

Coordination on bubbles in large-group asset pricing experiments

Te Bao^{a,b} Myrna Hennequin^b Cars Hommes^b Domenico Massaro^{b,c}

November 10, 2016

^a *Division of Economics, Nanyang Technological University, Singapore*

^b *CeNDEF, Amsterdam School of Economics, University of Amsterdam*

^c *Department of Economics and Finance, Catholic University of Milan*

Abstract

We present a large-group asset pricing experiment with the design of Hommes et al. (2008). Participants are asked to predict the price of a risky asset, whose realization depends on the aggregation of all individual forecasts. The asset markets consist of 21 to 32 participants, a group size larger than in most experiments. Multiple large price bubbles are observed in six out of seven markets. The bubbles emerge even faster and more frequently than in smaller markets. Individual forecast errors do not cancel out at the aggregate level and expectations cannot be called rational in the sense of Muth. Participants tend to coordinate on a trend-following prediction strategy that gives rise to the large bubbles. The price patterns observed in the experiment can be captured by a behavioral heuristics switching model. Heterogeneity in expectations seems crucial to explain the market dynamics.

JEL codes: D84, C91, C92, G12

Keywords: Experimental economics, Asset price bubbles, Heterogeneous expectations, Heuristics switching

1 Introduction

Expectations are a crucial part of many economic systems. Asset markets are an example of positive feedback systems: when many traders expect the price of an asset to rise, the demand for the asset will increase, leading to a higher market price. This price rise in turn affects expectations. If traders extrapolate a trend in the asset price, the positive feedback can give rise to a bubble, which occurs when the market price becomes significantly higher than the fundamental value of the asset. Such expectations-driven bubbles were observed in several asset pricing experiments (e.g. Hommes et al. (2005a) and Hommes et al. (2008)).

The main goal of this paper is to study the effect of group size on the (in)stability of experimental asset markets: will participants coordinate on bubbles in large groups? It is important to know whether the results of small-scale asset market experiments can be generalized to settings with larger groups. Are earlier observed bubbles in small markets perhaps caused by a few “irrational” subjects? Surprisingly little work focusing on group size in experiments has been done, maybe because larger experiments are costly. We conduct an asset pricing experiment with larger groups than usual and analyze both individual expectations and aggregate outcomes to empirically test if expectations-driven bubbles also occur in larger markets.

Theory does not provide a definite answer to this question and offers opposite views. For example, one may argue that the formation of bubbles is more unlikely in large groups, because individual predictions have less influence on the market price and coordination is likely to be harder. Therefore, individual forecast errors might cancel out at the aggregate level, which would be consistent with the formulation of the rational expectations hypothesis by Muth (1961). If this is correct, a large market could stabilize quickly. On the other hand, participants might see a rising price as a trend and adopt a trend-following strategy, just as in small-scale asset pricing experiments.¹ Once a large group coordinates on a bubble, the coordination

¹Trend-following behavior is also observed in real asset markets, even among more sophisti-

could be hard to break. This behavior may cause big bubbles in a larger market as well. Only a new experiment with larger group sizes can clarify which of these opposite effects will dominate in a controlled laboratory environment.

Any dynamic model of an asset market strongly depends on the underlying expectations hypothesis. It is therefore essential to develop a theory about how people form expectations and how they adapt their forecasting strategies over time. Laboratory experiments are well-suited to study individual expectation formation. Since the experimenter can control the underlying economic fundamentals, it is possible to obtain explicit observations on expectations, investigate how individual behavior shapes market outcomes and study whether aggregate outcomes deviate from fundamentals. The experimental data can be used to empirically validate different expectations hypotheses, from rational expectations to boundedly rational heuristics. Moreover, the experimental outcomes can provide insight into which forecasting strategies are more likely to be used, so that the “wilderness of bounded rationality” (Sims, 1980) can be disciplined.

Learning-to-forecast experiments (LtFEs) are especially useful to study expectations in dynamic feedback systems. This type of experiment was introduced by Marimon and Sunder (1993, 1994). In LtFEs, the participants’ only task is to submit forecasts in a particular economic setting. All other actions are computerized. Because LtFEs separate expectation formation from other choices, such as trading decisions, they provide “clean” data on expectations. Hommes (2011) provides a review of LtFEs in different economic settings.

In the asset pricing LtFE of Hommes, Sonnemans, Tuinstra and Van de Velden (2008) (henceforth HSTV08), participants have to predict the price of a risky asset for 50 periods. Each experimental asset market consists of six participants.

cated traders like hedge funds and mutual funds. Brunnermeier and Nagel (2004) establish that hedge fund managers were riding the technology bubble and escaped before the crash, thereby destabilizing the market instead of stabilizing it. Greenwood and Nagel (2009) note that young and therefore less experienced managers controlled a significant fraction of total mutual fund assets around the peak of the technology bubble. Young managers exhibited more trend-chasing behavior than old managers, leading to a larger bubble. Barberis et al. (2016) present a model of bubbles driven by extrapolation to explain the technology bubble and the recent housing bubble.




Figure 1: Market prices in all sessions of HSTV08

In every period, the market price is derived from mean-variance optimization and depends on the average price forecast. Earnings are based on prediction accuracy. Participants only have qualitative information about the asset market, but they know the risk free interest rate and the mean dividend of the risky asset. This is enough information to calculate the fundamental value of the asset. Nevertheless, large price bubbles occur in five out of six sessions, with prices rising up to fifteen times the fundamental value of 60 (see Figure 1). Once predictions reach an artificial upper bound of 1000, the trend reverses and the market crashes rapidly. There is no convergence to the fundamental.² Analysis of the individual expectations reveals that participants within the same market coordinate on a common trend-following prediction strategy. This result is remarkable, since participants do not observe the forecasts of others, so they could only coordinate through the realized market prices. An earlier asset pricing LtFE by Hommes et al. (2005a) is similar, although the design of that experiment inhibits the formation of large bubbles because of the presence of fundamental robot traders and an upper bound on price forecasts of 100.

In macroeconomic and financial market experiments, group size is an important

²Heemeijer et al. (2009) show that the positive feedback drives the non-convergence. In another treatment with negative feedback, typical for supply-driven commodity markets, prices are stable and quickly converge to the fundamental value of 60. In particular, a positive feedback system with a near unit root leads to coordination on bubbles (Sonnemans and Tuinstra, 2010; Bao and Hommes, 2015).

issue. The scale of many macroeconomic situations can often not be replicated in experiments because of the limited lab size and financial constraints. Yet, the use of macroeconomic models with microfoundations makes it possible to test the assumptions on individual behavior in small-scale lab experiments. Many examples of macroeconomic experiments are discussed by Duffy (2008, 2015) and Ricciuti (2008), who argue that lab experiments are useful to complement the theoretical and empirical research in macroeconomics. However, the question remains whether the results of such small-scale experiments are robust to increases in group size. The answer to this question is critical if the results from the experiment are compared to real macroeconomic situations with many interacting agents.

Experiments with large groups are rare.³ The vast majority uses small groups, often consisting of less than ten participants. Strategic interaction already becomes more difficult in a group of six participants than in a group of two or three participants. Still, this does not directly imply that behavior in small-scale experiments is representative of behavior in experiments with larger groups.

In this paper, we address the issue of group size by analyzing an asset pricing experiment with the same design as HSTV08, but with larger markets. The experiment consists of seven sessions with 21 to 32 participants per market. This is a group size that fits in most labs for economic experiments. Although the groups in this experiment are still small compared to the number of traders in a real asset market, the increase in market size compared to HSTV08 can shed light on the differences that group size could make in both individual expectations and market outcomes. Our experiment provides an empirical answer to the question whether bubbles will arise in large groups as well.

A large-scale classroom experiment based on the classical design of Smith et al. (1988) has shown that bubbles and crashes also occur in markets with 244, 304 and

³One of the few topics in the literature where large-group experiments are conducted is the provision of public goods. Several papers examine the effects of group sizes ranging from 4 to 100 participants (Isaac and Walker, 1988; Isaac et al., 1994; Weimann et al., 2012; Xu et al., 2013; Diederich, 2013). A large-scale financial market experiment with groups of 19 to 63 participants is conducted by Bossaerts and Plott (2004), who test the basic principles of asset pricing theory.

310 participants (Williams and Walker, 1993; Williams, 2008). In their experiment, participants buy and sell a risky asset in a double auction market. The fundamental value of the finitely lived asset is monotonically decreasing over time. Despite the fact that the fundamental value is explicitly given to all participants in each trading round, a large price bubble forms in all markets, followed by a crash. The main difference with our experiment is that Williams and Walker (1993) only consider trading decisions, while our LtFE specifically focusses on expectation formation. Futhermore, the fundamental value of the risky asset in our experiment is constant instead of declining, and not explicitly given. Lastly, the experiment of Williams and Walker (1993) was an extra-credit exercise for a microeconomics course, lasting for eight weeks. We use a standard laboratory setting with monetary incentives.

The results of our LtFE show that six out of seven markets exhibit large price bubbles. The typical price pattern shows multiple large bubbles with decreasing amplitude. Due to the high instability of the market, the forecasting performance of most participants is poor, resulting in very low earnings. Participants are able to coordinate on similar prediction strategies, but we also observe heterogeneity in expectations and strategy switching. Estimation of individual prediction strategies shows that the behavior of many participants can be captured by simple linear forecasting rules that resemble benchmark heuristics, such as trend-following rules and anchoring and adjustment. It is remarkable that the bubbles in our large groups occur even faster and more frequently than in the smaller markets of HSTV08. The reason for this is not entirely clear. It could be due to a combination of factors, such as higher initial prices and stronger trend extrapolation in large groups.

The individual prediction rules are further investigated with a behavioral heuristics switching model (HSM). This model takes account of heterogeneity in expectations and evolutionary selection among different forecasting heuristics. We use the same benchmark HSM as in Anufriev and Hommes (2012b), which is an extension of the model of Brock and Hommes (1997). The results of one-period-ahead simulations show that the bubbles are amplified by the use of a strong trend-

following prediction rule. The use of an anchoring and adjustment heuristic can explain the persistent price oscillations after the first crash. The flexibility of the HSM substantially improves the model fit compared to six homogeneous prediction rules. 50-period-ahead simulations of the HSM demonstrate that the typical price patterns can be replicated without the use of experimental data. Again, the importance of strong trend extrapolation for the formation of large bubbles is highlighted.

The remainder of this paper is organized as follows. Section 2 explains the experimental design in detail. The results of the experiment are discussed in Section 3. In Section 4, we estimate forecasting rules for each participant. Section 5 compares the formation of the first bubble in our large groups and the small groups of HSTV08. The specification of the HSM is described in Section 6 and simulations with the model are presented in Section 7. Section 8 summarizes and concludes.

2 Asset pricing experiment

2.1 Asset pricing model with heterogeneous expectations

The experiment is based on a standard asset pricing model with heterogeneous beliefs, as in Campbell et al. (1997) and Brock and Hommes (1998). The asset market consists of I traders. Each trader can invest in a risk free asset or a risky asset. The risk free asset (e.g. a savings account) pays a fixed interest rate r . Each share of the infinitely lived risky asset pays an uncertain dividend y_t in every period t . The dividends are independently and identically distributed with mean \bar{y} . The fundamental value of the risky asset is thus given by $p^f = \bar{y}/r$. The price of the risky asset in period t is denoted by p_t and the number of shares purchased by trader i in period t is denoted by z_{it} . Note that a negative value of z_{it} means that the trader sells a number of shares of the risky asset. The realized wealth of

the trader in the next period is given by

$$W_{i,t+1} = RW_{i,t} + (p_{t+1} + y_{t+1} - Rp_t)z_{it}, \quad (1)$$

where $R = 1 + r$ is the gross rate of return of the risk free asset.

Traders differ in their beliefs about the conditional mean E_{it} and the conditional variance V_{it} of the evolution of wealth. Traders are myopic mean-variance optimizers, so the demand for shares z_{it} corresponds to the solution of

$$\begin{aligned} \max_{z_{it}} \left\{ E_{it}(W_{t+1}) - \frac{1}{2}aV_{it}(W_{t+1}) \right\} = \\ \max_{z_{it}} \left\{ z_{it}E_{it}(p_{t+1} + y_{t+1} - Rp_t) - \frac{1}{2}az_{it}^2V_{it}(p_{t+1} + y_{t+1} - Rp_t) \right\}, \end{aligned} \quad (2)$$

where a measures the degree of risk aversion. It is assumed that traders believe that the conditional variance of excess returns is constant: $V_{it}(p_{t+1} + y_{t+1} - Rp_t) = \sigma^2$ for all i . This assumption can be made because the experiment only asks for point predictions, so traders beliefs about the distribution of returns are not considered.⁴ The solution of the maximization problem is then

$$z_{it} = \frac{E_{it}(p_{t+1} + y_{t+1} - Rp_t)}{a\sigma^2}. \quad (3)$$

An increase in the expected price of the risky asset in period $t + 1$ thus leads to an increase in demand in period t .

Assume that the outside supply of shares z^s is zero, which means that some traders buy the shares that other traders sell, so that z_{it} can be interpreted as excess demand. This is equivalent to assuming risk neutral traders ($a = 0$). The

⁴Bottazzi et al. (2011) conduct an experiment with a treatment where participants forecast future returns of an asset and provide confidence intervals, which are taken as a measure of perceived volatility of the returns. This somewhat reduces high volatility in price fluctuations and increases coordination on a common prediction strategy compared to the baseline treatment where no confidence intervals were elicited.

market equilibrium condition is then given by

$$\sum_{i=1}^I z_{it} = \frac{1}{a\sigma^2} \sum_{i=1}^I E_{it}(p_{t+1} + y_{t+1} - Rp_t) = z^s = 0. \quad (4)$$

Using that the mean dividend \bar{y} is common knowledge so that $E_{it}(y_{t+1}) = \bar{y}$ for all i and all t , we can solve the equation for the market equilibrium price:

$$p_t = \frac{1}{1+r} \left[\frac{1}{I} \sum_{i=1}^I p_{i,t+1}^e + \bar{y} \right] = \frac{1}{1+r} [\bar{p}_{t+1}^e + rp^f], \quad (5)$$

where $E_{it}(p_{t+1}) = p_{i,t+1}^e$ denotes the prediction by trader i in period t for the price in period $t+1$, and $\bar{p}_{t+1}^e = \frac{1}{I} \sum_{i=1}^I p_{i,t+1}^e$ is the average price prediction. Note that the price in each period depends on the predictions for the next period. When traders are in period t and they have to forecast the price for period $t+1$, they do not know p_t yet and thus can only use the past prices up to period $t-1$. Hence, they make a two-period-ahead price prediction.

Equation (5) shows that there is positive feedback: if the average price prediction is high, then the realized market price will also be high, and vice versa. This is an important characteristic of asset markets. The level of the interest rate r determines the strength of feedback, given by $\lambda = \frac{1}{1+r}$. The fundamental price $p^f = \bar{y}/r$ gets a weight of $1 - \lambda = \frac{r}{1+r}$ in the pricing mechanism. So the higher the interest rate, the weaker the expectations feedback and the stronger the market price is pushed towards the fundamental. On the other hand, for low interest rates (e.g. $r = 0.05$) each price prediction is almost self-fulfilling ($p_{t+1} \approx \bar{p}_{t+1}^e$), but the fundamental price is the only steady state equilibrium that is perfectly self-fulfilling ($p_{t+1} = \bar{p}_{t+1}^e$).

