - \bar{p}_{t-1}) \qquad (BH)$$

522 Table 8 gives the leave-one-period-out LOOCV RMSE for each session.

523 The leave-one-session-out RMSE is

$$\begin{aligned} RMSE^{S}_{BH} &= 23.97 \\ CV^{S}_{BH} &= 3.99\% \end{aligned}$$

Table 8. Forecast errors of the Brock and Homme model (BH) for the three sessions.

	Session 1		Session 2		Session 3	
Part.	$RMSE_{BH}^{P}$	$CV_{BH}^{P}(\%)$	$RMSE_{BH}^{P}$	$CV_{BH}^{P}(\%)$	$RMSE_{BH}^{P}$	$CV_{BH}^{P}(\%)$
1	32.88	5.48	20.46	3.41	6.67	1.11
2	29.75	4.96	17.67	2.95	13.85	2.31
3	27.91	4.65	12.06	2.01	5.62	0.94
4	47.03	7.84	28.97	4.83	6.04	1.01
5	39.53	6.59	21.39	3.56	5.94	0.99
6	36.56	6.09	23.83	3.97	4.25	0.71
7	25.31	4.22	25.05	4.18	6.65	1.11
8	47.88	7.98	19.19	3.20	6.48	1.08
9	38.10	6.35	20.05	3.34	7.09	1.18
10	42.57	7.10	18.37	3.06	5.42	0.90
11	33.99	5.67	16.98	2.83	10.87	1.81
12	21.75	3.62	22.79	3.80	15.55	2.59
13	37.31	6.22	22.03	3.67	8.47	1.41
14	35.11	5.85	21.32	3.55	5.02	0.84
15	27.19	4.53	21.47	3.58	5.93	0.99
16	30.09	5.02	27.89	4.65		
17			16.53	2.75		

6.3.5. Exponential Smoothing with Quadratic H (ESH2)

The correction function is given by

$$H(\theta) = \frac{\theta + \rho(\theta - 1)^2}{1 + (\theta - 1)^2}$$

In this case, the model function ψ_j is given by

$$\psi_j = \bar{f} H\left(\frac{X_{j-1} y_j}{\bar{f}}\right) \qquad (ESH2)$$
$$y_j = (1-\alpha) y_{j-1} + \alpha \frac{X_{j-1}}{X_{j-2}} \qquad j = 3, \ldots, n_p$$

There are four parameters:

$$\begin{aligned} \alpha &= \text{exponential smoothing parameter} \\ y_2 &= \text{initial price ratio to start the exponential smoothing} \\ \rho &= \text{parameter for the } H \text{ function} \\ \bar{f} &= \text{fundamental value} \end{aligned}$$

Initial numerical results showed that the objective function was very flat with respect to y_2. This is because the effect of y_2 on the forecast goes down exponentially at rate $1 - \alpha$, and the data of the first 10 periods are used to prime the exponential smoothing method, but as discussed before, the data of the first 10 periods are not used in the squared error calculations. So unless α is very close to 0, y_2 has

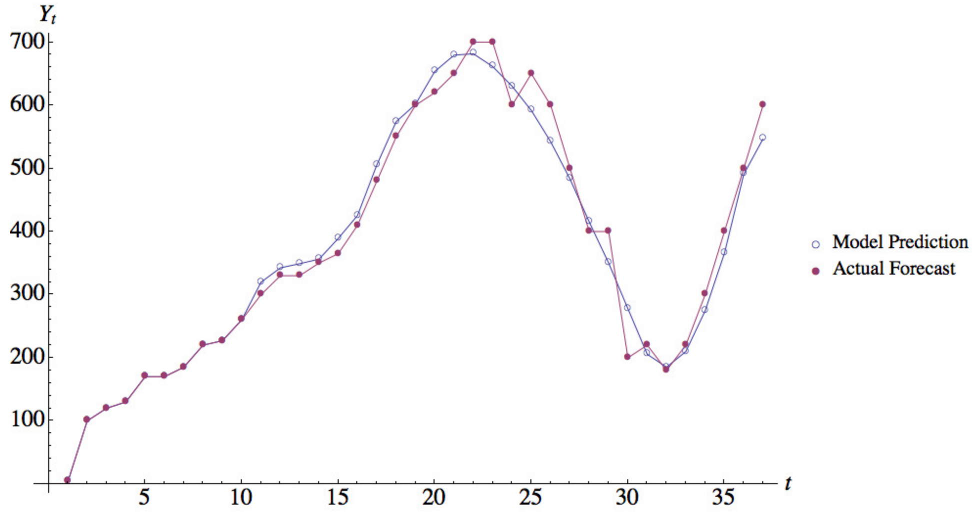

Figure 3. Data fit for Participant 14 in Session 1. Note that the first 10 periods were used for priming, hence are not included in the data fit.

very little effect on the objective. Therefore, y_2 was fixed at 1 and the remaining three parameters were estimated by minimizing the sum of squared errors.

Figure 3 shows the actual forecasts and the forecasts predicted by model ESH2 for Participant 14 in Session 1; this participant had the highest value of CV. It can be seen from the figure that even for this case, the forecasts predicted by model ESH2 match the actual forecasts reported by the participant quite well.

Table 9 gives the leave-one-period-out LOOCV RMSE for each session.

The leave-one-session-out LOOCV RMSE is

$$\begin{aligned} RMSE^S_{ESH2} &= 19.51 \\ CV^S_{ESH2} &= 3.25\% \end{aligned}$$

6.3.6. Exponential Smoothing with Non-Mononotic H (ESHE)

The correction function is given by

$$H(\theta) = \begin{cases} 1 + (1+\rho^2)\left(1 - \exp\left(-\frac{2\eta(\theta-1)(1-\rho)}{1+\rho^2}\right) + \rho\left(1 - \exp\left(-\frac{2\eta(\theta-1)}{(1+\rho^2)}\right)\right)\right) & \text{if } \theta \geq 1 \\ 1/H(1/\theta) & \text{if } \theta < 1 \end{cases}$$

The structure of the H function was chosen so that when $\rho \in [0,1]$, then H is increasing and when $\rho < 0$, then H is non-monotonic. Thus, a fitted value of ρ that is negative would indicate that the participant exhibits panicking behavior. Also, the parameter η is such that $H'(1) = 2\eta$.

Figure 4 shows members of this family of functions H.

Table 9. Forecast errors of the exponential smoothing with quadratic H model (ESH2) for the three sessions.