2.2 Benchmark expectation rules

To get a feeling for the dynamics of the asset pricing model, we simulate the evolution of the market price under homogeneous expectations, i.e. when $p_{i,t+1}^e =$

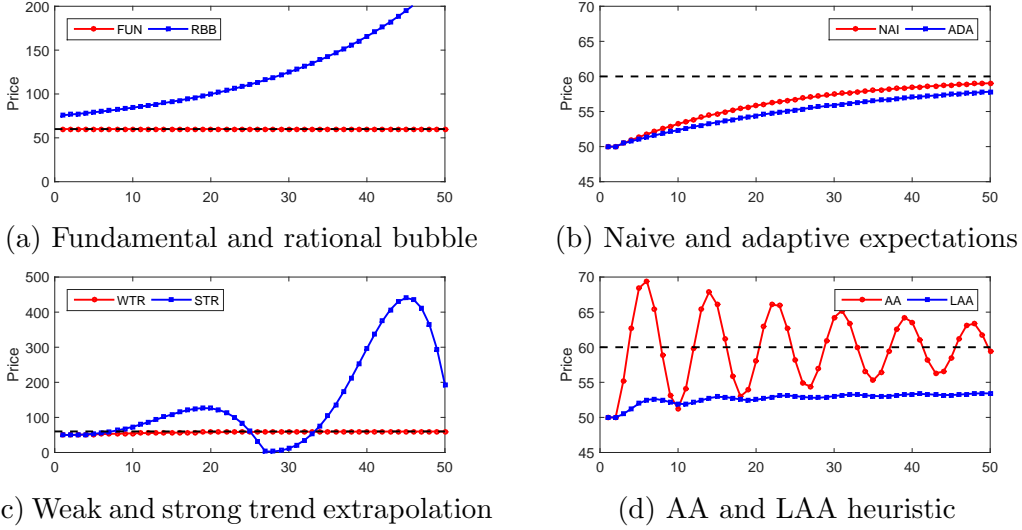


Figure 2: Simulations with homogeneous expectations

Initialization of the simulations: $p_1 = p_2 = 50$ and $p_3^e = 50$. The dashed line indicates the fundamental value of 60. Note that the scale of the vertical axis differs in the four figures.

p_{t+1}^e for all i . We consider some well-known heuristics: rational expectations, adaptive and naive expectations, trend-following rules and anchoring and adjustment rules. Figure 2 shows the price dynamics for these benchmark expectation rules.

In a rational expectations equilibrium (REE) of the model, the realized price in period t coincides with the prediction for that period. The unique REE steady state is where all traders predict the fundamental price of the asset:

$$p_{t+1}^e = p^f = \bar{y}/r. \quad (6)$$

This is called fundamental expectations. This solution satisfies a no-bubbles condition, which ensures that predictions and prices do not explode. However, as explained in Hommes et al. (2008), there is a continuum of other solutions of the form $p_{t+1}^e = p^f + R^{t+1}c$ with $c > 0$ that do not satisfy this condition. These so-called rational bubbles grow at rate R . Figure 2a shows the fundamental price and a rational bubble for $c = 15$.

The rational expectations hypothesis is quite demanding. It requires all traders to be able to calculate the fundamental price and coordinate on the REE steady

state. If traders instead use simple forecasting heuristics, it could be that predictions and prices converge to the fundamental value because of individual learning. This is the case when all traders form expectations adaptively:

$$p_{t+1}^e = wp_{t-1} + (1-w)p_t^e = p_t^e + w(p_{t-1} - p_t^e), \quad (7)$$

with weight $w \in [0, 1]$. The forecast is formed by adapting the previous prediction in the direction of the last observed price. Naive expectations are a special case of this rule, obtained for $w = 1$. The prediction is then simply the last observed price. The asset pricing model is stable under adaptive or naive expectations. Prices converge slowly but monotonically to the fundamental value. This is illustrated in Figure 2b, which shows prices for naive expectations and adaptive expectations with $w = 0.65$.

A trend-following rule bases the prediction on the last observed price and adjusts in the direction of the last observed price change:

$$p_{t+1}^e = p_{t-1} + \gamma(p_{t-1} - p_{t-2}) \quad (8)$$

with extrapolation coefficient $\gamma > 0$. The trend-following rule is called weak if $\gamma < 1$ and strong if $\gamma > 1$. The price dynamics depend on the extrapolation coefficient. For weak trend extrapolation with a small value of γ , there is monotonic convergence, as illustrated in Figure 2c for $\gamma = 0.4$. For γ close to but smaller than 1, convergence becomes oscillatory. For strong trend extrapolation, the oscillations increase in amplitude and prices diverge. Figure 2c shows an example of this with $\gamma = 1.3$.

The anchoring and adjustment (AA) rule is based on Tversky and Kahneman (1974) and has a clear behavioral interpretation. The average of the fundamental price and the last observed price serves as an anchor for the next prediction. The

prediction is adjusted in the direction of the last observed price change:

$$p_{t+1}^e = 0.5(p^f + p_{t-1}) + (p_{t-1} - p_{t-2}). \quad (9)$$

A variant of this heuristic is the learning anchoring and adjustment (LAA) rule, where traders do not have to know the fundamental value, but they might learn it through the sample average of past prices. For this rule, p^f is replaced by $p_{t-1}^{av} = \frac{1}{t-1} \sum_{j=1}^{t-1} p_j$. Under the AA rule, prices converge to the fundamental value, but convergence is slow and oscillatory. There is also convergence under the LAA rule, but even slower and with less pronounced oscillations. This can be seen from Figure 2d.

The evolution of prices can thus vary greatly with the prediction rule that is used. When expectations are heterogeneous instead of homogeneous, almost any type of price dynamics can occur. Brock and Hommes (1998) consider an evolutionary competition between simple heuristics, such as fundamentalists versus trend followers. They find price oscillations, bubble and crash dynamics and chaotic behavior of the asset pricing model with evolutionary switching.

2.3 Experimental design

The experiment is aimed at investigating expectation formation in a positive feedback environment. Therefore, the participants only have one task: predicting the price of the risky asset. Trading is computerized in the sense that the computer calculates the optimal demand for shares of each participant based on mean-variance optimization, as in the asset pricing model. Given the forecasts of all participants in the market, the realized market price is determined by Equation (5).

In the instructions (see Appendix A), the participants are told that they are a financial advisor to a large pension fund. It is explained that the pension fund can invest in a risk free asset and a risky asset. Participants have to forecast the price of the risky asset so that the pension fund (i.e. the computer) can calculate

the optimal demand for shares. It is a two-period-ahead prediction: in period t , the available information for forecasting the price in period $t + 1$ consists of past prices up to period $t - 1$ and past predictions up to period t . After all participants have entered their predictions for period $t + 1$, the realized market price in period t is revealed. This process continues for 51 periods. Participants do not have any information about past prices in the first two periods, but they are told that it is very likely that the price will be between 0 and 100 in those periods.

Participants receive only qualitative information about the asset market. The interest rate $r = 5\%$ and the mean dividend $\bar{y} = 3$ are given, so they have enough information to calculate the fundamental value of the risky asset: $p^f = \bar{y}/r = 60$. Participants are told that the higher their price forecast, the larger the demand for shares by the pension fund. Hence they could infer that there is positive feedback. They are also informed that supply is fixed and that the price is determined by equilibrium between demand and supply. However, they do not know the underlying market equilibrium equation (5). They also do not know the exact number of pension funds in the market or the identity of the other participants in their group.

Because of the participants' role of advisor to a pension fund, payoffs are based on their prediction accuracy: the better their price forecast, the higher their earnings.⁵ The number of points e_{it} earned by participant i in period t is determined by the quadratic scoring rule

$$e_{it} = \max \left\{ 1300 - \frac{1300}{49}(p_t - p_{it}^e)^2, 0 \right\}. \quad (10)$$

Note that participants could not lose money, because earnings are simply zero when the forecast error $|p_t - p_{it}^e|$ is larger than 7. The instructions include a table where participants could find the number of points they would earn for each value of the

⁵As explained in Hommes (2001), using the quadratic forecast error in the payoff function is equivalent to using realized risk-adjusted utility from wealth or profits. Moreover, Bao et al. (2014) run an experiment with a learning-to-optimize treatment, where participants trade and are paid by realized profits. They show that bubbles form even more easily than in the learning-to-forecast treatment. The findings of Hanaki et al. (2016) support the choice of incentivizing participants in the learning-to-optimize treatment of Bao et al. (2014).

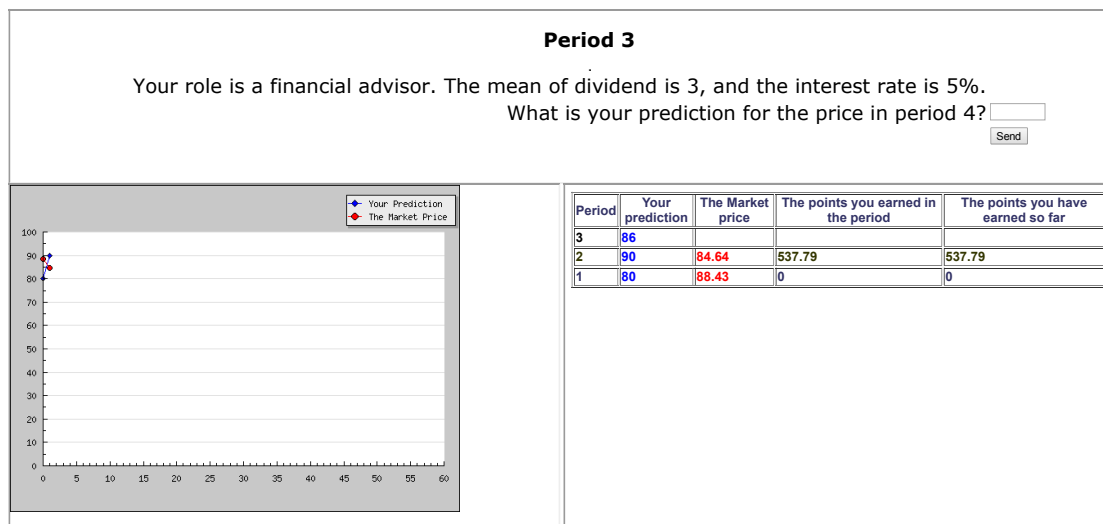


Figure 3: Example of the screen that participants see during the experiment

forecast error. At the end of the experiment, the participants receive €0.5 per 1300 points that they have earned.

To keep the experiment comparable with HSTV08, an upper bound on predictions of 1000 is imposed. The instructions do not mention this upper bound. If a participant tries to enter a forecast higher than 1000, the computer screen shows a message that predictions above 1000 are not accepted and that the participant has to submit a lower prediction.

At the beginning of the experiment, the participants receive instructions on paper. They could consult these instructions at any time. After reading the instructions, participants answer five control questions to check if they understand the experiment. The computer screen during the 51 periods shows a table and a graph of all past prices and own predictions that are available in that period, earnings in the previous period, total earnings, the interest rate and the mean dividend. An example of the screen is shown in Figure 3. After the last period, participants fill in a short questionnaire, with a single open question about the strategy they used in the experiment.

The main motivation for our experiment is to test whether the results of the LtFE of HSTV08 are robust to an increase in the number of participants in the

market. However, the group size was limited to the number of computers in the laboratory. The aim was to run sessions with groups of around 30 participants, a group size that typically fits into an experimental lab. To avoid having to cancel a session because not enough participants showed up, the group size was flexible. A session would start if 20 or more participants arrived at the lab on time. Seven sessions were run, with respectively 26, 26, 24, 32, 21, 21 and 32 participants in session 1–7. No sessions were cancelled.

The experiment was conducted with students. It took place in the CREED laboratory of the University of Amsterdam in April 2014. Each session lasted for about 1.5–2 hours. As discussed in Section 3, the number of points that participants earned during the experiment was very low in all sessions. To reward them for their time, it was announced after the session that a lump sum payment of €15 was added to their earnings. Including this additional payment, average earnings were €17.27.

3 Experimental results

3.1 Aggregate outcomes

Figure 4 shows the realized market prices in the seven sessions of the experiment. For the ease of comparison, our large-group sessions are named L1–L7 and the small-group sessions of HSTV08 will be called S1–S6. Summary statistics can be found in Table 7 in Appendix B. It is immediately clear that the large price bubbles that were observed in HSTV08 also occur in our large groups. In six out of seven sessions, the markets are highly unstable and the price approaches the upper bound of 1000.⁶ In these sessions, the fundamental value of 60 does not seem to play a role in the evolution of the market price at all.

In all sessions, the market price starts rising right from the beginning. The reason for this might be twofold. First, many participants make an initial prediction

⁶The maximum price that can be attained is actually $(1000 + 3)/(1 + 0.05) = 955.24$. The realized prices in session L1, L2, L3, L4, L6 and L7 come close to this maximum price.

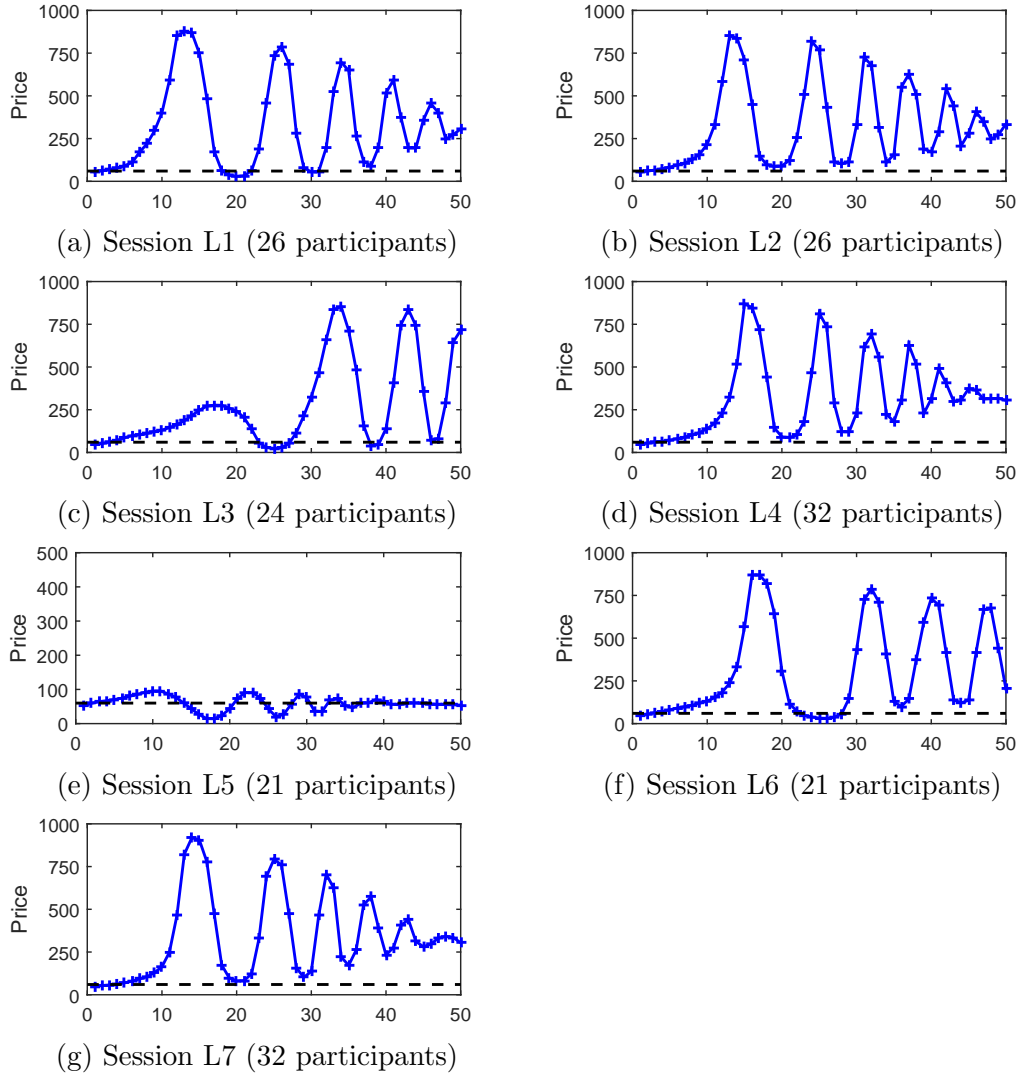


Figure 4: Realized market prices

The dashed line indicates the fundamental value of 60. Note that the vertical axis for session L5 runs up to 500 instead of 1000.

of around 50. This is probably because it is mentioned in the instructions that the price in the first two periods is likely to be between 0 and 100, and participants have no information about past prices in the first two periods. Since the fundamental price is 60, the initial market price turns out to be a little higher than the average prediction in the first two periods. Therefore, many participants increase their predictions after they learn the first market price, resulting in an upward trend. This might evoke trend extrapolating behavior, leading to a bubble. Second, it is notable that 132 out of 182 participants make a prediction for period 2 that is higher

than their prediction for period 1.⁷ They do this without having any information about past prices. This suggests that a majority of participants expects an upward trend from the beginning.

Session L1, L2, L4, L6 and L7 display the same general pattern. A large price bubble forms immediately. As the price approaches 1000 after about fifteen periods, most participants discover the upper bound and lower their predictions.⁸ The trend reverses and the market rapidly crashes back to a price below 100. After that, a new bubble quickly forms again and the process repeats itself, although the amplitude of the bubbles is decreasing. The price at which the bubble peaks becomes lower each time, and the price at which the bubble starts to grow again becomes higher. This could be because participants learn to anticipate a new crash (bubble), and therefore decrease (increase) their forecasts earlier each time. The price seems to stabilize around 300 towards the end of these sessions, except in session L6. Note that a price of 300 is five times the fundamental value, so prices are still far above the unique REE steady state.

It is remarkable that the bubbles in these sessions arise even faster than in the small-group experiment of HSTV08. In the large groups, the bubble already peaks between period 13 and 16. In the small groups, it takes more time before the trend really takes off: there is only a slow upward trend in the first fifteen to twenty periods, after which the price starts increasing rapidly. The bubble takes between 22 and 38 periods to reach its peak in the small groups (see Table 8 in Appendix B). In the typical large-group sessions, a total of four to six repeated bubbles and crashes occur. Only two or three bubbles per session are observed in the small groups of HSTV08. This difference also reflects that bubbles arise faster in our large-group experiment.

Session L3 and L5 are the exceptions to the general price pattern of multiple

⁷Furthermore, 30 participants submit equal predictions for the first two periods. Only 20 participants predict a lower price for period 2 than for period 1.

⁸In session L1, 21 out of 26 participants entered a prediction of 999 or 1000 in one or more periods around the time of the first peak. For session L2, L4, L6 and L7, these numbers were 16/26, 20/32, 14/21, and 23/32, respectively. This suggests that more than 60% of the participants found out about the upper bound.

large bubbles. Session L3 starts with a small bubble that forms slowly. The upward trend reverses when two participants predict 0 or 1 for two or three periods. After the first small bubble, the pattern in session L3 is similar to the sessions discussed above. The price starts rising quickly and a large bubble occurs, followed by two more bubbles that are only slightly smaller.

Large bubbles did not occur in session L5. In this session, the price stays between 0 and 100. This price range is mentioned in the instructions as the most likely price range in the first two periods. It might therefore serve as an anchor, causing some participants to lower their predictions as they approach 100.⁹ Some participants submit a forecast higher than 100 in some periods, but these forecasts do not become higher than 131.¹⁰ The price oscillates, but the amplitude of the oscillations decreases and the price stabilizes around 55, which is fairly close to the fundamental value of 60.

The bubbles in each session can be quantified and compared by the Relative Absolute Deviation (RAD) and Relative Deviation (RD), as proposed by Stöckl et al. (2010). For our experiment, these measures are defined as follows:

$$\text{RAD} = \frac{1}{50} \sum_{t=1}^{50} \frac{|p_t - p^f|}{p^f} \quad (11)$$

$$\text{RD} = \frac{1}{50} \sum_{t=1}^{50} \frac{p_t - p^f}{p^f} \quad (12)$$

The RAD measures the average level of mispricing in the market: a value of e.g. 1.5 means that on average the market price differs 150% from the fundamental. The RD measures the average level of overvaluation. Equal values of RAD and RD indicate that the asset is always overvalued, while differences imply that the

⁹Ten participants (48%) do not predict a price higher than 100 and lower their predictions in the two periods before they observe the price peak. They might think that the price can or will not become higher than 100 in all periods, not just the first two. In the other sessions, this behavior occurs less often, so that the upward trend just continues. There are at most four participants in each of the other sessions that lower their predictions in the periods before the price actually becomes higher than 100, after which they adjust their expectations upwards.

¹⁰The maximum prediction is actually 440 in period 40, but that is most likely a typing error. The last observed market price at that time was 60.81, and the participant predicted 50 in both period 39 and period 41.

Table 1: Relative Absolute Deviation (RAD) and Relative Deviation (RD)

Large groups	RAD	RD	Small groups	RAD	RD
Session L1	4.52	4.46	Session S1	0.98	0.93
Session L2	4.40	4.40	Session S2	3.67	3.56
Session L3	3.66	3.55	Session S3	2.17	1.96
Session L4	4.36	4.35	Session S4	4.27	4.10
Session L5	0.26	0.00	Session S5	2.15	1.78
Session L6	4.25	4.16	Session S6	2.40	2.05
Session L7	4.67	4.66			
Average	3.73	3.65	Average	2.60	2.40

market also has periods where the asset is undervalued.

Table 1 shows the bubble measures for both our large groups and the small groups of HSTV08. As expected, the bubble measures indicate that the deviations from the fundamental price are large in the markets where large bubbles occurred. In the large groups, RAD is between 3.66 and 4.67. In the small groups, this range is wider: RAD is between 2.15 and 4.27. Sessions S2 and S4 exhibited two bubbles where the price almost reached the upper bound. This results in a larger RAD than for sessions S3, S5 and S6, in which at least one of the bubbles stayed smaller. RAD is considerably lower in the markets without large bubbles, namely 0.26 in session L5 and 0.98 in session S1. In all sessions, RD is only slightly smaller than RAD, meaning that the asset is mostly overvalued. RD is virtually equal to zero in session L5, indicating that the price oscillated around the fundamental value. On average, both measures are higher in the large groups (RAD = 3.73 vs. RAD = 2.60 and RD = 3.65 vs. RD = 2.40), confirming that the large markets are even more unstable than the small markets.

3.2 Individual expectations

Figure 5 shows the individual predictions of all participants in the large groups in color, together with the realized market prices in black. After a few periods, it seems that the participants coordinate on a common trend-following prediction strategy, leading to the first bubble. The differences in predictions within markets appear to be somewhat larger than in the experiment of HSTV08. This is probably simply

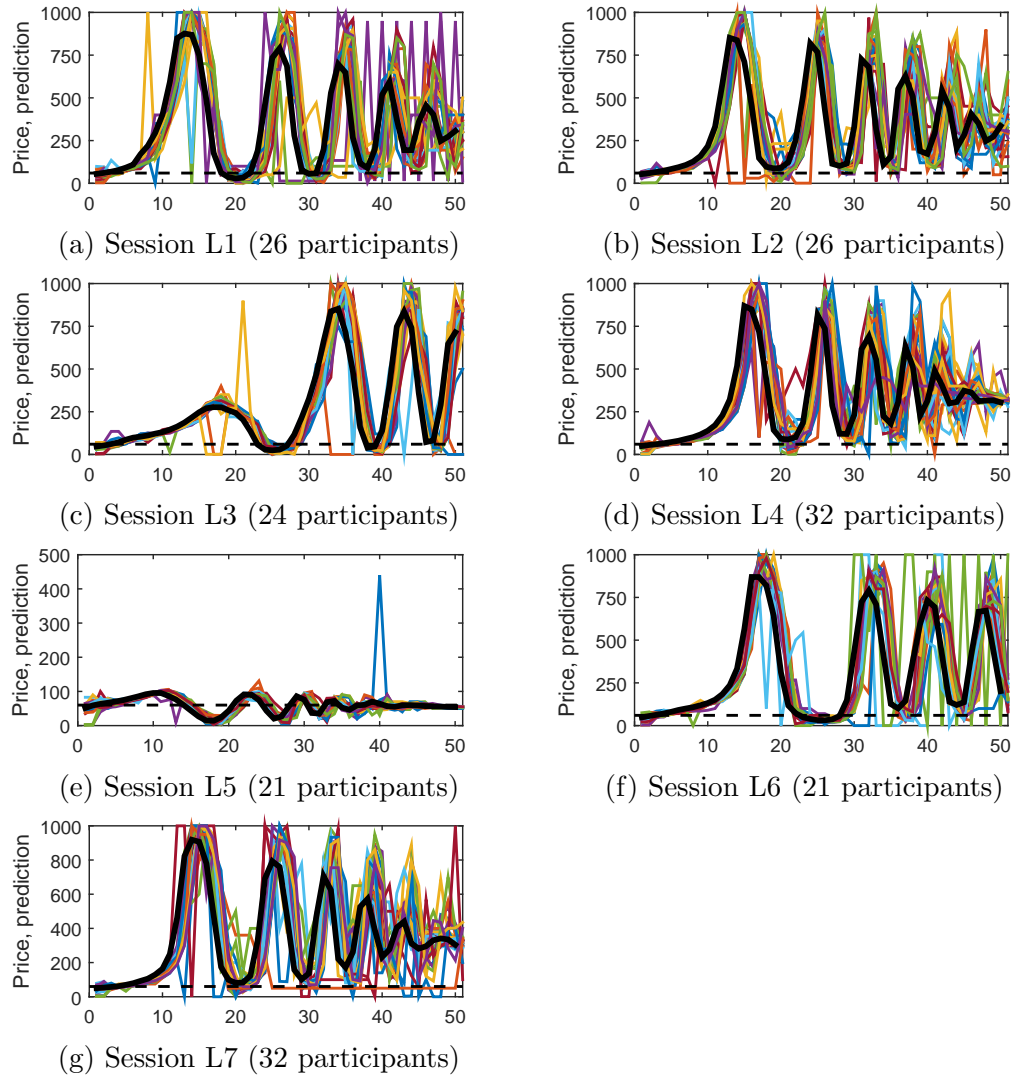


Figure 5: Individual predictions

The colored lines indicate the individual predictions, the thick black line indicates the market price. The dashed line indicates the fundamental value of 60. Note that the vertical axis for session L5 runs up to 500 instead of 1000.

because it is harder to achieve strong coordination in a larger group. Although most participants submit similar forecasts, there are also some participants that predict very different prices. The discovery of the upper bound of 1000 seems to break the coordination to a large extent. In the typical sessions L1, L2, L4, L6 and L7, the differences in predictions become larger after the first two bubbles, and forecasts are quite heterogeneous towards the end of the experiment.

A way to investigate the importance of coordination quantitatively is to compute the average individual quadratic forecast error. This is the individual quadratic

Table 2: Average individual quadratic forecast error

	Av. individual error	Av. dispersion error	(%)	Av. common error	(%)
Session L1	54,804	24,210	44%	30,594	56%
Session L2	53,749	16,917	31%	36,832	69%
Session L3	33,225	11,432	34%	21,793	66%
Session L4	46,764	14,326	31%	32,438	69%
Session L5	478	242	51%	236	49%
Session L6	50,127	21,095	42%	29,033	58%
Session L7	52,813	22,459	43%	30,354	57%

forecast error averaged over the last 45 periods and over the I participants in a session. The first five periods of the experiment are omitted in the computation to allow for a short learning phase. The average individual quadratic forecast error can be broken up in two terms, the average dispersion error and the average common error:

$$\frac{1}{45I} \sum_{i=1}^I \sum_{t=6}^{50} (p_{it}^e - p_t)^2 = \frac{1}{45I} \sum_{i=1}^I \sum_{t=6}^{50} (p_{it}^e - \bar{p}_t^e)^2 + \frac{1}{45} \sum_{t=6}^{50} (\bar{p}_t^e - p_t)^2, \quad (13)$$

where $\bar{p}_t^e = \frac{1}{I} \sum_{i=1}^I p_{it}^e$ is the average price prediction for period t . The average dispersion error measures the squared distance between the individual prediction p_{it}^e and the average prediction \bar{p}_t^e , averaged over time and participants. This term is zero if all participants use exactly the same prediction strategy. Hence, a relatively small average dispersion error indicates that there is coordination on a common strategy. The average common error measures the squared distance between the average prediction \bar{p}_t^e and the realized price p_t , averaged over time. Muth's (1961) formulation of the rational expectations hypothesis implies that individual expectations may be wrong, but expectations should be approximately correct in the aggregate. If this is the case, then the average common error is relatively small.

Table 2 presents the value of the average individual quadratic forecast error in each session, together with the values and percentages of the average dispersion error and the average common error. In all sessions with large bubbles, the average dispersion error is between 31% and 44%. This indicates that predictions are similar, but coordination is not perfect. The relatively large average common error

Table 3: Earnings (excluding the lump-sum payment)

	Average	Minimum	Maximum
Session L1	0.77	0.0	1.8
Session L2	1.30	0.2	3.4
Session L3	2.08	0.8	3.6
Session L4	1.57	0.1	3.2
Session L5	7.75	4.2	12.8
Session L6	2.14	0.4	3.6
Session L7	1.61	0.0	3.2

(between 56% and 69%) implies that individual prediction errors do not cancel out in the aggregate. Even in the stable session L5, the average common error is still 49%. It suggests that expectations in this experiment cannot be called rational. Compared to HSTV08, the size and proportions of the average dispersion error and the average common error are roughly the same. Apparently, the difference in coordination between the two experiments is not that large.

Clearly, the average individual quadratic forecast errors are very large, especially in the sessions with large bubbles. This is also reflected in the low earnings of the participants during the experiment (see Table 3). In the sessions with large bubbles, earnings are between €0.00 and €3.60, while the maximum possible earnings are €25.00. During a bubble, participants barely earn anything. Most points are earned in the first ten periods or just after a crash, when the price starts going up again.¹¹ In the relatively stable session L5, earnings were substantially higher, but still far from maximal. In this session, participants also earn points in the beginning, but most points are earned in the last ten periods, when the price stabilizes.

Figure 5 already indicates that there is heterogeneity in expectations. Of course, participants do not have to use the same prediction strategy during the whole experiment. As they learn more about the market, they might switch to a different strategy. A graphical inspection of individual predictions shows that this indeed

¹¹The low earnings led to frustration among the participants. Some started to predict almost random prices, hoping to be lucky and earn some points. But this behavior did not affect the formation of the first large bubble, since participants were still actively trying to make accurate predictions in the beginning. This is also supported by the comments in the questionnaire.

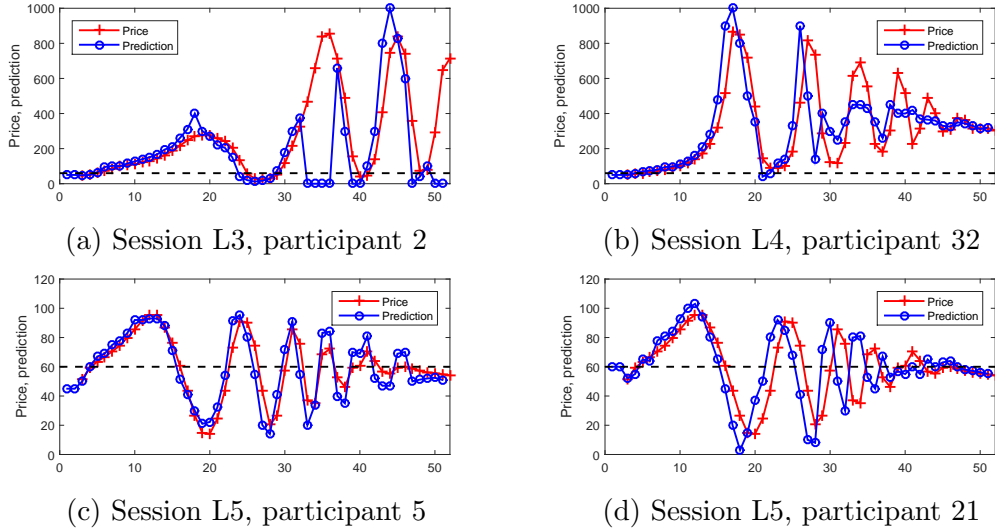


Figure 6: Examples of strategy switching

The red line indicates the market price lagged for two periods, the blue line indicates the individual prediction. The dashed line indicates the fundamental value of 60.

occurs. Figure 6 gives four examples of participants who switch between forecasting rules. The figure shows for each period t the lagged market price p_{t-2} and the individual prediction $p_{i,t}^e$. This thus illustrates how the last observed market price is used to form the two-period-ahead price forecast. For example, if the two time series coincide, the participant is using a naive prediction strategy.

In session L3, participant 2 starts with a strategy close to naive expectations (see Figure 6a). After some periods, she begins to extrapolate the trend in the market price. In period 31, she realizes that this strategy is not making any money during a bubble, so she attempts to bring the price down by predicting zero. In some periods, she tries to make an accurate prediction by extrapolating the trend.

Participant 32 in session L4 (Figure 6b) uses a trend-following strategy in the first half of the experiment. He then switches to a strategy that looks like an anchoring and adjustment rule, with an anchor of around 300 and a relatively small adjustment in the direction of the last price change.

Figure 6c shows that in session L5, a weak form of trend extrapolation is used by participant 5 until period 30. After that, she switches to a much stronger trend-following rule, despite the large forecast errors of this strategy. In the last five

periods, she uses an adaptive forecasting rule, predicting a price between the last realized price and her own last prediction.