	Session 1		Session 2		Session 3	
Part.	$RMSE^P_{ESH2}$	$CV^P_{ESH2}(\%)$	$RMSE^P_{ESH2}$	$CV^P_{ESH2}(\%)$	$RMSE^P_{ESH2}$	$CV^P_{ESH2}(\%)$
1	17.61	2.93	13.80	2.30	5.52	0.92
2	21.41	3.57	15.21	2.54	10.63	1.77
3	20.59	3.43	9.95	1.66	5.09	0.85
4	22.40	3.73	10.53	1.76	4.72	0.79
5	16.39	2.73	12.48	2.08	2.33	0.39
6	24.12	4.02	16.22	2.70	4.31	0.72
7	25.80	4.30	10.40	1.73	4.00	0.67
8	20.01	3.34	16.78	2.80	1.80	0.30
9	20.91	3.48	19.04	3.17	2.51	0.42
10	29.96	4.99	11.62	1.94	2.46	0.41
11	25.87	4.31	10.74	1.79	9.38	1.56
12	27.55	4.59	15.36	2.56	14.02	2.34
13	16.84	2.81	13.98	2.33	5.06	0.84
14	39.95	6.66	16.96	2.83	3.55	0.59
15	26.45	4.41	9.34	1.56	4.29	0.72
16	16.77	2.79	9.04	1.51		
17			12.95	2.16		

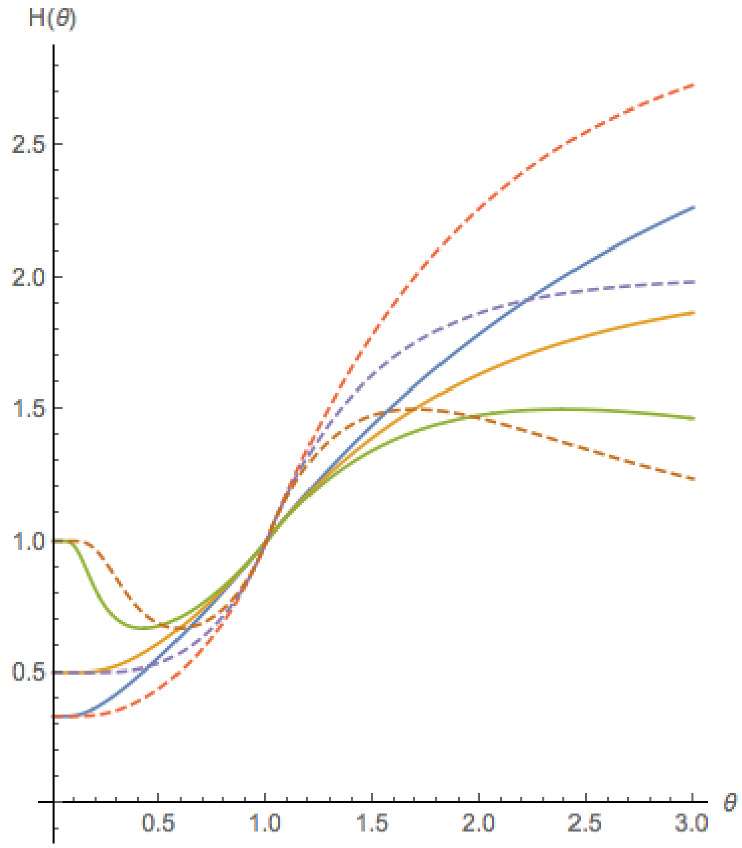

Figure 4. Members from the $ESHE$ family of functions. The solid lines are for $\eta = 1/2$ ($H'(1) = 1$) and the dotted lines are for $\eta = 1$ ($H'(1) = 2$). For each η, functions are plotted for $\rho = 1, 0, -1$. When $\rho < 0$, the function is non-monotonic, as defined in the main text.

The model function ψ_j is given by

$$\begin{aligned} \psi_j &= \bar{f}H\left(\frac{X_{j-1}y_j}{\bar{f}}\right) && (ESHE) \\ y_j &= (1-\alpha)y_{j-1} + \alpha\frac{X_{j-1}}{X_{j-2}} \qquad j = 3,\ldots,n \end{aligned}$$

There are five parameters:

$$\begin{aligned} \alpha &= \text{exponential smoothing parameter} \\ y_2 &= \text{initial price ratio to start the exponential smoothing} \\ \rho &= \text{monotonicity parameter for the } H\text{function} \\ \eta &= \text{slope paramter for the } H\text{function} \\ \bar{f} &= \text{fundamental value} \end{aligned}$$

In order to be consistent with fitting the ESH2 model, the y_2 parameter was fixed to 1. We used a trust region type method to fit the four parameters. Unfortunately the objective function for the parameter fitting problem has multiple local minima. We started from 10 random starting points and picked the solution that gave the best objective value. (For session 3, participant 13, we started from 100 random starting points as the algorithm terminated with a local solution for only some of the starting points.)

Table 10 gives the leave-one-period-out LOOCV RMSE for each session.

Table 10. Forecast errors of the exponential smoothing with non-monotonic H model (ESHE) for the three sessions.

	Session 1		Session 2		Session 3	
Part.	$RMSE^P_{ESHE}$	$CV^P_{ESHE}(\%)$	$RMSE^P_{ESHE}$	$CV^P_{ESHE}(\%)$	$RMSE^P_{ESHE}$	$CV^P_{ESHE}(\%)$
1	18.70	3.12	12.50	2.08	6.18	1.03
2	21.99	3.66	15.45	2.58	10.75	1.79
3	20.88	3.48	9.52	1.59	4.91	0.82
4	23.91	3.98	10.87	1.81	4.80	0.80
5	20.23	3.37	12.31	2.05	2.65	0.44
6	35.25	5.88	16.69	2.78	4.70	0.78
7	36.74	6.12	10.65	1.78	3.47	0.58
8	26.11	4.35	16.02	2.67	1.83	0.30
9	20.35	3.39	18.53	3.09	2.74	0.46
10	36.36	6.06	11.11	1.85	2.66	0.44
11	30.39	5.06	9.53	1.59	10.16	1.69
12	33.54	5.59	12.81	2.13	15.26	2.54
13	17.99	3.00	12.08	2.01	6.77	1.13
14	40.86	6.81	14.48	2.41	3.91	0.65
15	24.64	4.11	8.80	1.47	4.32	0.72
16	17.73	2.95	9.94	1.66		
17			12.51	2.08		

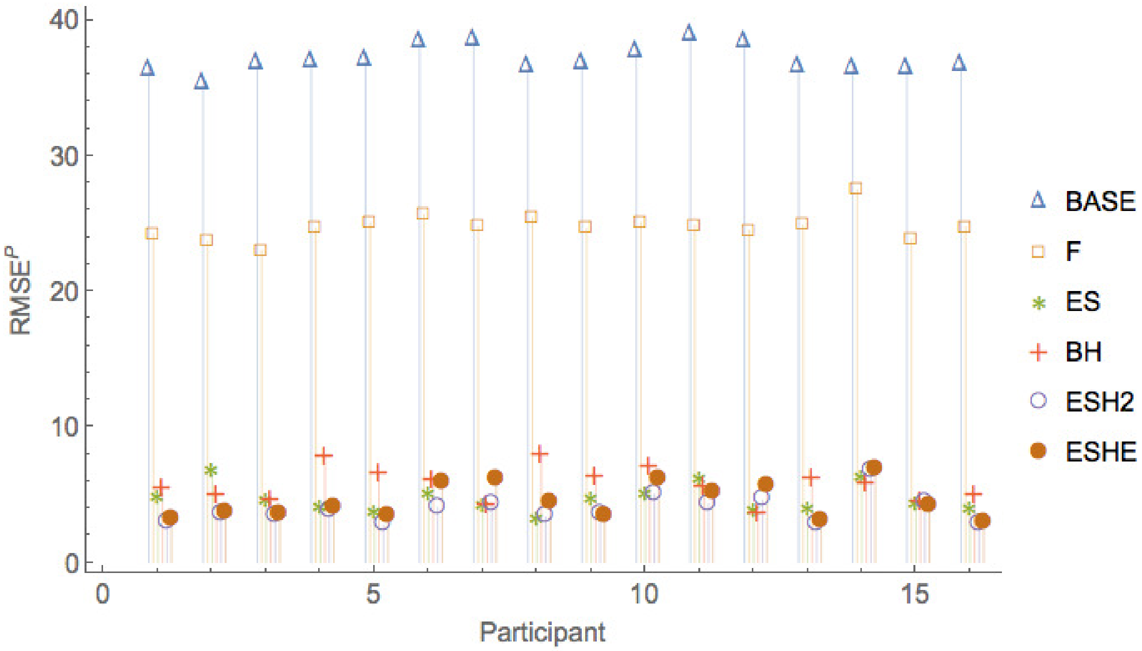

Figure 5. Comparison of leave-one-period-out LOOCV RMSE of all models for Session 1

The leave-one-session-out LOOCV RMSE is

$$\begin{aligned} RMSE^S_{ESHE} &= 19.25 \\ CV^S_{ESHE} &= 3.21\% \end{aligned}$$

It is also interesting to note that fitting a single set of parameters to all the participants in all the sessions gave an $RMSE = 16.07$ with the corresponding CV of 2.68%. Thus, the behavior of the group can be described fairly well by a single set of parameters.

6.4. Comparison of Various Models

Figures 5, 6 and 7 show a comparison of the leave-one-period-out LOOCV RMSE of the various models. It can be seen that the pure rational expectations model (model $BASE$) has the highest errors throughout. When we allow the fundamental value to be fitted (model F), the errors are reduced, but the errors are still much larger than for the other models. The simple exponential smoothing model (model ES) captures the participant forecasts remarkably well for a one-parameter model.

Figures 8, 9 and 10 show a comparison of the leave-one-period-out LOOCV RMSE of models ES, BH, $ESH2$, $ESHE$ only, that is, leaving out the two worst performing models $BASE$ and F.

The leave-one-session-out LOOCV RMSE of the various models are given in Table 11. As is the case for the leave-one-period-out LOOCV RMSE, the rational expectations model ($BASE$) has the highest leave-one-session-out LOOCV RMSE, and model F has slightly smaller leave-one-session-out LOOCV RMSE. It can be seen that an exponential smoothing model with a single parameter for all sessions and all participants captures much of the variation in the observed forecasts. More

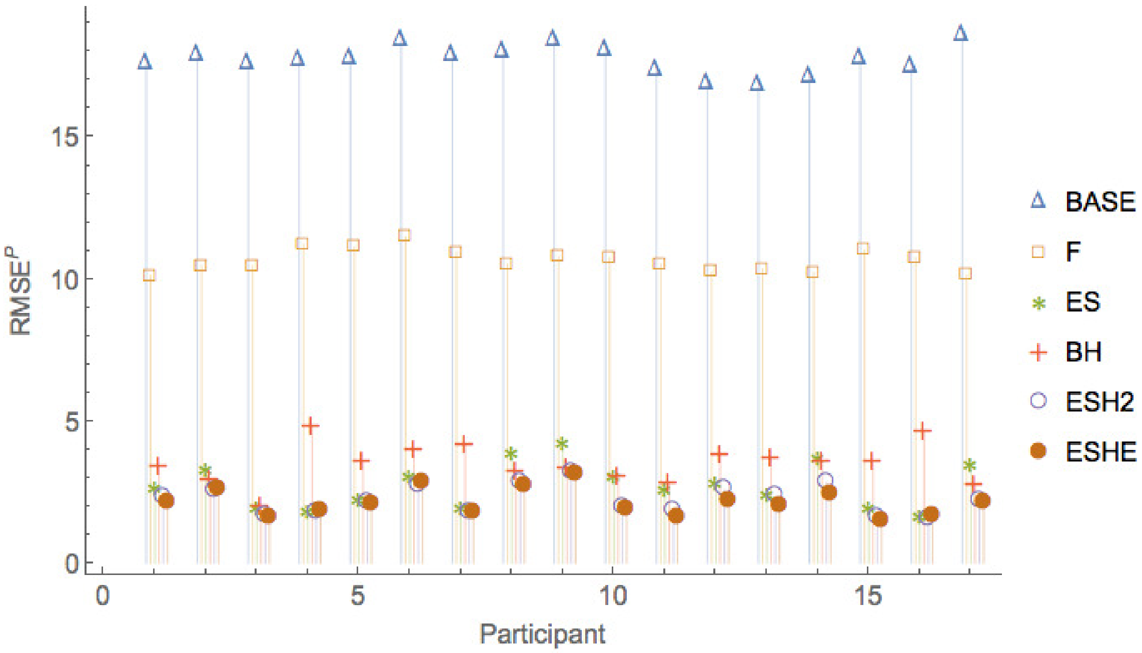

Figure 6. Comparison of leave-one-period-out LOOCV RMSE of all models for Session 2

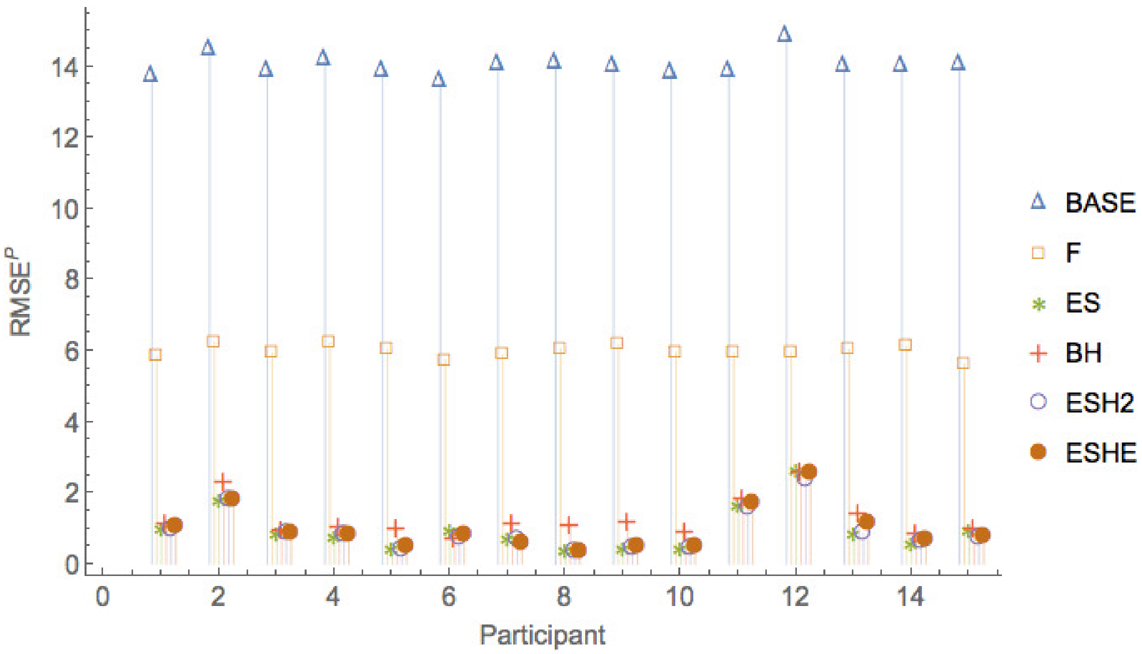

Figure 7. Comparison of leave-one-period-out LOOCV RMSE of all models for Session 3