In the same session, participant 21 realizes that the fundamental price of the asset is 60 and predicts this price in the first two periods (Figure 6d). However, the market price turns out to be lower than 60, so she lowers her prediction. She uses an almost naive strategy for some periods and then begins to extrapolate the price trends. After two bubbles, she tries to anticipate the trend reversals, in which she is quite successful. When the price stabilizes in the last five periods, she goes back to a strategy close to naive.

4 Estimating individual prediction strategies

4.1 General linear forecasting rules

It seems that the prediction strategies of participants could be described by fairly simple heuristics. Estimating some simple linear specifications of forecasting rules can give insight in the way participants form expectations.¹² Similar analyses for different experiments are done in Hommes et al. (2005a), Heemeijer et al. (2009) and Assenza et al. (2014). Starting with a general specification, a simple linear rule of the form

$$p_{i,t+1}^e = \alpha + \sum_{k=1}^4 \beta_k p_{t-k} + \sum_{l=0}^3 \gamma_l p_{i,t-l}^e + u_t \quad (14)$$

is estimated for each participant i .¹³ This rule uses the last four observations of the market price and the last four own predictions to form the new prediction. The forecasting rule is estimated from period $t = 5$, which is the first period that all regressors are available. The first five periods can also be considered as a learning

¹²An alternative way to find out which forecasting rules are used is to do a strategy experiment as in Hommes et al. (2005b). As participants formulate a complete prediction strategy, there is no need to estimate their strategies from the observed predictions.

¹³The forecasting rule is estimated after removing outliers, i.e. predictions that differ substantially from what would be expected from the general pattern of the other observations. A total of 24 outliers for 19 participants were removed by linear interpolation (0.3% of all observations).

phase. Insignificant regressors are removed one by one, highest p -value first, until all remaining regressors are significant at a 5% level. The final prediction rule is tested for autocorrelation (Breusch-Godfrey test, two lags), heteroskedasticity (White test, no cross terms) and misspecification (Ramsey RESET test, one fitted term).

Table 10 in Appendix B presents the estimation results. It seems that the general linear specification can describe the forecasting rules fairly well for most participants. 139 out of the 182 estimated rules (76%) pass all three diagnostic tests. In all treatments, the most popular regressors are the last two observations of the market price: p_{t-1} is used by 89% of the participants, p_{t-2} by 71%. For 20% of the participants, these are the only two significant regressors. In all these cases, the coefficient of p_{t-2} is negative, so that it is a trend-following rule. The last own prediction p_t^e is significant in 36% of the estimated rules. 11% of the participants use a form of anchoring and adjustment, basing their predictions on p_{t-1} , p_{t-2} and p_t^e .

Although these are the most commonly used regressors, a majority of participants (69%) also use higher lags of the market price and their own predictions in their forecasting rules. The use of more past observations indicates learning behavior by participants. Particularly in session L5, p_{t-3} and p_{t-4} are significant in more than half of the estimated rules. These third and fourth lags of the market price are also used more often in session L3 than in the other sessions. Furthermore, the estimation results show 22 different combinations of significant regressors across all sessions. Within each session, ten to fourteen different forecasting rules are estimated.

We also estimated these forecasting rules for the experiment of HSTV08. Relatively less rules pass all three diagnostic tests, namely 15 out of 36 estimated rules (42%). Again, the last two observations of the market price are the most widely used: p_{t-1} is used by 97% of the participants, p_{t-2} by 75%. 11% of the participants only used these two regressors and the coefficient of p_{t-2} is always negative,

so that it is a trend-following rule. In 33% of the estimated rules, the last own prediction p_t^e is significant. 8% of the participants base their predictions on p_{t-1} , p_{t-2} and p_t^e and thus use a form of anchoring and adjustment. Just as in the large groups, a majority of participants (75%) also base their forecasts on higher lags of the market price and their own predictions. Sixteen different combinations of significant regressors across all sessions are found in total, with four to five different rules within each session. All in all, the results in the small groups of HSTV08 are remarkably similar to the results in our large-group experiment.

Hommes et al. (2005a) estimated the same specification for the individual predictions in their LtFE with a robot trader and found different results. In their experiment, p_{t-1} and p_{t-2} are the only two significant regressors for a majority of 63%. For 30% of the participants, the last own prediction p_t^e is significant as well, but only 2% use a form of anchoring and adjustment. This is because almost all participants that base their forecast on p_{t-1} , p_{t-2} and p_t^e also use higher lags of prices and predictions in their forecasting rule. In total, 23% of the participants makes use of higher lags. In their experiment, fourteen different combinations of regressors are observed in total, with two to five different rules per session. Comparing the results of Hommes et al. (2005a) to ours, we see that the prediction rules that we observe in our experiment are more diverse and use more information on past prices and predictions. This could be due to the larger instability of the markets in our experiment, which calls for more complicated forecasting rules. Of course, we should keep in mind the differences in group size: Hommes et al. (2005a) only have six participants per session and sixty participants in total. Therefore, we cannot directly conclude that the larger variety in forecasting heuristics implies that there is more heterogeneity in expectations in our LtFE with larger groups.

4.2 First-order heuristics

Since the last two observed market prices and the last own prediction are the most popular regressors, it is interesting to consider a simplified forecasting rule based

on these variables, nested in the general specification (14):

$$p_{i,t+1}^e = \alpha_1 p_{t-1} + \alpha_2 p_{i,t}^e + (1 - \alpha_1 - \alpha_2) \bar{p} + \beta(p_{t-1} - p_{t-2}) + v_t, \quad (15)$$

where $\bar{p} = \frac{1}{50} \sum_{t=1}^{50} p_t$ is the average market price in the session. This rule is called a first-order heuristic. It is a more general form of the anchoring and adjustment rule. The anchor is now a weighted average of the last observed market price, the last own prediction and the sample average price.¹⁴ The adjustment is based on the difference between the last two observed prices.

The first-order heuristic has a couple of benchmark expectation rules as special cases. For $\alpha_1 = 1, \alpha_2 = \beta = 0$, the rule reduces to naive expectations. If $\alpha_1 + \alpha_2 = 1, \beta = 0$, expectations are adaptive. In case $\alpha_1 = \alpha_2 = \beta = 0$, the forecast is simply the sample average price. This case is referred to as fundamental expectations, although the average price is not near the fundamental price of 60 in this experiment, except in session L5. The sample average price just serves as a proxy for the long-run price level, even though it is not close to the steady state level in all but one session. Finally, when $\beta > 0$, the adjustment is called trend-following.

The first-order heuristic is estimated for those participants whose general forecasting rule of type (14) can be restricted to a heuristic of type (15). This is verified by an F -test on joint parameter restrictions.¹⁵ The first-order heuristics are again estimated from period $t = 5$, to allow for some learning. The regressor with the highest p -value is removed until there are no more insignificant regressors. F -tests

¹⁴The sample average price of the entire session is of course not available to the participants at the moment they form their predictions. It is included in the estimation as a proxy for the equilibrium price level that they are trying to learn. In the HSM of Section 6, the LAA heuristic uses the observable sample average of past prices available at each point in time. This generally converges quickly to the sample average price.

¹⁵Note that separate t -tests were used to remove insignificant regressors from the general forecasting rules of type (14), while the F -test used here is a joint significance test. The two ways of testing are not equivalent. It is possible that the individual t -tests indicate that all regressors are significant, while the F -test cannot reject the null hypothesis that a number of these coefficients are jointly zero. Hence, the number of rules that can be restricted to a heuristic of type (15) is not the same as the number of rules of type (14) that only have p_{t-1} , p_{t-2} and p_i^e as significant regressors.

are used to see if the first-order heuristic can be restricted to one of the benchmark expectation rules.

The estimation results can be found in Table 11 in Appendix B. For 58 out of 182 participants (32%), the estimated general forecasting rule can be successfully restricted to a first-order heuristic. In session L5, there is only one general rule that can be restricted, since many participants in that session used higher lags of the market price and their own predictions. For the same reason, only three first-order heuristics are estimated in session L3.

The last two columns of the tables indicate if the anchor and the adjustment can be restricted to a benchmark case. Eight heuristics are classified with a naive anchor, one of those heuristics is purely naive ($\beta = 0$). One heuristic has a fundamental anchor. For the other heuristics, the sample average price generally gets little weight: $1 - \alpha_1 - \alpha_2$ ranges from 0.01 to 0.6 and is on average 0.3. The anchors of 49 heuristics are not equivalent to a benchmark and are therefore classified as “mixed”. The adjustment is trend-following ($\beta > 0$) for all but one first-order heuristic. The nonzero values of β range from 0.4 to 1.3 and the average is 0.8.¹⁶ For ten heuristics, β is larger than 1, indicating a strong trend-following adjustment.

Figure 7 illustrates the estimated coefficients of the first-order heuristics in the three-dimensional space $(\alpha_1, \alpha_2, \beta)$. Each dot in the prism represents one of the estimated heuristics.¹⁷ Even though most first-order heuristics cannot be fully restricted to a benchmark rule, the prism shows a clear clustering of heuristics. A majority of 33 out of 58 heuristics (57%) has a relatively high value of α_1 and a low value of α_2 . These anchors are close to naive, giving a high weight to the last observed market price and some weight to the sample average price. Combined with a positive value for β , these rules are almost pure trend-following. A group of 19 heuristics (33%) shows the reverse pattern: a low value of α_1 and a relatively

¹⁶Remarkably, 0.4 and 1.3 are exactly the values of the extrapolation coefficient that are used for the two trend-following rules in HSM of Section 6. The rules in this model were based on estimates of forecasting behavior in Hommes et al. (2005a).

¹⁷Five coefficient vectors with a negative estimate for α_2 fall outside the prism. They are excluded from the plot to make it easier to distinguish the dots for the coefficient vectors that lie within the prism.

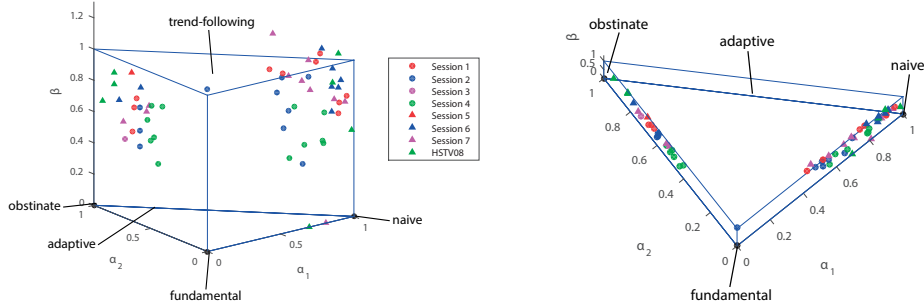


Figure 7: Estimated coefficient vectors of first-order heuristics (excluding outliers)

high value of α_2 . These participants base their anchor on their last own prediction and the sample average price, ignoring the last observed market price. Adaptive expectations are not observed, probably because this forecasting rule does not work well when markets are as unstable as in the experiment. There are no clear differences between sessions in the estimated first-order heuristics.

The first-order heuristics that we estimated for the individual predictions in HSTV08 are also largely comparable to the heuristics in our large-group experiment. The estimated general forecasting rule can be successfully restricted to a first-order heuristic for 10 out of 36 rules (28%). Two first-order heuristics were estimated in session S1, S2, S5 and S6, and one first-order heuristic was estimated in session S3 and S4. Two heuristics are classified with a naive anchor ($\alpha_1 = 1$). One heuristic has an obstinate anchor ($\alpha_2 = 1$), meaning that the participant only considered his own past prediction and not the last observed market price or the sample average price. The other seven heuristics have a mixed anchor. Again, the weight of the sample average price is generally low, with an average of 0.15. For eight out of ten first-order heuristics, the adjustment is trend-following. The nonzero values of β range from 0.6 to 1.2 and the average is 0.7. Both the range and the average of β are thus somewhat smaller than in the large groups. The prism in Figure 7 also shows the estimated coefficients of the ten first-order heuristics of

HSTV08.¹⁸ Three heuristics are part of the cluster of almost pure trend-following rules (a relatively high value of α_1 , a low value of α_2 and a positive β). Three heuristics are part of the other cluster with a low value of α_1 and a relatively high value of α_2 . Adaptive expectations are not found, just as in the large groups.

5 Formation of the first bubble

In our large-group experiment, bubbles emerge faster than in the smaller markets of HSTV08. The bubbles in large markets already peak after about 15 periods, while the bubbles in small markets take around 30 periods to reach the highest price. Still, we did not find substantial differences in the estimated heuristics for both experiments. This could be due to the fact that these estimates are based on the last 45 periods. The price oscillations that occur after the first bubble bursts are a large part of that. These oscillations call for a different prediction strategy than the trend-following behavior that causes the first bubble. Therefore, we now examine the formation of the first bubble in more detail.

5.1 Growth rates

If prices evolve according to a rational bubble, $p_t = p^f + R^t c$, the growth rate of prices is $R = 1.05$. Defining $q_t = \ln(p_t - p^f)$, we must have that $q_{t+1} - q_t = \ln(R)$. We calculate $q_{t+1} - q_t$ and its mean \bar{q} for the first large bubble¹⁹. The implied growth rate is then $\hat{R} = e^{\bar{q}}$. Note that q_t only exists when the price is above the fundamental value. Moreover, a price between $p^f = 60$ and 61 gives a negative value of q^t , leading to a relatively large positive value of $q_{t+1} - q_t$, causing \bar{q} and \hat{R} to be very high. Therefore, the first observation that we take into account satisfies $p_t > (p^f + 1)$. Because the upper bound on predictions flattens the last part of the bubble, the last observation that we take into account satisfies $p_t - p_{t-1} > p_{t-1} - p_{t-2}$.

¹⁸Three coefficient vectors with an estimate of $\alpha_1 > 1$ fall outside the prism and are excluded from the plot.

¹⁹Session L5 does not exhibit large bubbles, so we simply take the first (small) bubble. We disregard session S1, because there are no clear bubbles in that session.

Table 4: Implied growth rates of first (large) bubbles

	\bar{q}	\hat{R}	p -value	Sample
Large groups				
Session L1	0.659	1.932	0.004	2-12
Session L2	0.486	1.626	0.000	3-13
Session L3	0.524	1.690	0.022	28-33
Session L4	0.449	1.567	0.000	4-15
Session L5	0.331	1.392	0.006	3-10
Session L6	0.422	1.525	0.000	3-16
Session L7	0.657	1.928	0.001	4-13
Average	0.504	1.666		
Small groups (HSTV08)				
Session S2	0.225	1.253	0.000	6-27
Session S3	0.292	1.339	0.000	6-29
Session S4	0.310	1.363	0.001	6-21
Session S5	0.718	2.049	0.045	29-37
Session S6	0.307	1.359	0.013	22-28
Average	0.370	1.473		

Note: \bar{q} is the mean of $q_{t+1} - q_t$ for the given sample, and $\hat{R} = e^{\bar{q}}$ is the implied growth rate. The p -value corresponds to a t -test of whether \hat{R} differs significantly from $R = 1.05$.

Table 4 shows the results of this calculation. The implied growth rates \hat{R} differ significantly from $R = 1.05$ at the 5% level, so the bubbles cannot be called rational. Furthermore, the implied growth rates in the large groups are larger than in the small groups, except for session S5, which has a relatively high growth rate ($\hat{R} = 2.049$) because the first half of the experiment is ignored in the calculation since $p_t < (p^f + 1)$. Unsurprisingly, the bubble in the stable session L5 grows slower ($\hat{R} = 1.392$) than the bubbles in the unstable large groups. However, the implied growth rate of the stable session L5 is similar to the bubbles in the unstable small groups, as these take longer to form. Because of these two exceptions, i.e. the fast growth in session S5 and the slow growth in session L5, the difference in means of the large and small groups is not statistically significant (p -value = 0.226), even though it seems that bubbles grow faster in large groups than in small groups.²⁰

²⁰Disregarding the exceptional sessions L5 and S5, the mean growth rate is 1.711 in the large groups and 1.329 in the small groups. This difference is statistically significant (p -value = 0.001).

5.2 Initial predictions and prices

It is interesting to look at the two initial predictions p_1^e and p_2^e , because they are made without any information about past prices. Therefore, they may be informative about the participants initial expectations of the price dynamics. Furthermore, the third prediction p_3^e is submitted after the first price p_1 becomes known. Note that the first two prices p_1 and p_2 actually result from the predictions p_2^e and p_3^e . These predictions determine the first price increase and could thus initiate an upward trend. Tables 7 and 8 in Appendix B give the average initial predictions, their standard deviations, and the initial prices for our large-group experiment and the small-group experiment of HSTV08.