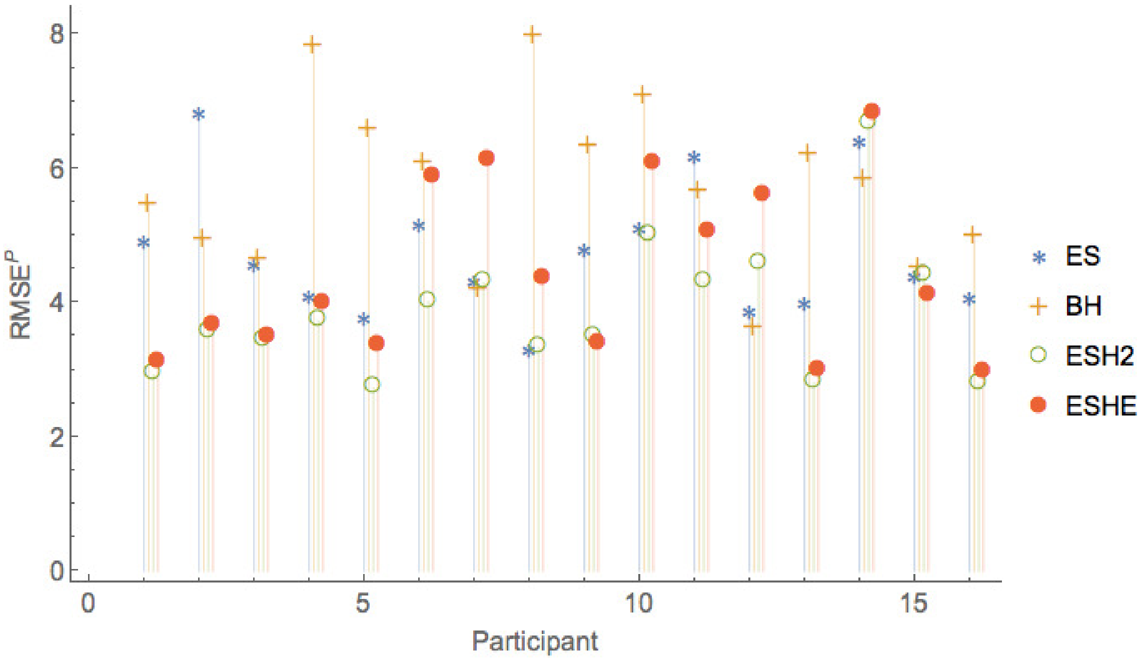

Figure 8. Comparison of leave-one-period-out LOOCV RMSE of selected models for Session 1

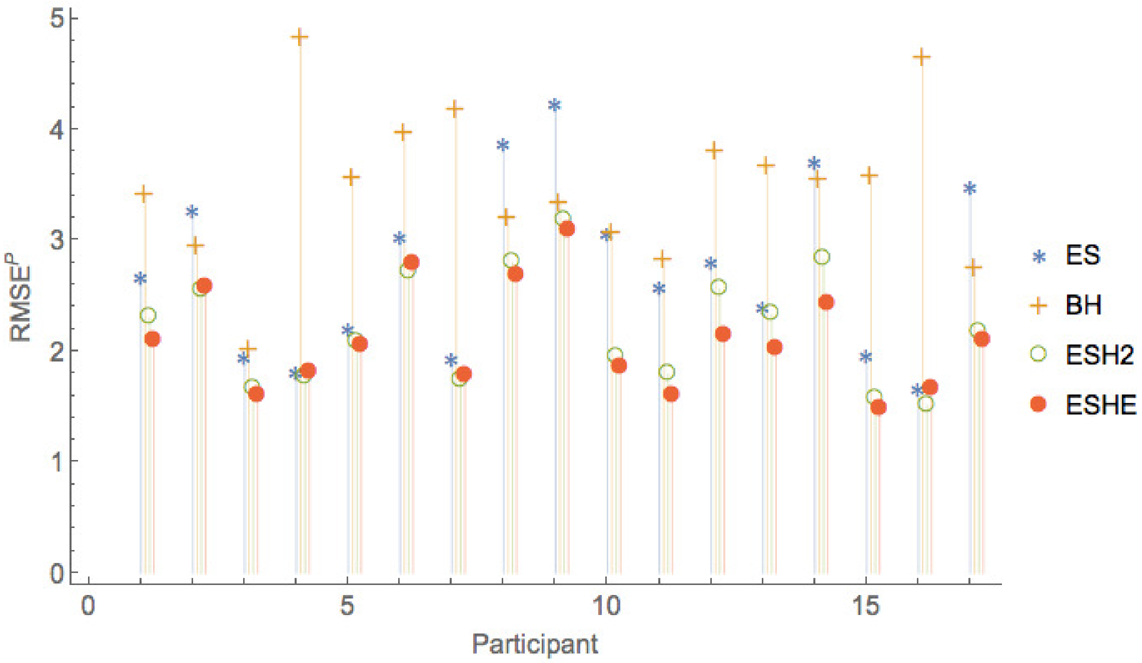

Figure 9. Comparison of leave-one-period-out LOOCV RMSE of selected models for Session 2

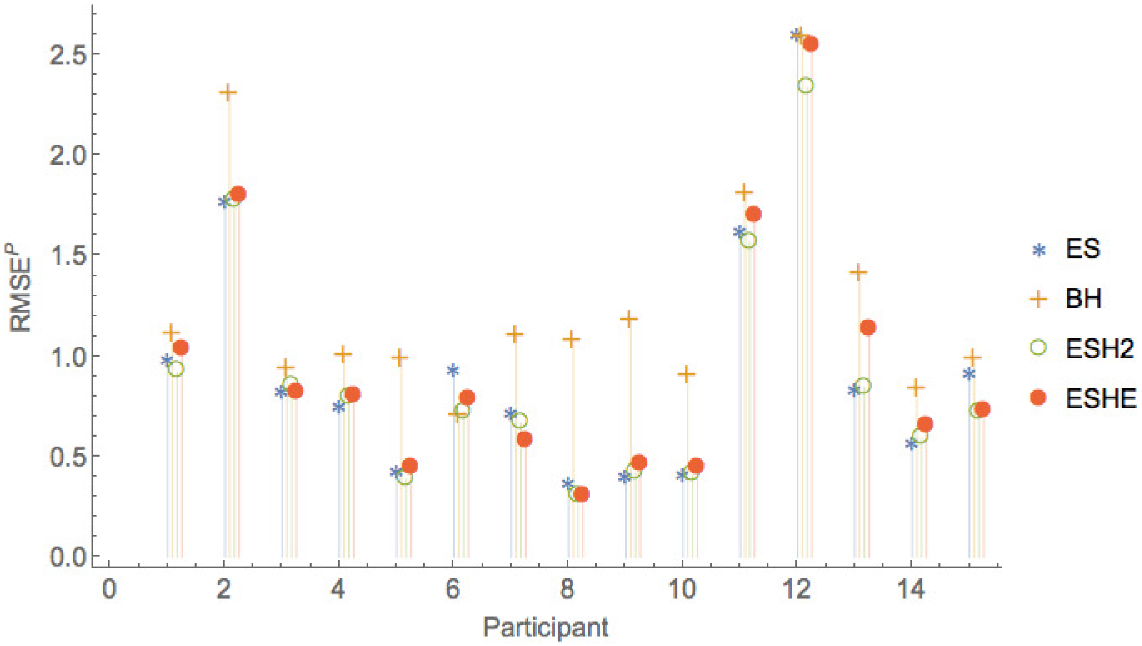

Figure 10. Comparison of leave-one-period-out LOOCV RMSE of selected models for Session 3

sophisticated models with additional parameters ($ESH2$ and $ESHE$) provide only small improvements in the leave-one-session-out LOOCV RMSE over what can be obtained with the ES model.

Table 11. Comparison of leave-one-session-out LOOCV RMSE of various models.

Model	No. of Parameters	$RMSE^S$	$CV^S(\%)$
BASE	0	150.34	25.06
F	1	109.04	18.17
ES	1	22.87	3.81
BH	2	23.97	3.99
ESH2	3	19.51	3.25
ESHE	4	19.25	3.21

7. Interpretation and Implications of Results

In Cheriyan and Kleywegt [12] the dynamical system associated with the price process was studied both analytically and numerically. It was shown there that the qualitative behavior of the trajectories of this dynamical system is determined by the parameter α and the nature of the H function. These can also be related to the investors' behavioral characteristics. Although the purpose of our experiment was to calibrate models of investor forecasting behavior and not to calibrate the dynamical system itself, in this section we make observations regarding the fitted models of investor forecasting behavior and the qualitative behavior of the dynamical system.

The parameter α captures the investors' memory — the weight that they put on the most recent observed price ratio. Given an H function and a fundamental dividend to price ratio δ (or price-earnings ratio $1/\delta$), a critical value of α is given by

$$\alpha^* = \frac{1 + \delta - H'(1)}{H'(1)}.$$

The slope of the H function at 1 captures investor confidence. If $H'(1) < 1$, it denotes ***cautious confidence***, and if $H'(1) > 1$, it denotes ***excessive exuberance***. It can be shown that if $\alpha < \alpha^*$, then the fundamental value corresponds to a stable attracting point, that is, price trajectories that start in a neighborhood of the fundamental value, converge to the fundamental value. Numerical evidence suggests that the same is true for all trajectories, that is, if $\alpha < \alpha^*$, then all price trajectories eventually converge to the fundamental value. If $\alpha > \alpha^*$, then numerical results in Cheriyan and Kleywegt [12] show that a price cycle appears that attracts all trajectories. This bifurcation process is continuous in the sense that the price cycles are small if α is slightly larger than α^*, and the price cycles gradually grow larger as α increases from α^*.

For α not too far from α^*, the price cycles appear to be predictable, that is, the price trajectories converge to a smooth curve, and any two trajectories that start close to each other will remain close to each other at all times. If the H function is monotonic, this behavior persists for $\alpha > \alpha^*$. However, if H is non-monotonic, then for larger values of α, the price trajectories can be non-predictable. This means that the trajectories do not converge to a smooth curve, and they exhibit sensitive dependence on initial conditions, that is, two trajectories that start close to each other will grow apart exponentially fast. A non-monotonic H captures what we called panicking behavior of the investor — as the extrapolation forecast increases beyond the fundamental value, the price forecast initially increases (*confidence* or *exuberance*), but beyond some value the price forecast decreases and moves closer to the fundamental value (*panic*).

Table 12 shows the fitted parameters for the $ESH2$ model. For this function $H'(1) = 1$ and $\alpha^* = \delta = 1/60$. The fitted α-values are all greater than the critical α^*, and this is consistent with the price cycles observed in the three sessions.

Table 13 shows the fitted parameters for the $ESHE$ model. It can be seen that the α parameter is much larger than the critical α^*. Once again, this is consistent with the observed price cycles. The slopes of the fitted H functions at the fundamental value, $H'(1)$, are all close to 1. Some participants exhibit less exuberance ($H'(1) \in (0, 1)$), and some exhibit more excessive exuberance ($H'(1) > 1$). All the fitted values for ρ are positive — thus we do not find evidence of panicking behavior in the data.

Table 14 summarizes the fitted parameter values for the Brock and Homme model (BH). Brock and Hommes [23] calls b_0 and b_1 the "bias" and "trend" parameters respectively. In particular, if $b_1 < 0$, then the investor is called a "contrarian", and if $b_1 > 0$, then the investor is called a "trend chaser". If $b_0 = 0$ and $b_1 = 0$, then the investor is a "fundamentalist". The fitted parameter values indicate that all the participants in our experiment were trend chasers.

8. Conclusions

We designed an experiment to study investors' price forecast formation in the context of a market for an investment asset. Many experiments in the literature use trading runs with a pre-announced

Table 12. Fitted parameters for quadratic H model ($ESH2$) for the three sessions.

	Session 1				Session 2				Session 3			
Part.	α	ρ	$\bar{f}$	CV(%)	α	ρ	$\bar{f}$	CV(%)	α	ρ	$\bar{f}$	CV(%)
1	0.35	0.01	172.32	2.6	0.73	0.18	365.69	1.94	0.6	2.26	543.41	0.78
2	0.24	0.52	407.91	3.03	0.43	0.36	395.39	2.34	0.49	0.03	312.63	1.7
3	0.29	0.01	141.31	3.12	0.15	0.07	317.4	1.54	0.56	1.69	493.54	0.76
4	0.99	0.01	186.84	3.52	0.99	0.11	413.81	1.71	0.99	0.61	488.64	0.72
5	0.49	0.01	179.32	2.51	0.62	0.19	405.06	1.9	0.72	0.08	697.33	0.37
6	0.39	0.01	167.4	3.86	0.5	0.01	214.67	2.48	0.09	1.69	452.22	0.6
7	0.17	0.01	146.57	3.87	0.99	0.03	298.27	1.64	0.84	1.37	475.69	0.58
8	0.99	0.01	250.59	3.23	0.44	0.23	370.67	2.52	0.49	0.89	472.42	0.25
9	0.38	0.28	371.52	2.89	0.49	0.1	321.78	2.83	0.94	0.01	816.76	0.37
10	0.47	0.01	168.01	4.01	0.35	0.06	302.52	1.78	0.71	0.19	643.85	0.31
11	0.25	0.01	142.26	3.96	0.42	0.36	414.8	1.62	0.51	1.18	569.84	1.53
12	0.19	0.01	143.75	4	0.99	1.06	473.49	2.24	0.99	3.17	460.64	2.21
13	0.35	0.01	184.62	2.58	0.96	1.08	490.54	2.04	0.99	1.95	498.14	0.78
14	0.19	0.08	448.71	5.43	0.51	0.79	444.85	2.43	0.47	0.06	419.71	0.53
15	0.23	0.42	394.28	3.55	0.57	0.12	370.91	1.37	0.52	3.45	482.54	0.6
16	0.35	0.01	175.67	2.67	0.99	0.05	351.82	1.47				
17					0.34	0.09	291.93	1.99				

finite number of periods. However, the known end-of-horizon seems to affect participants' forecasts, apparently via reasoning involving backward induction from the end of the horizon, that is not representative of forecasting in actual asset markets. In our experiment, we emulated an infinite horizon with discounting by stopping the trading run in each period with a pre-announced stopping probability. We conducted three experimental sessions with one trading run each. The equilibrium prices in all three trading runs exhibited cycles.

We fit a number of models of expectation formation to the data. The fit for the rational base case $(BASE)$ indicates that the rational expectations model does not provide an accurate model of investor forecasting behavior. Even when the fundamental value was replaced by a parameter that was fitted with the data (model F), the accuracy of the model did not improve much. (For example, the CV of the fit decreased from 22% to 18%.) In contrast, a one-parameter exponential smoothing model (ES) gave remarkably accurate predictions of investor forecasts for such a simple model (with CV around 4%). Thus the evidence indicates that the investors, despite being reminded of the importance of the fundamental value of an investment asset and being told explicitly what the fundamental value was, resorted mostly to extrapolating from the past price data. Models with a larger number of parameters provided a slightly better fit. Moreover, it can be shown that these models are able to explain price cycles and more complex price trajectories in addition to price bubbles, see Cheriyan and Kleywegt [12] for details. For every participant, the fitted value of α was larger than the critical value α^* for the associated dynamical system, which is consistent with the price cycles observed in the sessions. The parameter fits also indicated some amount of overconfidence. The data did not provide evidence of panicking behavior.