On average, the initial predictions and prices are higher in the large groups than in the small groups. However, the difference is not observed in all sessions. In the large groups, p_1 is around 50 and p_2 is around 55. This is comparable to session S1, S3 and S5 of the small groups. Initial prices are lower in the other small groups (session S2, S4 and S6): p_1 is around 35 and p_2 is around 40. This difference does not seem to explain the faster formation of bubbles in the larger groups.

Many participants seem to expect an upward trend from the beginning. In the large groups, p_2^e is higher than p_1^e for 71% of the participants, in the small groups this holds for 69%. There is however no clear difference between sessions where the bubble forms fast and sessions where it does not. For example, in session S6, all six participants predicted a higher price for the second period than for the first period, but the large bubble does not form until after twenty periods.

The expectation of an upward trend is immediately fulfilled in all large groups. On average, predictions increase in the first three periods: $\bar{p}_1^e < \bar{p}_2^e < \bar{p}_3^e$. Moreover, the increase from \bar{p}_2^e to \bar{p}_3^e is larger than the increase from \bar{p}_1^e to \bar{p}_2^e . This indicates that the trend takes off right away. The same pattern is not observed in the small groups: $\bar{p}_1^e < \bar{p}_2^e < \bar{p}_3^e$ does not hold in session S1 and S2 and $\bar{p}_3^e - \bar{p}_2^e$ is not larger than $\bar{p}_2^e - \bar{p}_1^e$ in session S1, S5 and S6. As a result, the initial price increase $p_2 - p_1$ is bigger in the large groups than in the small groups, with the exception of session

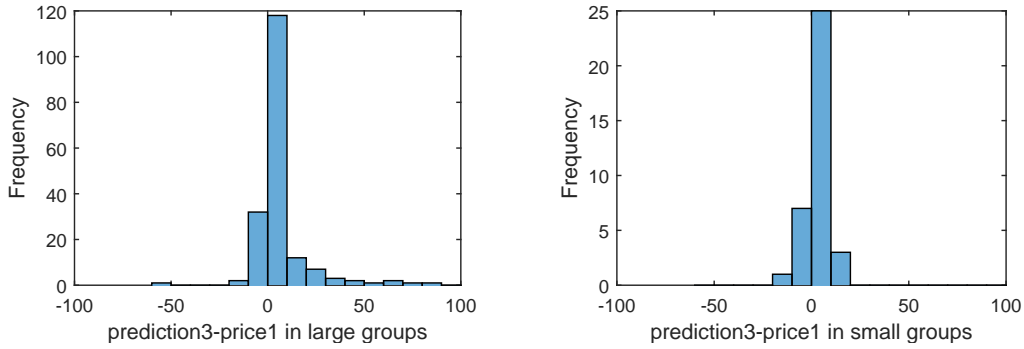


Figure 8: Histograms of $p_3^e - p_1$ in large groups (left) and small groups (right)

L7. It would be an intuitive idea that a large bubble can form faster if the first price increase is bigger, but this is not completely in line with the results of the experiments. For instance, session L5 has a relatively large first price increase, but the bubbles remain small throughout the whole session. In contrast, the first price increase in session L7 is relatively small, but a large bubble forms quickly.

It appears that once the first price becomes known, it is used as a coordination device. In all sessions, the individual predictions for the first two periods are much more dispersed than the predictions for the third period. This is also reflected in the standard deviation of the predictions, which is generally lower for the third period than for the first and second period. Many participants adjust their expectations for the third period upwards when they learn that the price in the first period was higher than they expected, and vice versa.²¹

Despite the coordination in period 3, there are one or two participants in each session that predict a relatively high price. Figure 8 shows histograms of $p_3^e - p_1$ in large and small groups. Most predictions in period 3 are around the first realized price p_1 , usually slightly higher. But in the large groups, there are also some predictions with $p_3^e - p_1 > 20$. These predictions are 1.5–3 times higher than the

²¹It is noteworthy that there are always some participants in the large groups (between one and six per session) that start out with two very low predictions. These predictions are anchored around the mean dividend, $\bar{y} = 3$, so the participants probably misunderstand the value of the asset. Once p_1 turns out to be much higher, they have a large upward adjustment in their prediction for the third period. This seems to make the trend in the large groups stronger from the beginning. This behavior is not observed in the small groups.

average prediction without these outliers. They increase the realization of p_2 and thus the first price increase $p_2 - p_1$ by 2.8 on average, and the increase per session ranges from 0.7 to 5. If this stronger trend is extrapolated, it might contribute to the faster formation of a bubble. These relatively high predictions in period 3 do not occur in the small groups.²² Although it looks like there are more extreme predictions in the large groups, a two-sample Kolmogorov-Smirnov test does not reject the null hypothesis that both samples are drawn from the same distribution (p -value = 0.461). This suggests that the faster formation of bubbles is also not explained by more extreme predictions in large groups.

5.3 Trend-following rules

Looking at individual predictions, it seems that the first bubble is caused by trend extrapolation. For each participant, we thus estimate the trend extrapolation coefficient γ in the following trend-following heuristic:

$$p_{i,t+1}^e = p_{t-1} + \gamma(p_{t-1} - p_{t-2}) + \varepsilon_t. \quad (16)$$

All regressors are defined from period $t = 3$. The last period we take into account is the peak of the first large bubble. We focus on sessions where a large bubble is formed immediately: session L1, L2, L4, L6 and L7 of our large-group experiment and session S2, S3, S4 and S6 of HSTV08. The estimation results are presented in Table 9 in Appendix B.

On average, the estimated trend extrapolation coefficient γ is higher in the large groups than in the small groups. However, the difference in means is not statistically significant (p -value = 0.616). The result does not hold true for each session individually. For instance, the average γ in session S3 is larger than in all

²²Relatively high predictions are also observed occasionally in later periods in both the large and the small groups. However, these predictions do not cause the bubbles. In the large groups, the extreme predictions that have a substantial effect on the market price only occur just before the first peak, when the price is already rising rapidly, or after the crash. In the small groups, individual predictions have more impact on the market price, so extreme predictions cause price jumps, but do not lead to bubbles.

large groups except session L4.

Furthermore, simulations with the estimated trend-following rules do not always correspond to the price patterns in the experiment. Taking the initial prices and predictions from the large-group experiment, γ needs to be about 1.45–1.5 to produce a bubble that is similar to the typical first bubble in the experiment. For session L1 and L7, the average estimated γ is smaller than this, while the average estimated γ for session L4 is larger. For the small-group experiment, γ should be around 1.43. This is larger than the average estimated γ in session S2, S4 and S6, but smaller than in session S3. There is no clear explanation for these differences. For example, the estimated γ 's in session L1 and L7 are almost all lower than in session L2, L4 and L6. Still, a large bubble forms quickly in all these sessions.

All in all, the results do not fully correspond to the intuition that a larger trend extrapolation coefficient leads to larger and faster bubbles. Hence, this trend-following rule alone cannot explain the faster formation of the first bubbles that we observed in the experiments.

6 Heuristics switching model

6.1 A learning model with evolutionary selection

The previous sections discussed some important observations about the individual forecasting behavior in the experiment. The prediction strategies of many participants can be described by a simple linear forecasting rule. However, as participants learn during the experiment, they sometimes switch to a different strategy. Moreover, there is heterogeneity in prediction strategies, both between and within sessions. A heterogeneous expectations model such as the HSM is based on simple heuristics, but takes adaptive learning and strategy switching into account. Therefore, the HSM is a very suitable model to analyze the experimental data.

The main idea of the HSM is that participants can choose between a number of simple heuristics to make a price prediction. In each period, participants evaluate

the past performance of all heuristics, measured by the quadratic forecast error. Then evolutionary selection takes place, meaning that participants tend to switch to better performing rules. Hence, the impact that the different rules have on the market price changes over time. Of course, the realized market price in turn affects the impacts of the heuristics. The HSM is thus a dynamic model that exhibits path dependence, since a different pattern of the market price or a different distribution of impacts in the first couple of periods leads to a different course of both prices and impacts in later periods. This property makes it possible for the HSM to explain very different market outcomes in the same market setting.

We use the same model setup as in Anufriev and Hommes (2012b), which extends the switching model of Brock and Hommes (1997). Anufriev and Hommes (2012b) use the HSM for one-period-ahead simulations of the asset pricing LtFEs of Hommes et al. (2005a) and HSTV08. In Anufriev and Hommes (2012a), they reproduce the price patterns in those experiments with 50-period-ahead simulations of the HSM. Slightly different versions of the HSM are used to analyze LtFEs in different settings, i.e. positive versus negative expectations feedback (Bao et al., 2012; Anufriev et al., 2013) and a New-Keynesian macro-model (Assenza et al., 2014). There are large differences in the observed behavior and the resulting market outcomes in all these experiments. Still, the HSM is able to capture these different patterns and fit the experimental data quite well. In particular, the HSM outperforms models with homogeneous expectations. This indicates that the flexibility of a heterogeneous expectations model such as the HSM is an important advantage. By using the same specification of the HSM as in Anufriev and Hommes (2012b), it is tested if this model is also suitable for the asset pricing experiment with larger markets.

6.2 Model specification

Participants can choose among H heuristics that give a two-period-ahead price prediction. The prediction for time $t + 1$ of each heuristic h is a function of the

available information at time t :

$$p_{h,t+1}^e = f_h(p_{t-1}, p_{t-2}, \dots; p_{h,t}^e, p_{h,t-1}^e, \dots). \quad (17)$$

The realized market price depends on the population-weighted average $\bar{p}_{t+1}^e = \sum_{h=1}^H n_{h,t} p_{h,t+1}^e$ of the heuristics. The weight $n_{h,t}$ is called the impact of heuristic h , which indicates the percentage of participants using this heuristic at time t . The market price is then determined in the same way as in the experiment:

$$p_t = \frac{1}{1+r} (\bar{p}_{t+1}^e + \bar{y}). \quad (18)$$

The impact of each heuristic depends on its past performance. The performance measure is based on the quadratic forecast error. Note that the quadratic forecast error was also used to determine the earnings in the experiment, so that the performance measure is in line with the incentive structure in the experiment. The performance measure up to time $t-1$ is defined as

$$U_{h,t-1} = -(p_{t-1} - p_{h,t-1}^e)^2 + \eta U_{h,t-2}. \quad (19)$$

The parameter $\eta \in [0, 1]$ can be interpreted as the memory. For $\eta = 0$, the performance is only based on the most recent forecast error. If $\eta > 0$, all past forecast errors affect the performance, with exponentially decaying weights.

The impact of heuristic h at time t is determined by a discrete choice model with asynchronous updating:

$$n_{h,t} = \delta n_{h,t-1} + (1 - \delta) \frac{\exp(\beta U_{h,t-1})}{\sum_{h=1}^H \exp(\beta U_{h,t-1})}. \quad (20)$$

The parameter $\delta \in [0, 1]$ gives some inertia in the evolution of the impacts. This reflects the observation in the experiment that not all participants change their prediction strategy in every period or at the same time. δ represents the percentage

of participants that stick to their previous strategy, regardless of the performance. If $\delta = 1$, the initial impacts of the heuristics never change. If $\delta = 0$, all participants consider to change their strategy based on past performance. In general, $1 - \delta$ is the fraction of participants that update their strategy according to the discrete choice model. Note that this does not necessarily mean that they all switch to the most successful strategy. This depends on how sensitive participants are to differences in performance, measured by the intensity of choice $\beta \geq 0$. The higher the intensity of choice, the faster they switch to more successful heuristics. If $\beta = 0$, the impacts converge to an equal distribution. In the extreme case that $\beta = \infty$, all participants that update their strategy switch to the best performing heuristic.

6.3 Implementing the model

Before the HSM can be used, the H different heuristics should be chosen, the parameters β , η and δ should be fixed, and the model should be initialized. Anufriev and Hommes (2012b) based their selection of heuristics on the forecasting rules that were observed in the experimental data. To keep the model relatively simple, they include only four heuristics: adaptive expectations (ADA), a weak trend-following (WTR) and a strong trend-following (STR) rule, and a learning anchoring and adjustment (LAA) heuristic. The heuristics are defined as follows:

$$\text{ADA: } p_{1,t+1}^e = 0.65p_{t-1} + 0.35p_{1,t}^e, \quad (21a)$$

$$\text{WTR: } p_{2,t+1}^e = p_{t-1} + 0.4(p_{t-1} - p_{t-2}), \quad (21b)$$

$$\text{STR: } p_{3,t+1}^e = p_{t-1} + 1.3(p_{t-1} - p_{t-2}), \quad (21c)$$

$$\text{LAA: } p_{4,t+1}^e = 0.5(p_{t-1}^{av} + p_{t-1}) + (p_{t-1} - p_{t-2}), \quad (21d)$$

where p_{t-1}^{av} is the sample average of all past prices up to (and including) time $t - 1$.²³

The parameter values for the benchmark HSM are $\beta = 0.4$, $\eta = 0.7$ and $\delta = 0.9$.

²³Just as in the experiment, the predictions should be between 0 and 1000. If heuristic h leads to a prediction lower than 0, $p_{h,t+1}^e$ is set to 0. Similarly, if the heuristic leads to a prediction higher than 1000, $p_{h,t+1}^e$ is set to 1000.


Four periods are needed to initialize the model. Given two initial prices p_1 and p_2 , the four heuristics generate predictions for period 4. The initial prediction of the ADA heuristic for period 3 is the fundamental price, $p^f = 60$. The initial distribution of impacts in period 3 is equal for all heuristics: $n_{h,3} = 0.25$ for $h = \{1, 2, 3, 4\}$. The performance measure for period 3 is zero. The predictions for period 4 and the impacts in period 3 are then used to calculate the price in period 3. In period 4, the same initial impacts are used, since past performance is not well defined yet. After these four periods, the model is fully initialized. In every period $t \geq 5$, first the performance measure $U_{h,t-1}$ is updated, then the new impacts $n_{h,t}$ are determined, the predictions $p_{h,t+1}^e$ are computed, and the new market price p_t is realized.

7 Simulations with the heuristics switching model


7.1 One-period-ahead simulations

To get insight in the adaptive learning and strategy switching, we use the experimental data for one-period-ahead simulations of the HSM. For each session, the two initial prices are chosen to correspond to the first two prices in that session. The prices in the experiment are also used to compute the performance measure and the predictions of the four heuristics. Each period, the model thus simulates a price for period t , but this price is not used in period $t+1$ to simulate the next price. By using the experimental prices instead, the model uses exactly the same information about past prices that was available to the participants in the experiment.


Figure 9 visualizes the results of the HSM. For each session, the figures show the experimental and simulated prices, the forecasts and the forecast errors of the four heuristics, and the evolution of impacts. The figures suggest that the fit of the HSM is reasonable. Not surprisingly, the simulated prices follow the same bubble-and-crash pattern as in the experiment. However, when the price rises rapidly during a bubble, the simulated price is in most cases lower than the experimental




(a) Session L1



(b) Session L2




(c) Session L3




(d) Session L4


Figure 9: One-period-ahead simulations of the HSM (session L1-L4)



(a) Session L5



(b) Session L6



(c) Session L7

Figure 9: One-period-ahead simulations of the HSM (session L5-L7)

Left: experimental prices and simulated prices (top panel), forecasts of the four heuristics (bottom left panel) and corresponding forecast errors (bottom right panel). The dashed line in the top and bottom left panel indicates the fundamental value of 60. Note that the vertical axis for session L5 runs up to 200 instead of 1000. *Right:* impacts of the four heuristics.

price, and vice versa during a crash. It seems that the four heuristics are not able to fully capture the extremely large price changes. The model fit is better when prices are relatively stable, mainly in the beginning of a session and if there is some stabilization in the last periods. The simulated prices are then quite close to the realized prices in the experiment.

In all sessions, we see that the four heuristics make similar predictions. This is in line with coordination of individual expectations. Forecast errors are defined as the experimental price minus the predicted price, so that a positive error means that the prediction is too low. The plots of the forecast errors also indicate that the heuristics produce forecasts that are generally too low during a bubble and too high during a crash. The forecast errors are correlated: all heuristics either overpredict or underpredict at the same time. The order of magnitude of the errors is in line with the forecast errors of the participants in the experiment.