An interesting observation was the gap between theoretical knowledge and internalized knowledge. For example, based on correct answers to questions before the experiment, we concluded

Table 13. Fitted parameters for exponential smoothing with non-monotonic H model (ESHE) for the three sessions.

	Session 1				Session 2				Session 3			
Part.	α	α^*	ρ	$H'(1)$	α	α^*	ρ	$H'(1)$	α	α^*	ρ	$H'(1)$
1	0.349	-0.003	0.874	1.019	0.604	-0.101	0.800	1.131	0.012	-0.122	1.000	1.158
2	0.246	-0.051	0.365	1.071	0.433	-0.086	1.000	1.113	0.445	0.002	0.916	1.014
3	0.300	0.031	0.871	0.986	0.172	0.102	0.928	0.922	0.550	-0.069	0.905	1.092
4	0.990	-0.060	0.862	1.082	0.990	-0.036	0.915	1.055	0.990	-0.023	0.393	1.041
5	0.470	-0.052	0.872	1.072	0.595	-0.046	0.876	1.065	0.610	-0.028	0.253	1.046
6	0.463	-0.369	0.810	1.610	0.550	-0.163	0.847	1.215	0.336	-0.117	0.807	1.152
7	0.182	-0.008	0.872	1.024	0.990	0.002	0.917	1.015	0.874	-0.074	0.362	1.098
8	0.990	-0.087	0.866	1.114	0.439	-0.070	0.568	1.094	0.478	-0.013	0.484	1.030
9	0.384	-0.048	0.945	1.068	0.475	-0.034	0.731	1.053	0.901	-0.022	0.961	1.040
10	0.457	-0.035	0.873	1.054	0.355	0.080	0.908	0.942	0.667	0.020	0.946	0.997
11	0.248	-0.042	0.858	1.061	0.423	-0.110	0.811	1.142	0.393	-0.065	0.570	1.088
12	0.205	0.003	0.873	1.013	0.990	-0.159	0.116	1.209	0.990	-0.131	0.846	1.170
13	0.342	-0.048	0.868	1.067	0.686	-0.146	0.805	1.190	0.010	-0.135	0.937	1.176
14	0.194	-0.161	1.000	1.212	0.479	-0.148	1.000	1.194	0.422	-0.012	0.948	1.029
15	0.230	0.016	0.886	1.001	0.547	-0.040	0.849	1.059	0.568	-0.113	0.255	1.146
16	0.347	-0.037	0.870	1.056	0.990	0.040	0.938	0.977				
17					0.343	-0.022	0.746	1.039				

Table 14. Summary of fitted slope parameters for model BH.

Session	min b_0	avg b_0	max b_0	min b_1	avg b_1	max b_1
1	-12.87	4.78	45.11	0.95	1.02	1.14
2	-12.74	-1.69	4.35	0.93	0.98	1.05
3	-6.63	4.61	8.59	1.01	1.09	1.15

that the participants understood the meaning of the memoryless property of the geometric distribution and the computation of fundamental value. Nevertheless, their answers to questions and their trading behavior during the experiment seemed to indicate that they did not believe the theoretical properties. For example, the total number of periods in a trading run was a geometric random variable, and this was explained to the participants together with a reminder of the memoryless property of the geometric distribution. However, at the beginning of each period, each participant was asked to give the expected number of periods remaining in the trading run, and few participants gave the correct answer. It would be interesting to design an experiment that could lead to a better understanding of this apparent gap between theoretical and internalized knowledge.

Bibliography

1. Armstrong, J.S. *Long-Range Forecasting*, 2nd. ed.; John Wiley & Sons, 1985.
2. Ramnath, S.; Rock, S.; Shane, P. The Financial Analyst Forecasting Literature: A Taxonomy with Suggestions for Further Research. *International Journal of Forecasting* **2008**, *24*, 34–75.
3. LeRoy, S.F.; Porter, R.D. The Present-Value Relation: Tests Based on Implied Variance Bounds. *Econometrica* **1981**, *49*, 555–574.
4. Shiller, R.J. Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends? *The American Economic Review* **1981**, *71*, 421–436.

5. Shiller, R.J. The Use of Volatility Measures in Assessing Market Efficiency. *The Journal of Finance* **1981**, *36*, 291–304.
6. Shiller, R.J. The Marsh-Merton Model of Managers' Smoothing of Dividends. *The American Economic Review* **1986**, *76*, 499–503.
7. Campbell, J.Y.; Shiller, R.J. Cointegration and Tests of Present Value Models. *The Journal of Political Economy* **1987**, *95*, 1062–1088.
8. Campbell, J.Y.; Shiller, R.J. The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors. *The Review of Financial Studies* **1988**, *1*, 195–228.
9. West, K.D. Dividend Innovations and Stock Price Volatility. *Econometrica* **1988**, *56*, 37–61.
10. West, K.D. Bubbles, Fads and Stock Price Volatility Tests: A Partial Evaluation. *The Journal of Finance* **1988**, *43*, 639–656.
11. LeRoy, S.F. Efficient Capital Markets and Martingales. *Journal of Economic Literature* **1989**, *27*, 1583–1621.
12. Cheriyan, V.; Kleywegt, A.J. A dynamical systems model of price bubbles and cycles. *Quantitative Finance* **2016**, *16*, 309–336.
13. Smith, V.; Suchanek, G.; Williams, A. Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica: Journal of the Econometric Society* **1988**, *56*, 1119–1151.
14. Ball, S.B.; Holt, C.A. Classroom games: Speculation and bubbles in an asset market. *The Journal of Economic Perspectives* **1998**, *12*, 207–218.
15. Camerer, C.F.; Weigelt, K. Convergence in Experimental Double Auctions in Stochastically lived Assets. In *The Double Auction Market: Institutions, Theories and Evidence*; 1993; pp. 355–359.
16. Hirota, S.; Sunder, S. Price bubbles sans dividend anchors: Evidence from laboratory stock markets. *Journal of Economic Dynamics and Control* **2007**, *31*, 1875–1909.
17. Haruvy, E.; Lahav, Y.; Noussair, C.N. Traders' Expectations in Asset Markets: Experimental Evidence. *The American Economic Review* **2007**, *97*, 1901–1920.
18. Lugovskyy, V.; Puzzello, D.; Tucker, S. An Experimental Study of Bubble Formation in Asset Markets Using the Tâtonnement Pricing Mechanism. Working Paper.
19. Brock, W.A.; Hommes, C.H. A Rational Route to Randomness. *Econometrica* **1997**, *65*, 1059–1095.
20. Tversky, A.; Kahneman, D. Judgment under Uncertainty: Heuristics and Biases. *Science* **1974**, *185*, 1124–1131.
21. Scheinkman, J.A.; Xiong, W. Overconfidence and Speculative Bubbles. *Journal of Political Economy* **2003**, *111*, 1183–1219.
22. Chuang, W.I.; Lee, B.S. An empirical evaluation of the overconfidence hypothesis. *Journal of Banking & Finance* **2006**, *30*, 2489–2515.
23. Brock, W.A.; Hommes, C.H. Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model. *Journal of Economic Dynamics and Control* **1998**, *22*, 1235–1274.