The evolution of impacts in the six sessions with large bubbles is similar. The impact of the STR rule immediately starts increasing from 25% up to 80% or more, while the impacts of the other heuristics gradually drop. The dominance of the STR rule reflects the strong upward trend in prices. But this heuristic misses the trend reversals when the market crashes and when a new bubble starts to form. The LAA rule also gives weight to the average price in the market and is therefore better at predicting trend reversals and price oscillations. Hence, the impact of STR decreases after the first crash and the LAA rule gains impact, rising to more than 80%. In session L1, L2, L4 and L7, the amplitude of the bubbles decreases and prices slowly seem to stabilize. The impact of the ADA heuristic increases at the expense of LAA, explaining the price stabilization. There is no stabilization in session L3 and L6, where the impact of the LAA rule keeps increasing and enforces persistent price oscillations.

In session L5, the pattern in the first thirty periods is comparable to the sessions with large bubbles. The STR rule dominates in the beginning, but is taken over by the LAA heuristic when the price starts oscillating. The impacts grow to a

maximum of around 65% for STR and 75% for LAA. These impacts are not as large as in the other sessions, because the bubbles are smaller and the price stabilizes more quickly. This causes the impact of the ADA rule to rise up to 85% and dominate the other rules in the last ten periods, reinforcing price stabilization.

Our results with large groups are in line with the results of Anufriev and Hommes (2012b) for small groups. Session L5 is similar to the oscillating markets of HSTV05, and the other sessions are comparable to the markets with large bubbles of HSTV08.

7.2 Forecasting performance

The model fit can be evaluated by the mean squared error (MSE). We compare six homogeneous expectation rules and three versions of the HSM. The homogeneous rules are fundamental expectations, naive expectations, and the four heuristics of the HSM. The versions of the HSM are the benchmark model, a HSM with fixed fractions ($\delta = 1$), and a HSM that is fitted to the experimental data.

The fitted HSM is found by using a grid search to minimize the MSE over the parameters β , η and δ . The optimal parameter values can be found in Table 5. In most sessions, the optimal value of β is close to or slightly higher than the benchmark value of 0.4. Recall that a larger β means that participants switch faster to more successful rules. Still, the difference with the benchmark case is small even for the highest value of $\beta = 1.72$, bearing in mind that the grid search considered values for β in the interval $[0, 10]$.²⁴ The optimal value of η is lower than the benchmark value of 0.7 in all but one session. This indicates that the memory of the performance measure is slightly shorter, so that more distant prediction errors get less weight. In all sessions, the optimal value of δ is somewhat lower than the benchmark value of 0.9, suggesting that the inertia is somewhat smaller

²⁴Anufriev and Hommes (2012b) found that $\beta = 10$ is the optimal value for four of the fourteen sessions of Hommes et al. (2005a) and for one of the six sessions of HSTV08. In addition, they found a relatively high value of $\beta = 5.93$ for another session of HSTV08. The other values of β and the optimal values of η and δ that they found for the sessions of HSTV08 are comparable to our results.

Table 5: MSE and parameters fitted HSM

Session	L1	L2	L3	L4	L5	L6	L7
Homogeneous models							
FUN	150,194	136,948	117,585	131,006	456	147,455	152,423
NAI	27,275	32,837	20,217	28,928	209	25,825	27,067
ADA	42,515	45,084	33,072	38,881	295	41,589	40,376
WTR	18,184	25,325	11,907	23,311	157	15,917	19,006
STR	17,878	23,630	6,182	22,199	238	11,721	18,518
LAA	12,732	16,613	9,140	17,209	119	11,994	13,425
Heuristics switching models							
Fixed frac.	13,475	18,254	8,344	17,559	120	10,789	13,978
Benchmark	6,673	11,530	3,781	13,028	93	5,144	8,352
Fitted	<i>5,735</i>	<i>10,851</i>	<i>3,055</i>	<i>9,688</i>	<i>60</i>	<i>4,874</i>	<i>6,869</i>
Optimal parameter values of the fitted HSM							
$\beta \in [0, 10]$	0.49	0.59	0.38	1.05	0.37	0.86	1.72
$\eta \in [0, 1]$	0.39	0.69	0.93	0.56	0.36	0.38	0.57
$\delta \in [0, 1]$	0.77	0.81	0.87	0.48	0.66	0.78	0.72

Note: For each session, the lowest MSE is shown in italic and it is always the MSE of the fitted HSM. The benchmark HSM has the second lowest MSE, which is shown in bold. The MSE of the best performing homogeneous model is also shown in bold for comparison.

and a somewhat larger fraction of participants considers to update their rule in each period.

For each model, the MSE is calculated for $t = 5, \dots, 50$. This is because the HSM takes four periods to fully initialize. It can also be seen as a learning phase. Table 5 shows the MSE of the nine models in every session. The fundamental prediction performs much worse than all other models, especially in the sessions with large bubbles where prices deviate far from the fundamental. The best performing homogeneous model in the sessions where prices stabilize in the last periods (session L1, L2, L4, L5 and L7) is the LAA heuristic. In the sessions without stabilization (session L3 and L6), the STR rule performs better.

The fitted HSM has the minimal MSE of the HSMs by definition. Ignoring the fitted model, the benchmark HSM always has the lowest MSE. The differences of the benchmark HSM compared to the homogeneous models and the fixed fractions HSM are substantial. In most sessions, the improvement of fit of the fitted HSM compared to the benchmark HSM is not as large as when the benchmark HSM is compared to the best homogeneous model. This suggests that the results of the HSM are not very sensitive to the parameters. Nonetheless, the HSM with

fixed fractions has a higher MSE than the best performing homogeneous model, except in session L6. It seems that the model fit significantly improves due to the flexibility of the strategy switching.

Compared to the other sessions, the MSE of all models is much lower in session L5. This illustrates that it is easier to predict prices in a relatively stable market. The very high MSEs in the other sessions show that it is extremely hard to forecast the large bubbles and crashes.

When we compare our results to the findings of Anufriev and Hommes (2012b), we see that the fit of the HSM in our experiment is generally worse than in HSTV08. This is due to the rapid price increases and decreases in our large groups, that even the STR rule cannot accurately capture. Hence the more large bubbles occur in a session, the higher the MSE.

7.3 50-period-ahead simulations

Next, we consider 50-period-ahead simulations of the HSM. The goal is to qualitatively replicate the price patterns in the experiment as accurately as possible. These simulations do not use experimental data. We only specify values for β , η and δ , initial prices for period 1 and 2, and initial impacts for period 3 and 4. These parameters and initial values are found by trial and error.

Figure 10 shows 50-period-ahead simulations of the typical price pattern with large bubbles (similar to session L1, L2, L4, L6 and L7) and the price pattern with small bubbles (similar to session L5). The initializations that were used for these simulations can be found in Table 6.

The formation of large bubbles strongly depends on the initial prices and impacts. Choosing two prices that are relatively far apart gives a strong upward trend and triggers a quick formation of a large bubble. The initial impacts of the heuristics determine the course of the market price. A very large impact of the STR rule reinforces the strong trends and leads to large bubbles. Starting with a small impact of the LAA heuristic allows this rule to gain impact over time, leading to




Figure 10: 50-period-ahead simulations of the HSM
Left: typical price pattern with large bubbles (similar to session L1, L2, L4, L6 and L7). *Right:* price pattern with small bubbles (similar to session L5). The dashed line indicates the fundamental value of 60.

Table 6: Parameters and initial values for 50-period-ahead simulations

	Parameters			Prices		Impacts ($t = 3, 4$)			
	β	η	δ	p_1	p_2	$n_{1,t}$	$n_{2,t}$	$n_{3,t}$	$n_{4,t}$
Large bubbles	0.4	0.7	0.9	30	80	0	0	0.99	0.01
Small bubbles	0.4	0.5	0.9	40	45	0	0.2	0.7	0.1

persistent oscillations with decreasing amplitude. We simply used the benchmark parameter values to simulate the large bubbles. Changing β barely has any effect. Slightly lowering η and δ makes the bubbles form a bit faster.

To simulate the small bubbles, the initial prices are relatively low and close to each other, to give rise to a weak upward trend. The STR rule still dominates in the beginning so that a bubble forms, but now the WTR rule and the LAA heuristic have larger initial impacts to make the oscillations smaller and converging. We used the benchmark values for β and δ , but a lower value of η . The shorter memory ensures that the bubbles remain small.

The HSM is thus able to qualitatively match the typical long run price patterns. However, it is difficult to perfectly replicate the frequency and amplitude of the oscillations. Unfortunately, we were not able to replicate the atypical price pattern of session L3. With this specification of the HSM, it seems impossible to simulate a market that starts with a small bubble and then suddenly forms several large bubbles.²⁵

²⁵Recall that the use of a strong trend-following rule leads to increasing bubbles (see Figure 2c). It is possible to simulate a price pattern that is qualitatively similar to session L3 by using a trend-following rule with a slightly larger trend extrapolation coefficient as in the HSM, i.e. $\gamma = 1.4$ instead of $\gamma = 1.3$.

8 Conclusion

Previous studies have shown that large bubbles caused by coordination on a trend-following prediction strategy often occur in asset pricing experiments with small groups of ten participants or less. Will coordination on these bubbles also occur in larger groups? We increase the market size in the learning-to-forecast experiment of HSTV08 to the size of a typical computer lab. The experiment consists of seven sessions with 21 to 32 participants per market, a group size larger than in most laboratory experiments.

Our paper shows that the empirical answer to the above question is yes: coordination on large bubbles occurs in six out of seven markets. Prices increase rapidly up to fifteen times the fundamental value. Participants seem to coordinate on a trend-following prediction strategy, causing the first large bubble. We also observe heterogeneity in expectations and switching between prediction strategies. Simple linear forecasting rules provide a good description of the behavior of many participants. The strategies resemble simple benchmark heuristics, such as trend-following rules and anchoring and adjustment.

The bubbles in our large groups occur even faster and more frequently than in the smaller markets of HSTV08. The estimated heuristics over the whole course of both experiments are however not substantially different. Focusing on the formation of the first bubble, we find that the growth rate of prices and the estimated trend extrapolation coefficients are on average higher in large groups, but the differences are not statistically significant. We also observe on average higher initial predictions and prices in large groups. Still, these observations on their own cannot explain the differences in price dynamics. But a combination of factors, such as higher initial prices and stronger trend extrapolation in large groups, might contribute to the faster formation of large bubbles.

Heuristics switching models are suitable to analyze the experimental data, because these heterogeneous expectations models incorporate adaptive learning and

switching between simple strategies. We illustrate with both one-period-ahead and 50-period-ahead simulations that the HSM is able to capture the price patterns observed in the experiment. The evolution of the impacts of the four heuristics highlights that strong trend extrapolation is important for the formation of a large bubble, that the flexible anchoring and adjustment rule promotes persistent oscillations, and that adaptive expectations lead to price stabilization and convergence. The HSM has better forecasting performance than several benchmark homogeneous expectation rules. It appears that allowing for learning and strategy switching substantially improves the model fit. Heterogeneity in expectations seems crucial to explain the price dynamics.

Coordination, heterogeneity in expectations, the use of simple forecasting heuristics and strategy switching are all observed in multiple LtFEs (see e.g. Hommes et al., 2005a, 2008; Heemeijer et al., 2009; Anufriev and Hommes, 2012b; Assenza et al., 2014). The larger markets in our experiment did not result in individual forecast errors cancelling out at the aggregate level. Hence, expectations cannot be called rational in the sense of Muth (1961). The results of our large-group experiment are thus largely comparable to the results of HSTV08 and other related experiments with small groups.

These results are reassuring for financial market experiments with small groups. However, a caveat is in order as this conclusion is only based on seven large-group observations in a simple asset pricing environment. More experimental data on aggregate and individual behavior in large groups are necessary for more accurate empirical evidence and to test the robustness of bubbles in large groups.

References

- Anufriev, M. and Hommes, C. (2012a). Evolution of market heuristics. *The Knowledge Engineering Review*, 27(2):255–271.
- Anufriev, M. and Hommes, C. (2012b). Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments. *American Economic Journal: Microeconomics*, 4(4):35–64.
- Anufriev, M., Hommes, C., and Philipse, R. (2013). Evolutionary selection of expectations in positive and negative feedback markets. *Journal Of Evolutionary Economics*, 23(3):663–688.
- Assenza, T., Heemeijer, P., Hommes, C., and Massaro, D. (2014). Managing self-organization of expectations through monetary policy: a macro experiment. CeNDEF Working Paper 2014-07, University of Amsterdam.
- Bao, T. and Hommes, C. (2015). When speculators meet constructors: Positive and negative feedback in experimental housing markets. CeNDEF Working Paper 2015-10, University of Amsterdam.
- Bao, T., Hommes, C., and Makarewicz, T. (2014). Bubble formation and (in)efficient markets in learning-to-forecast and -optimize experiments. CeNDEF Working Paper 2014-01, University of Amsterdam.
- Bao, T., Hommes, C., Sonnemans, J., and Tuinstra, J. (2012). Individual expectations, limited rationality and aggregate outcomes. *Journal of Economic Dynamics and Control*, 36(8):1101–1120.
- Barberis, N., Greenwood, R., Jin, L., and Shleifer, A. (2016). Extrapolation and bubbles. NBER Working Paper 21944, National Bureau of Economic Research.
- Bossaerts, P. and Plott, C. (2004). Basic principles of asset pricing theory: Evidence from large-scale experimental financial markets. *Review of Finance*, 8:135–169.

- Bottazzi, G., Devetag, G., and Pancotto, F. (2011). Does volatility matter? expectations of price return and variability in an asset pricing experiment. *Journal of Economic Behavior and Organization*, 77:124–146.
- Brock, W. and Hommes, C. (1997). A rational route to randomness. *Econometrica*, 65(5):1059–1095.
- Brock, W. and Hommes, C. (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control*, 22(4):1235–1274.
- Brunnermeier, M. and Nagel, S. (2004). Hedge funds and the technology bubble. *Journal of Finance*, 59(5):2013–2040.
- Campbell, J., Lo, A., and MacKinlay, A. (1997). *The Econometrics of Financial Markets*. Princeton, NJ: Princeton University Press.
- Diederich, J. (2013). *Determinants of the Voluntary Provision of Public Goods: Experimental Investigations*. PhD thesis, Heidelberg University.
- Duffy, J. (2008). Experimental macroeconomics. In Durlauf, S. and Blume, L., editors, *New Palgrave Dictionary of Economics*. New York, NY: Palgrave Macmillan.
- Duffy, J. (2015). Macroeconomics: a survey of laboratory research. In Kagel, J. and Roth, A., editors, *Handbook of Experimental Economics Vol. 2 (forthcoming)*. Princeton, NJ: Princeton University Press.
- Greenwood, R. and Nagel, S. (2009). Inexperienced investors and bubbles. *Journal of Financial Economics*, 93:239–258.
- Hanaki, N., Akiyama, E., and Ishikawa, R. (2016). A methodological note on eliciting price forecasts in asset market experiments. GREDEG Working Paper 2016-02, University of Nice Sophia Antipolis.

- Heemeijer, P., Hommes, C., Sonnemans, J., and Tuinstra, J. (2009). Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. *Journal of Economic Dynamics and Control*, 33:1052–1072.
- Hommes, C. (2001). Financial markets as nonlinear adaptive evolutionary systems. *Quantitative Finance*, 1(1):149–167.
- Hommes, C. (2011). The heterogeneous expectations hypothesis: Some evidence from the lab. *Journal of Economic Dynamics and Control*, 35:1–24.
- Hommes, C., Sonnemans, J., Tuinstra, J., and Van de Velden, H. (2005a). Coordination of expectations in asset pricing experiments. *Review of Financial Studies*, 18(3):955–980.
- Hommes, C., Sonnemans, J., Tuinstra, J., and Van de Velden, H. (2005b). A strategy experiment in dynamic asset pricing. *Journal of Economic Dynamics and Control*, 29:823–843.
- Hommes, C., Sonnemans, J., Tuinstra, J., and Van de Velden, H. (2008). Expectations and bubbles in asset pricing experiments. *Journal of Economic Behavior and Organization*, 67:116–133.
- Isaac, R. and Walker, J. (1988). Group size effects in public goods provision: The voluntary contributions mechanism. *The Quarterly Journal of Economics*, 103(1):179–199.
- Isaac, R., Walker, J., and Williams, A. (1994). Group size and the voluntary provision of public goods: Experimental evidence utilizing large groups. *Journal of Public Economics*, 54:1–36.
- Marimon, R. and Sunder, S. (1993). Indeterminacy of equilibria in a hyperinflationary world: Experimental evidence. *Econometrica*, 61(5):1073–1107.