Appendix A. Fundamental Value Computations Corresponding to Different Dividend Processes

Appendix A.1. Deterministic Constant Dividend

Recall that in each period of the market, trades are made first and then the dividend is paid out. The dividend for each time period is a constant d. If the realization of a Geometric(p) random variable is a success, the salvage value s is paid out for each unit of stock held, otherwise, a new period starts. Algorithm 1 gives the details of the market algorithm for this case. Since the actual duration of the experiment is of the order of hours, we assume the discount factor is 0.

Algorithm 1 Market Algorithm with Deterministic Constant Dividend

1. $t \leftarrow 1$
2. For each participant, initialize account with initial stock and initial cash
3. Begin period t
4. Capture forecast data for period
5. Conduct trades
   (a) Participants enter buy and sell orders
   (b) Perform market clearing to determine market clearing price
   (c) Execute trades and update cash and stock accounts for each participant
6. For each participant, distribute dividend d for each unit of stock
7. Generate $\xi_{t+1} \sim Bern(p)$
   (a) IF $\xi_{t+1} = 1$
      i. For each participant, distribute s for each unit of stock
      ii. STOP
   (b) ELSE
      i. End period t. $t \leftarrow t + 1$
      ii. Proceed to step 3

Lemma 1. *In the setting of Market Algorithm 1, the fundamental value of a unit of stock is constant at each period and is given by*

$$\bar{p} = \frac{d}{p} + s$$

Proof. Let ξ_t be the random variable that is 1 if the market is running in period t and is 0 otherwise. Then $\{\xi_t\}_{t=1}^{\infty}$ is a sequence of iid Bernouilli($1 - p$) random variables. Let D_t be the dividend in period t. We have that

$$D_t = \xi_1 \xi_2 \ldots \xi_t d$$

Note that if any of the ξ_i's are zero, then the entire expression is zero, this automatically captures the fact that if the market stops in period i, then for $t \geq i$, $D_t = 0$. Also, for the present period, we know that the dividend is certain to be d, that is

$$D_0 = d$$

The dividend stream at period t is given by

$$\begin{aligned} D &= \sum_{i=0}^{\infty} D_i \\ \implies \mathbb{E}[D] &= \sum_{i=0}^{\infty} \mathbb{E}[D_i] \end{aligned}$$

Now,

$$\begin{aligned} \mathbb{E}[D_i] &= \mathbb{E}[\xi_1 \xi_2 \ldots \xi_i d] \\ &= \mathbb{E}[\xi_1]^i d \qquad (\xi_i \text{ iid}) \\ &= (1-p)^i d \end{aligned}$$

Therefore,

$$\mathbb{E}[D] = \sum_{i=0}^{\infty}(1-p)^i d = \frac{d}{p}$$

Now the dividend stream is

$$\begin{aligned} D &= \sum_{i=0}^{\infty} D_i \\ \implies \mathbb{E}[D] &= \sum_{i=0}^{\infty} \mathbb{E}[D_i] \end{aligned}$$

Now,

$$\begin{aligned} \mathbb{E}[D_i] &= \mathbb{E}[\xi_1\xi_2\ldots\xi_i d] \\ &= \mathbb{E}[\xi_1]^i d \qquad (\xi_i \text{ iid}) \\ &= (1-p)^i d \end{aligned} \tag{4}$$

Therefore,

$$\mathbb{E}[D] = \sum_{i=0}^{\infty}(1-p)^i d = \frac{d}{p}$$

which is independent of t. Since a unit of stock held for ever necessarily will result in a final payout of s, and since the discount factor is 0, the fundamental value of a unit of stock is given by

$$\overline{p} = \frac{d}{p} + s$$

□

Remark 1. Equation (4) says that, in the case of deterministic dividends, probabilistic stopping of the market with stopping probability p is equivalent to an infinitely lived marked with discount factor $1-p$.

Appendix A.2. Markov Dividends

In the case of Markov Dividends, the support of the dividend distribution in a period t depends on the market state in that period. The market states $X_t \in \{1, 2\}$ evolve according to a Markov chain with transition matrix P. Let π denote the stationary distribution corresponding to transition matrix P. In our experiment, the initial state X_0 was drawn from the distribution π.

Let the matrix Q denote the conditional p.m.f. for the dividend

$$q_{ij} = P(\mathfrak{d}_t = j | X_t = i)$$

and

Algorithm 2 gives the details of the market algorithm in the case of Markov dividends.

Algorithm 2 Market Algorithm with Deterministic Constant Dividend

Given:

- p, the stopping probability
- π, the initial probability distribution for market state
- $P = [p_{ij}]$, the state transition matrix for the Markov chain; $p_{ij} = P(X_t = j | X_{t-1} = i)$
- $Q = [q_{ij}]$, the dividend distribution; $q_{ij} = P(\mathfrak{d}_t = j | X_t = i)$

Algorithm:

1. $t \leftarrow 1$
2. For each participant, initialize account with initial stock and initial cash
3. Begin period t
4. Generate X_t. IF $t = 1$, generate $X_1 \sim \pi$, ELSE generate $X_t \sim p_{X_{t-1}}$
5. Capture for forecast data for period
6. Conduct trades
   (a) Participants enter buy and sell orders
   (b) Perform market clearing to determine market clearing price
   (c) Execute trades and update cash and stock accounts for each participant
7. Generate dividends $\mathfrak{d}_t \sim q_{X_t}$
8. For each participant, distribute dividend $\mathfrak{d}_t$ for each unit of stock
9. Generate $\xi_{t+1} \sim Bern(p)$
   (a) IF $\xi_{t+1} = 1$
      i. For each participant, distribute s for each unit of stock
      ii. STOP
   (b) ELSE
      i. End period t. $t \leftarrow t + 1$
      ii. Proceed to step 3

Participants have complete information about the parameters of the Markov Chain (P, Q, π) and but they do not know the underlying state process $\{X_t\}$. They know that the market started in the steady state $\pi = [\pi_1, \pi_2]$, that is X_0 is chosen according to π.

They observe the prices $p_{t-1}, p_{t-2}, \ldots$ and the dividends $d_{t-1}, d_{t-2}, \ldots$ Let $\mathbf{d}_{t-1} = (d_1, \ldots, d_{t-1})$ denote the history of dividends up to and including period $t-1$.

Let $\hat{v}_{x,t}(\mathbf{d}_t) = P(X_{t+1} = x | \mathbf{d}_t)$. That is, $\hat{v}_{x,t}(\mathbf{d}_t)$ is the estimate at the end of period t (i.e. beginning of period $t+1$) that the probability of the state X_{t+1} will be $x \in \{1, 2, \}$. Let $\hat{v}_t(\mathbf{d}_t) = [\hat{v}_{1,t}(\mathbf{d}_t), \hat{v}_{2,t}(\mathbf{d}_t)]$. In each period, the investor updates $\hat{v}_t$ and uses it to compute the fundamental value for the next period.

The estimate of the probability distribution of the state is given by

$$
\begin{aligned}
\hat{v}_0 &= \left[\begin{array}{cc} P(X_1 = 1) & P(X_1 = 2) \end{array} \right] = [\pi_1 \pi_2] \\
\hat{v}_{x,t}(\mathbf{d}_t) &= \frac{\sum\limits_{x_{t-1}} p_{x_{t-1} x} q_{x d_t} \hat{v}_{x_{t-1}, t-1}(\mathbf{d}_{t-1})}{\sum\limits_{x_t} \sum\limits_{x_{t-1}} p_{x_{t-1} x_t} q_{x_t d_t} \hat{v}_{x_{t-1}, t-1}(\mathbf{d}_{t-1})} \\
&= \frac{\left(\hat{v}_{t-1}(\mathbf{d}_{t-1})^T P\right)_x q_{x d_t}}{\left(\hat{v}_{t-1}(\mathbf{d}_{t-1})^T P Q\right)_{d_t}}
\end{aligned}
$$

723 Then, the fundamental value at period t can be computed as follows:

$$\begin{aligned}\bar{p}_t(\mathbf{d}_{t-1}) &= \sum_{k=0}^{\infty} \mathbb{E}[D_{t+k}|\zeta_{t+k}=1,\mathbf{d}_{t-1}] + s \\ &= \sum_{k=0}^{\infty} \mathbb{E}[\mathfrak{d}_{t+k}|\mathbf{d}_{t-1}] + s\end{aligned}$$

724 Now, for $k \geq 0$,

$$\begin{aligned}\mathbb{E}[\mathfrak{d}_{t+k}|\mathbf{d}_{t-1}] &= \sum_{x_{t+k}} \mathbb{E}\left[\mathfrak{d}_{t+k}|X_{t+k}=x_{t+k}\right] P(X_{t+k}=x_{t+k}|\mathbf{d}_{t-1}) \\ &= \sum_{x_{t+k}} \mathbb{E}\left[\mathfrak{d}_{t+k}|X_{t+k}=x_{t+k}\right] \sum_{x_t} P(X_{t+k}=x_{t+k}|X_t=x_t) P(X_t=x_t|\mathbf{d}_{t-1}) \\ &= \sum_{x_{t+k}} d_{x_{t+k}} \sum_{x_t} p^{(k)}_{x_t x_{t+k}} \hat{v}_{x_t,t-1}(\mathbf{d}_{t-1}) \\ &= \sum_{x_{t+k}} d_{x_{t+k}} (\hat{v}_{t-1}(\mathbf{d}_{t-1}) P^k)_{x_{t+k}} \\ &= \hat{v}_{t-1}(\mathbf{d}_{t-1}) P^k \begin{bmatrix} d_1 \\ d_2 \end{bmatrix}\end{aligned}$$

725 Therefore

$$\begin{aligned}\bar{p}_t(\mathbf{d}_{t-1}) &= \sum_{k=0}^{\infty} \mathbb{E}[\mathfrak{d}_{t+k}|\mathbf{d}_{t-1}] + s \\ &= \sum_{k=0}^{\infty} (1-p)^k \hat{v}_{t-1}(\mathbf{d}_{t-1}) P^k \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + s \\ &= \hat{v}_{t-1}(\mathbf{d}_{t-1}) \left(\sum_{k=0}^{\infty} (1-p)^k P^k \right) \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + s \\ &= \hat{v}_{t-1}(\mathbf{d}_{t-1}) \left(I - (1-p)P \right)^{-1} \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + s\end{aligned}$$

The next period's fundamental value is given by

$$\bar{p}_{t+1}(\mathbf{d}_{t-1}) = \sum_{k=0}^{\infty} (1-p)^k \mathbb{E}[\mathfrak{d}_{t+1+k}|\mathbf{d}_{t-1}] + s$$

726 where

$$\mathbb{E}[\mathfrak{d}_{t+1+k}|\mathbf{d}_{t-1}] \quad = \mathbb{E}[\mathfrak{d}_{t+(k+1)}|\mathbf{d}_{t-1}] = \quad \hat{v}_{t-1}(\mathbf{d}_{t-1}) P^{k+1} \begin{bmatrix} d_1 \\ d_2 \end{bmatrix}$$

727 following similar calculations as above. Thus,

$$
\begin{aligned}
\bar{p}_{t+1}(\mathbf{d}_{t-1}) &= \sum_{k=0}^{\infty}(1-p)^k \hat{v}_{t-1}(\mathbf{d}_{t-1}) P^{k+1} \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + s \\
&= \hat{v}_{t-1}(\mathbf{d}_{t-1}) P \left(\sum_{k=0}^{\infty}(1-p)^k P^k \right) \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + s \\
&= \hat{v}_t(\mathbf{d}_t) P \left(I - (1-p)P\right)^{-1} \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + s
\end{aligned}
$$

728 This was the fundamental value displayed to the participants.

729 Next,

$$
\begin{aligned}
\bar{p}_{t-1}(\mathbf{d}_{t-1}) &= \sum_{k=0}^{\infty}(1-p)^k \mathbb{E}[\mathfrak{d}_{t-1+k}|\mathbf{d}_{t-1}] + s \\
&= d_{t-1} + \sum_{k=1}^{\infty}(1-p)^k \mathbb{E}[\mathfrak{d}_{t+k-1}|\mathbf{d}_{t-1}] + s \\
&= d_{t-1} + \sum_{k=1}^{\infty}(1-p)^k \hat{v}_{t-1}(\mathbf{d}_{t-1}) P^{k-1} \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + s \\
&= d_{t-1} + (1-p) \sum_{k=1}^{\infty}(1-p)^{k-1} \hat{v}_{t-1}(\mathbf{d}_{t-1}) P^{k-1} \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + s \\
&= d_{t-1} + (1-p) \hat{v}_{t-1}(\mathbf{d}_{t-1}) \left(\sum_{k=0}^{\infty}(1-p)^k P^k \right) \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + s \\
&= d_{t-1} + (1-p) \hat{v}_{t-1}(\mathbf{d}_{t-1}) \left(I - (1-p)P\right)^{-1} \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + s
\end{aligned}
$$

730 and

$$
\begin{aligned}
\bar{p}_{t-2}(\mathbf{d}_{t-1}) &= \sum_{k=0}^{\infty}(1-p)^k \mathbb{E}[\mathfrak{d}_{t-2+k}|\mathbf{d}_{t-1}] + s \\
&= d_{t-2} + (1-p)d_{t-1} + \sum_{k=2}^{\infty}(1-p)^k \hat{v}_{t-1}(\mathbf{d}_{t-1}) P^{k-2} \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + s \\
&= d_{t-2} + (1-p)d_{t-1} + (1-p)^2 \sum_{k=2}^{\infty}(1-p)^{k-2} \hat{v}_{t-1}(\mathbf{d}_{t-1}) P^{k-2} \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + s \\
&= d_{t-2} + (1-p)d_{t-1} + (1-p)^2 \hat{v}_{t-1}(\mathbf{d}_{t-1}) \left(\sum_{k=0}^{\infty}(1-p)^k P^k \right) \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + s \\
&= d_{t-2} + (1-p)d_{t-1} + (1-p)^2 \hat{v}_{t-1}(\mathbf{d}_{t-1}) \left(I - (1-p)P\right)^{-1} \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + s
\end{aligned}
$$

Thus, at the beginning of period t, the investor forms the expectation $\hat{p}_{t+1}$ as follows:

$$
\begin{aligned}
\hat{y}_t(\mathbf{d}_{t-1}) &= (1-\alpha)\hat{y}_{t-1}(\mathbf{d}_{t-2}) + \alpha \frac{p_{t-1}}{p_{t-2}} \frac{\bar{p}_{t-2}(\mathbf{d}_{t-1})}{\bar{p}_{t-1}(\mathbf{d}_{t-1})} \\
\Longrightarrow \hat{p}_{t+1}(\mathbf{d}_{t-1}) &= \bar{p}_{t+1}(\mathbf{d}_{t-1}) H \left(\frac{p_{t-1}}{\bar{p}_{t-1}(\mathbf{d}_{t-1})} \hat{y}_t^2 \right)
\end{aligned}
$$