- Marimon, R. and Sunder, S. (1994). Expectations and learning under alternative monetary regimes: An experimental approach. *Economic Theory*, 4:131–162.
- Muth, J. (1961). Rational expectations and the theory of price movements. *Econometrica*, 29(3):315–335.
- Ricciuti, R. (2008). Bringing macroeconomics into the lab. *Journal of Macroeconomics*, 30:216–237.
- Sims, C. (1980). Macroeconomics and reality. *Econometrica*, 48:1–48.
- Smith, V., Suchanek, G., and Williams, A. (1988). Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica*, 56(5):1119–1151.
- Sonnemans, J. and Tuinstra, J. (2010). Positive expectations feedback experiments and number guessing games as models of financial markets. *Journal of Economic Psychology*, 31:964–984.
- Stöckl, T., Huber, J., and Kirchler, M. (2010). Bubble measures in experimental asset markets. *Experimental Economics*, 13:284–298.
- Tversky, A. and Kahneman, D. (1974). Judgement under uncertainty: Heuristics and biases. *Science*, 185:1124–1130.
- Weimann, J., Brosig-Koch, J., Hennig-Schmidt, H., Keser, C., and Stahr, C. (2012). Public-good experiments with large groups. FEMM Working Paper 120009, Otto-von-Guericke University Magdeburg.
- Williams, A. (2008). Price bubbles in large financial asset markets. In C.R., P. and Smith, V., editors, *Handbook of Experimental Economics Results*. Amsterdam: Elsevier.
- Williams, A. and Walker, J. (1993). Computerized laboratory exercises for microeconomics education: Three applications motivated by experimental economics. *The Journal of Economic Education*, 24(4):291–315.

Xu, B., Cadsby, C., Fan, L., and Song, F. (2013). Group size, coordination, and the effectiveness of punishment in the voluntary contributions mechanism: An experimental investigation. *Games*, 4:89–105.

Appendix

A Experimental instructions

General information

You are a **financial advisor** to a pension fund that wants to optimally invest a large amount of money. The pension fund has two investment options: a risk free investment and a risky investment. The risk free investment is putting all money on a bank account paying a fixed and known interest rate. The alternative risky investment is an investment in the stock market with uncertain return. In each time period the pension fund has to decide which fraction of its money to put on the bank account and which fraction of the money to spend on buying stocks. In order to make an optimal investment decision the pension fund needs an accurate prediction of the price of the stock. As their financial advisor, you have to predict the stock market price during 51 subsequent time periods. Your earnings during the experiment depend upon your forecasting accuracy. The smaller your forecasting errors in each period, the higher your total earnings.

Forecasting task of the financial advisor

The only task of the financial advisors in this experiment is to forecast the stock market index in each time period as accurate as possible. The stock price has to be predicted **two** time periods ahead. At the beginning of the experiment begins, you have to predict the stock price in the first **two** periods. It is very likely that the stock price will be between 0 and 100 in the first two periods. After all participants have given their predictions for the first two periods, the stock market price for the first period will be revealed and based upon your forecasting error your earnings for period 1 will be given. After that you have to give your prediction for the stock market index in the third period. After all participants have given their predictions for period 3, the stock market index in the second period will be revealed and, based upon your forecasting error your earnings for period 2 will be given. This process continues for 51 time periods.

The available information for forecasting the stock price in period t consists of

- all past prices up to period $t - 2$, and
- past predictions up to period $t - 1$, and
- total earnings up to period $t - 2$

Information about the stock market

The stock market price is determined by equilibrium between demand and supply of stocks. The stock market price in period t will be that price for which aggregate demand equals supply. The supply of stocks is fixed during the experiment. The demand for stocks is determined by the aggregate demand of a number of large pension funds active. Each pension fund is advised by a participant of the experiment.

Information about the investment strategies of the pension funds

The precise investment strategy of the pension fund that you are advising and the investment strategies of the other pension funds are unknown. The bank account of the risk free investment pays a fixed interest rate of 5% per time period. The holder of the stock receives a dividend payment in each time period. These dividend payments are uncertain however and vary over time. Economic experts of the pension funds have computed that the average dividend payments are 3 euro per time period. The return of the stock market per time period is uncertain and depends upon (unknown) dividend payments as well as upon price changes of the stock. As the financial advisor of a pension fund you are **not** asked to forecast dividends, but you are only asked to forecast the price of the stock in each time period. Based upon your stock market price forecast, your pension fund will make an optimal investment decision. The higher your price forecast the larger will be the fraction of money invested by your pension fund in the stock market, so the larger will be their demand for stocks.

Earnings

Earnings will depend upon forecasting accuracy only. The better you predict the stock market price in each period, the higher your aggregate earnings. Your earnings in each period will be determined by the following scoring rule:

$$e_t = \max\left\{1300 - \frac{1300}{49}(p_t - p_t^e)^2, 0\right\},$$

where e_t is the number of points you earn in period t , p_t is the realized stock price in period t , and p_t^e is your prediction of the stock price in period t . The earnings table below gives the number of points corresponding to each value of the forecasting error. At the end of the experiment, your total earnings in points will be converted into euros, at an exchange rate of **0.5 euro for 1300 points**.

Payoff Table									
1300 points equal 0.5 euro									
error	points	error	points	error	points	error	points	error	points
0.1	1300	1.5	1240	2.9	1077	4.3	809	5.7	438
0.15	1299	1.55	1236	2.95	1069	4.35	798	5.75	423
0.2	1299	1.6	1232	3	1061	4.4	786	5.8	408
0.25	1298	1.65	1228	3.05	1053	4.45	775	5.85	392
0.3	1298	1.7	1223	3.1	1045	4.5	763	5.9	376
0.35	1297	1.75	1219	3.15	1037	4.55	751	5.95	361
0.4	1296	1.8	1214	3.2	1028	4.6	739	6	345
0.45	1295	1.85	1209	3.25	1020	4.65	726	6.05	329
0.5	1293	1.9	1204	3.3	1011	4.7	714	6.1	313
0.55	1292	1.95	1199	3.35	1002	4.75	701	6.15	297
0.6	1290	2	1194	3.4	993	4.8	689	6.2	280
0.65	1289	2.05	1189	3.45	984	4.85	676	6.25	264
0.7	1287	2.1	1183	3.5	975	4.9	663	6.3	247
0.75	1285	2.15	1177	3.55	966	4.95	650	6.35	230
0.8	1283	2.2	1172	3.6	956	5	637	6.4	213
0.85	1281	2.25	1166	3.65	947	5.05	623	6.45	196
0.9	1279	2.3	1160	3.7	937	5.1	610	6.5	179
0.95	1276	2.35	1153	3.75	927	5.15	596	6.55	162
1	1273	2.4	1147	3.8	917	5.2	583	6.6	144
1.05	1271	2.45	1141	3.85	907	5.25	569	6.65	127
1.1	1268	2.5	1134	3.9	896	5.3	555	6.7	109
1.15	1265	2.55	1127	3.95	886	5.35	541	6.75	91
1.2	1262	2.6	1121	4	876	5.4	526	6.8	73
1.25	1259	2.65	1114	4.05	865	5.45	512	6.85	55
1.3	1255	2.7	1107	4.1	854	5.5	497	6.9	37
1.35	1252	2.75	1099	4.15	843	5.55	483	6.95	19
1.4	1248	2.8	1092	4.2	832	5.6	468	error \geq 7	0
1.45	1244	2.85	1085	4.25	821	5.65	453		

B Tables

B.1 Experimental results

Table 7: Summary statistics large-group experiment

Session	L1	L2	L3	L4	L5	L6	L7
Average price	327.32	323.98	273.05	321.28	59.96	309.42	339.48
Minimum price	28.59	53.74	24.17	49.26	14.12	31.02	50.41
Period min. price	20	1	25	1	18	26	1
Maximum price	876.29	851.28	852.56	868.59	95.48	869.72	918.55
Period max. price	13	13	34	15	11	16	14
Initial prices							
Period 1	55.28	53.74	46.73	49.26	51.55	48.38	50.41
Period 2	61.09	59.54	52.65	55.50	59.19	56.11	53.64
Average initial pred.							
Period 1	50.87	48.00	42.24	44.36	46.80	40.91	42.33
Period 2	55.19	53.57	46.21	48.87	51.28	47.94	50.08
Period 3	61.27	59.64	52.40	55.40	59.28	56.04	53.45
St. dev. initial pred.							
Period 1	21.26	12.91	17.79	24.06	21.42	20.21	21.13
Period 2	23.02	14.35	18.24	34.19	24.05	22.63	24.57
Period 3	16.26	19.21	19.70	14.18	13.58	10.00	6.17

Table 8: Summary statistics small-group experiment (HSTV08)

Session	S1	S2	S3	S4	S5	S6
Average price	115.52	273.45	177.45	306.25	167.07	182.98
Minimum price	48.4	32.86	9.68	9.21	12.56	23.78
Period min. price	44	48	39	31	46	18
Maximum price	228.70	934.54	931.11	954.75	940.16	749.62
Period max. price	28	28	29	22	38	29
Initial prices						
Period 1	52.06	41.53	52.54	36.71	54.52	30.49
Period 2	51.43	45.35	54.68	41.08	59.37	34.52
Average initial pred.						
Period 1	42.50	41.83	51.58	33.43	47.58	24.17
Period 2	51.67	40.61	52.17	35.54	54.25	29.01
Period 3	51.00	44.62	54.42	40.13	59.33	33.25
St. dev. initial pred.						
Period 1	14.05	14.63	15.49	15.33	6.17	20.41
Period 2	14.38	13.07	10.46	15.68	5.69	18.76
Period 3	6.79	3.68	1.28	6.22	6.56	3.01

B.2 Estimated forecasting rules

Table 9: Estimated trend extrapolation coefficient γ

P	Large-group experiment					Small-group experiment (HSTV08)			
	L1	L2	L4	L6	L7	S2	S3	S4	S6
1	1.603	1.774	2.201	1.357	0.578	1.077	2.113	0.868	0.647
2	0.587	1.789	2.380	1.920	0.178	0.527	2.148	0.991	1.094
3	0.642	1.800	1.779	1.638	1.112	1.097	1.656	1.138	1.020
4	1.227	1.769	0.517	2.235	0.888	1.147	2.078	1.119	1.355
5	0.512	0.826	1.512	1.259	0.904	0.891	2.006	0.866	2.058
6	0.323	1.348	1.886	1.403	1.010	1.066	1.795	1.081	0.508
7	1.372	1.346	2.448	2.232	0.924				
8	1.287	1.039	1.577	0.957	0.799				
9	0.916	-2.056	1.673	2.037	1.082				
10	0.053	2.103	1.287	1.911	1.072				
11	1.226	0.854	2.461	1.898	1.113				
12	0.989	1.825	2.390	1.867	1.276				
13	1.084	1.763	2.106	1.361	0.931				
14	1.120	0.617	2.473	1.741	1.249				
15	1.275	2.244	2.579	2.033	0.456				
16	1.189	1.934	2.829	1.718	0.710				
17	1.200	1.840	1.913	1.256	0.737				
18	1.278	1.874	2.027	1.863	1.127				
19	1.108	1.075	2.067	1.316	-0.188				
20	1.103	1.593	2.310	1.552	0.743				
21	0.918	1.672	2.392	1.754	1.117				
22	1.000	1.422	0.866		1.061				
23	1.411	2.098	1.711		1.023				
24	0.985	2.106	2.393		1.158				
25	1.424	1.639	2.838		1.181				
26	1.182	2.232	1.983		0.999				
27			2.114		1.114				
28			2.367		1.486				
29			1.492		1.135				
30			1.688		0.934				
31			3.156		0.638				
32			1.807		0.556				
$\bar{\gamma}$	1.039	1.482	2.038	1.681	0.909	0.968	1.966	1.01	1.114

Note: ‘P’ indicates the number of the participant. The last row gives the average estimated trend extrapolation coefficient $\bar{\gamma}$ for each session.

Table 10: General linear forecasting rules

S	P	α	β_1	β_2	β_3	β_4	γ_0	γ_1	γ_2	γ_3	R^2	BG	W	RR
1	1	141.0	2.198	-2.359	0	0	0	0	0.765	0	0.716	0.098	0.059	0.027
1	2	32.9	1.422	0	0	0	0	-0.476	0	0	0.952	0.821	0.557	0.070
1	3	40.8	1.727	-0.780	0	0	0	0	0	0	0.959	0.859	0.118	0.352
1	4	188.5	0	0	0	-0.392	0	0.868	0	0	0.669	0.412	0.009	0.642
1	5	45.5	0.701	-0.439	0	0	0.479	0	0	0	0.555	0.725	0.014	0.410
1	6	73.5	1.568	-0.680	0	0	0	0	0	0	0.910	0.516	0.701	0.672
1	7	69.2	1.901	-1.459	0	0	0	0	0.349	0	0.782	0.988	0.006	0.181
1	8	70.2	2.829	-1.046	0	0	-0.880	0	0	0	0.853	0.599	0.046	0.458
1	9	137.8	0.836	0	0	0	0.720	-0.913	0	0	0.789	0.622	0.424	0.197
1	10	6.2	0.505	0	-0.169	0	0.490	0	0	0	0.828	0.387	0.057	0.006
1	11	180.1	1.535	-1.024	0	0	0	0	0	0	0.772	0.757	0.504	0.022
1	12	81.1	1.255	-1.871	1.115	-0.413	0.699	0	0	0	0.892	0.917	0.081	0.334
1	13	128.8	0.913	-0.795	0	0	0.522	0	0	0	0.762	0.503	0.210	0.272
1	14	126.5	1.781	-1.048	0	0	0	0	0	0	0.854	0.706	0.219	0.268
1	15	153.3	0	0	0	0	1.115	-0.510	0	0	0.667	0.457	0.210	0.819
1	16	47.2	2.155	-1.004	0.499	0	0	-0.696	0	0	0.810	0.875	0.679	0.912
1	17	119.0	1.702	-1.033	0	-0.363	0	0	0.435	0	0.861	0.410	0.718	0.555
1	18	162.3	0	0	0	0	1.222	-0.653	0	0	0.744	0.730	0.100	0.643
1	19	142.0	0	0	0	0	1.254	-0.625	0	0	0.757	0.238	0.426	0.277
1	20	119.5	1.606	0	0	0	0	-0.932	0	0	0.636	0.077	0.171	0.349
1	21	47.3	1.895	-0.786	-0.489	0	0	0	0	0.355	0.923	0.653	0.568	0.332
1	22	178.9	1.637	0	0	0	0	-0.920	0	0	0.748	0.908	0.032	0.254
1	23	118.0	0.875	0	0	0	0.554	-0.632	0	0	0.790	0.507	0.112	0.487
1	24	190.0	1.366	-0.717	0	0	0	0	0	0	0.717	0.092	0.145	0.055
1	25	248.5	1.478	-1.068	0	0	0	0	0	0	0.475	0.557	0.192	0.443
1	26	100.9	1.855	-1.091	0	0	0	0	0	0	0.864	0.556	0.134	0.110

S	P	α	β_1	β_2	β_3	β_4	γ_0	γ_1	γ_2	γ_3	R^2	BG	W	RR
2	1	127.1	1.082	-1.293	0	0	0.525	0	0.319	0	0.817	0.382	0.071	0.052
2	2	190.6	1.500	-0.998	0	0	0	0	0	0	0.644	0.676	0.009	0.981
2	3	105.2	1.502	-0.925	0	0.292	0	0	0	0	0.776	0.505	0.009	0.136
2	4	136.2	1.593	-1.285	0	0	0	0	0.343	0	0.688	0.319	0.007	0.824
2	5	171.7	0	-0.438	0	0	0.954	0	0	0	0.594	0.384	0.479	0.299
2	6	159.3	1.643	-1.371	1.433	-0.871	0	0	-0.954	0.679	0.732	0.046	0.297	0.491
2	7	167.4	1.184	-0.546	0	0	0	0	0	0	0.747	0.077	0.145	0.920
2	8	72.7	1.321	0	0	0	0	-0.649	0	0.261	0.745	0.139	0.198	0.656
2	9	171.4	0	0	0	0	0.496	0	-0.502	0.346	0.378	0.732	0.504	0.092
2	10	157.1	1.584	-1.397	0.494	0	0	0	0	0	0.705	0.116	0.203	0.995
2	11	169.9	1.187	-0.673	0	0	0	0	0	0	0.665	0.713	0.681	0.924
2	12	195.3	1.378	-0.910	0	0	0	0	0	0	0.666	0.824	0.254	0.399
2	13	169.5	0	0	0	0.644	0.905	-0.415	0	-0.624	0.628	0.129	0.085	0.358
2	14	116.9	1.087	0	0	0	0	-0.314	0	0	0.663	0.948	0.135	0.967
2	15	146.1	0.563	-0.607	0	0	0.621	0	0	0	0.637	0.378	0.013	0.415
2	16	141.3	1.643	-0.962	0	0	0	0	0	0	0.799	0.762	0.124	0.628
2	17	136.2	1.519	-1.100	0	-0.581	0	0	0	0.806	0.650	0.161	0.673	0.981
2	18	130.3	1.672	-1.477	0.477	0	0	0	0	0	0.666	0.661	0.016	0.833
2	19	84.6	1.528	-0.972	0	0	0	0.324	0	0	0.822	0.595	0.371	0.019
2	20	88.3	1.632	-1.376	0.501	0	0	0	0	0	0.607	0.722	0.081	0.230
2	21	141.2	1.328	-0.775	0	0	0	0	0	0	0.679	0.160	0.139	0.076
2	22	119.5	1.382	-0.756	0	-0.829	0	0	0	0.855	0.779	0.044	0.477	0.435
2	23	306.7	0	-0.524	0	0	0.610	0	0	0	0.441	0.869	0.019	0.033
2	24	187.5	0	0	0	0	1.190	-0.701	0	0	0.745	0.536	0.457	0.228
2	25	136.2	1.060	-0.415	0	0	0	0	0	0	0.558	0.761	0.006	0.095
2	26	71.7	1.907	-1.489	0.435	0	0	0	0	0	0.769	0.856	0.091	0.512

Note: ‘S’ indicates the number of the session, ‘P’ indicates the number of the participant. A bold value of α indicates that the constant is significant at the 5% level. The columns ‘BG’, ‘W’ and ‘RR’ report p -values of the Breusch-Godfrey test, White test and Ramsey RESET test. Bold values indicate significance at the 5% level.

S	P	α	β_1	β_2	β_3	β_4	γ_0	γ_1	γ_2	γ_3	R^2	BG	W	RR
3	1	47.1	1.358	0	0	0.823	0	-0.664	0	-0.564	0.940	0.346	0.001	0.824
3	2	46.0	0.542	-0.484	0	0	0.658	0	0	0	0.548	0.280	0.002	0.035
3	3	-11.4	2.836	0	0	0.224	-0.696	-1.138	0	0	0.940	0.280	0.004	0.794
3	4	26.1	2.702	-2.809	1.999	0	0	0	-0.787	0	0.902	0.398	0.001	0.616
3	5	32.3	1.531	0	-1.211	0	0	0	0	0.599	0.911	0.419	0.475	0.103
3	6	29.9	2.486	-2.307	0	0	0	0	0	0.769	0.898	0.556	0.157	0.945
3	7	31.4	2.376	-0.604	0	0	-0.815	0	0	0	0.955	0.772	0.018	0.098
3	8	55.2	2.383	-2.3	0.793	0	0	0	0	0	0.864	0.205	0.000	0.651
3	9	57.0	2.324	0	-3.640	0.675	-0.576	0	1.982	0	0.964	0.155	0.044	0.956
3	10	9.1	1.647	-0.902	0.514	-0.584	0	0	0	0.340	0.982	0.137	0.086	0.660
3	11	42.3	1.889	-1.117	0	1.540	0	0	0	-1.243	0.898	0.101	0.605	0.094
3	12	-16.8	3.340	0	0	0.644	-0.946	-1.041	-0.637	0	0.915	0.338	0.021	0.325
3	13	58.6	0	-0.790	0	0.301	1.284	0	0	0	0.874	0.270	0.050	0.407
3	14	23.1	2.876	-1.571	0.851	0	0	-1.168	0	0	0.880	0.733	0.053	0.222
3	15	48.7	1.310	-2.436	2.117	-1.002	0.773	0	0	0	0.845	0.273	0.027	0.091
3	16	3.5	2.375	-0.896	0	0.388	0	-0.739	0	0	0.906	0.017	0.244	0.255
3	17	-5.0	2.242	0	-1.068	0.959	0	-0.932	0	0	0.877	0.230	0.021	0.185
3	18	-5.4	1.958	-0.385	0	-0.762	-0.480	0	0	0.711	0.984	0.430	0.000	0.231
3	19	28.8	2.756	-0.692	0	0	-1.038	0	0	0	0.963	0.863	0.330	0.768
3	20	54.0	1.008	0	-0.846	0.610	0	0	0	0	0.531	0.930	0.013	0.875
3	21	46.1	1.332	-1.353	0	0	0.436	0.437	0	0	0.949	0.128	0.009	0.723
3	22	56.3	2.734	-1.979	0	0	0	-0.676	0.773	0	0.889	0.544	0.066	0.678
3	23	10.5	1.399	-1.029	0	0.877	1.018	-1.225	0.820	-0.780	0.936	0.018	0.002	0.282
3	24	5.5	1.962	-1.068	0	0.533	0.506	-0.829	0	0	0.901	0.163	0.009	0.026

S	P	α	β_1	β_2	β_3	β_4	γ_0	γ_1	γ_2	γ_3	R^2	BG	W	RR
4	1	123.9	1.602	-0.983	0	0.205	0	0	0	0	0.755	0.362	0.559	0.473
4	2	113.4	1.780	-1.440	0	0	0	0	0.394	0	0.788	0.612	0.558	0.772
4	3	132.3	1.524	-1.208	0.392	0	0	0	0	0	0.657	0.447	0.220	0.306
4	4	98.4	1.307	-0.531	0	0	0	0	0	0	0.797	0.757	0.022	0.036
4	5	146.4	1.201	0	0	0	0	-0.581	0	0	0.652	0.890	0.455	0.257
4	6	183.8	1.075	-0.617	0	0	0	0	0	0	0.516	0.086	0.022	0.799
4	7	139.2	0.739	-0.939	0	0	0.538	0	0	0.304	0.699	0.879	0.821	0.340
4	8	123.7	1.169	-0.609	0	0	0	0	0	0	0.699	0.335	0.001	0.265
4	9	76.9	1.305	-0.516	0	0	0	0	0	0	0.851	0.809	0.152	0.586
4	10	37.1	2.106	0	0	0	-0.879	-0.449	0	0	0.800	0.338	0.077	0.776
4	11	130.0	1.606	-1.503	0.946	0	0	0	0	-0.381	0.694	0.908	0.596	0.824
4	12	186.2	0	0	0	0	0.949	-0.461	0	0	0.549	0.620	0.768	0.151
4	13	110.5	1.494	-0.709	0	0	0	0	0	0	0.774	0.231	0.641	0.922
4	14	195.7	1.330	-0.871	0	0	0	0	0	0	0.572	0.322	0.036	0.854
4	15	171.7	0	0	0	0	1.005	-0.489	0	0	0.589	0.537	0.078	0.446
4	16	153.9	1.410	0	0	0	0	-0.715	0	0	0.548	0.536	0.292	0.864
4	17	189.4	0	0	0	0	0.777	-0.315	0	0	0.417	0.887	0.124	0.302
4	18	156.2	1.378	-0.793	0	0	0	0	0	0	0.653	0.344	0.055	0.846
4	19	162.0	1.022	-0.469	0	0	0	0	0	0	0.485	0.556	0.055	0.765
4	20	144.3	1.855	-1.554	0	0	0	0	0.351	0	0.722	0.193	0.049	0.375
4	21	124.0	1.506	-1.576	0.639	0	0	0	0	0	0.617	0.019	0.648	0.361
4	22	32.6	1.677	0	0.822	0	-0.613	-0.594	-0.484	0	0.899	0.450	0.842	0.733
4	23	92.8	1.205	-0.537	0	0	0	0	0	0	0.699	0.882	0.144	0.118
4	24	101.4	1.733	-1.098	0	0	0	0	0	0.172	0.775	0.919	0.252	0.882
4	25	178.9	1.346	-0.811	0	0	0	0	0	0	0.643	0.004	0.541	0.479
4	26	195.8	0	0	0	0	0.918	-0.447	0	0	0.526	0.733	0.053	0.107
4	27	193.2	0.899	-0.524	0	0	0	0	0	0	0.376	0.088	0.007	0.718
4	28	174.5	1.262	0	0	0	0	-0.642	0	0	0.623	0.294	0.216	0.502
4	29	52.2	1.701	-0.867	0	0	0	0	0.393	-0.276	0.797	0.971	0.418	0.282
4	30	122.5	1.580	-1.337	0.470	0	0	0	0	0	0.730	0.239	0.647	0.902
4	31	202.5	0	-0.485	0	0	0.903	0	0	0	0.591	0.396	0.265	0.256
4	32	95.8	1.063	0	0	0.742	0	-0.509	0	-0.568	0.642	0.048	0.011	0.624

Note: 'S' indicates the number of the session, 'P' indicates the number of the participant. A bold value of α indicates that the constant is significant at the 5% level. The columns 'BG', 'W' and 'RR' report p -values of the Breusch-Godfrey test, White test and Ramsey RESET test. Bold values indicate significance at the 5% level.

S	P	α	β_1	β_2	β_3	β_4	γ_0	γ_1	γ_2	γ_3	R^2	BG	W	RR
5	1	25.5	0.806	-1.357	0	0	0.806	0	0.331	0	0.799	0.106	0.227	0.732
5	2	7.6	2.668	-1.806	0	0	-0.469	0	0.517	0	0.908	0.151	0.065	0.075
5	3	17.1	1.123	-0.865	0	0	0.464	0	0	0	0.881	0.347	0.230	0.565
5	4	4.8	2.333	-2.526	1.471	-0.881	0	0	0	0.524	0.879	0.105	0.240	0.467
5	5	3.1	1.396	-0.480	0	0	0	0	-0.182	0.213	0.958	0.452	0.686	0.287
5	6	20.2	1.994	-2.466	1.813	-0.694	0	0	0	0	0.729	0.400	0.748	0.343
5	7	14.8	1.353	-1.503	1.243	-0.686	0.358	0	0	0	0.883	0.211	0.399	0.354
5	8	10.4	1.935	-2.097	1.400	-0.446	0	0	0	0	0.756	0.224	0.181	0.815
5	9	20.2	1.657	-2.091	1.848	-0.771	0.526	-0.468	0	0	0.825	0.053	0.017	0.137
5	10	3.3	1.993	-1.826	0.751	0	0	0	0	0	0.824	0.170	0.997	0.643
5	11	14.6	1.661	-1.656	1.096	-0.421	0	0	0	0	0.763	0.359	0.215	0.114
5	12	4.0	2.452	-2.719	2.111	-0.688	0	-0.238	0	0	0.928	0.082	0.409	0.109
5	13	21.5	2.238	-2.690	1.417	-0.781	0	0	0.456	0	0.835	0.841	0.914	0.660
5	14	13.1	2.010	-1.982	1.191	-0.465	0	0	0	0	0.860	0.944	0.156	0.290
5	15	9.7	1.926	-1.447	0.352	0	0	0	0	0	0.833	0.960	0.730	0.300
5	16	10.0	1.344	-2.077	1.086	0	0.409	0.358	0	-0.291	0.829	0.459	0.838	0.000
5	17	24.8	2.050	-2.571	1.895	-0.783	0	0	0	0	0.764	0.205	0.652	0.262
5	18	23.8	1.291	-2.337	1.942	-0.865	0.564	0	0	0	0.749	0.080	0.093	0.581
5	19	7.1	1.704	-1.082	0.485	0	0	0	0	-0.235	0.845	0.663	0.106	0.227
5	20	24.0	0	-0.993	0.956	-0.514	1.147	0	0	0	0.765	0.565	0.382	0.319
5	21	13.5	2.492	-2.195	1.522	-0.713	0	-0.977	0.658	0	0.729	0.063	0.047	0.048

S	P	α	β_1	β_2	β_3	β_4	γ_0	γ_1	γ_2	γ_3	R^2	BG	W	RR
6	1	39.9	2.117	-0.776	0	0	-0.356	0	0	0	0.940	0.651	0.269	0.729
6	2	40.8	1.552	-0.704	0	0	0	0	0	0	0.836	0.311	0.008	0.883
6	3	43.0	2.336	-0.899	0	0	-0.491	0	0	0	0.936	0.327	0.394	0.847
6	4	59.3	1.831	0	-0.866	0.573	0	-0.591	0	0	0.822	0.217	0.058	0.970
6	5	14.5	2.129	-1.703	0.610	0	0	0	0	0	0.887	0.076	0.010	0.172
6	6	155.4	1.179	0	-0.492	0	0	-0.479	0	0	0.471	0.069	0.152	0.018
6	7	56.2	2.237	-1.858	0.484	0	0	0	0	0	0.865	0.797	0.043	0.682
6	8	47.4	1.992	-1.280	0	0.221	0	0	0	0	0.894	0.236	0.134	0.960
6	9	36.2	1.931	-0.851	0	0.657	0	0	-0.707	0	0.876	0.732	0.006	0.579
6	10	148.7	1.666	-1.057	0	0	0	0	0	0	0.745	0.144	0.873	0.727
6	11	54.5	2.530	-1.199	0	0.198	-0.644	0	0	0	0.909	0.803	0.344	0.877
6	12	299.1	0.866	0	0	0	0	0	-0.555	0	0.283	0.676	0.548	0.893
6	13	82.6	1.941	-1.350	0	0.218	0	0	0	0	0.855	0.196	0.071	0.982
6	14	85.7	2.665	-1.146	0	0	-0.628	0	0	0	0.899	0.685	0.006	0.210
6	15	26.8	0.858	-0.646	0	0	0.601	0	0	0	0.765	0.083	0.002	0.547
6	16	38.6	2.030	-1.685	1.281	0	0	0	-0.614	0	0.910	0.071	0.994	0.708
6	17	75.8	1.699	-0.857	0	0	0	0	0	0	0.887	0.559	0.416	0.623
6	18	69.7	1.728	-0.844	0	0	0	0	0	0	0.893	0.198	0.023	0.090
6	19	23.5	2.135	-0.450	0	0	-0.686	0	0	0	0.970	0.504	0.368	0.288
6	20	62.9	1.605	-1.207	0	0.431	0	0	0	0	0.612	0.579	0.068	0.772
6	21	105.4	0	0	-0.894	0.510	1.059	0	0	0	0.770	0.545	0.513	0.106

Note: 'S' indicates the number of the session, 'P' indicates the number of the participant. A bold value of α indicates that the constant is significant at the 5% level. The columns 'BG', 'W' and 'RR' report p -values of the Breusch-Godfrey test, White test and Ramsey RESET test. Bold values indicate significance at the 5% level.

S	P	α	β_1	β_2	β_3	β_4	γ_0	γ_1	γ_2	γ_3	R^2	BG	W	RR
7	1	19.3	1.689	-0.748	0	0	0	0	0	0	0.868	0.758	0.395	0.075
7	2	22.1	1.393	-0.585	0	0	0	0	-0.299	0.508	0.828	0.573	0.138	0.877
7	3	152.5	0	0	0	0	1.072	-0.470	0	0	0.641	0.790	0.555	0.702
7	4	82.9	1.681	-1.328	0.452	0	0	0	0	0	0.707	0.841	0.397	0.407
7	5	127.5	0	0	-0.449	0	0.772	0	0	0.327	0.549	0.645	0.067	0.886
7	6	117.8	0.796	0	0	0	0	0	0	0	0.524	0.689	0.425	0.683
7	7	96.1	0.561	-0.633	0	0	0.790	0	0	0	0.768	0.289	0.019	0.199
7	8	128.8	1.463	-1.501	0.633	0	0	0	0	0	0.462	0.074	0.204	0.830
7	9	68.1	2.355	-0.887	0	0	-0.576	0	0	0	0.885	0.505	0.542	0.164
7	10	168.6	0.876	0	0	0	0.506	-0.725	0	0	0.634	0.789	0.857	0.485
7	11	143.9	1.583	-0.944	0	0	0	0	0	0	0.786	0.734	0.412	0.125
7	12	220.9	0	-0.501	0	0	0.909	0	0	0	0.588	0.842	0.452	0.066
7	13	83.6	1.633	-0.775	0	0	0	0	0	0	0.906	0.931	0.442	0.830
7	14	138.7	1.524	0	0	0	0	-0.789	0	0	0.634	0.238	0.528	0.861
7	15	50.1	1.404	0	0	0	-0.206	-0.260	0	0	0.952	0.831	0.019	0.720
7	16	67.4	2.690	0	0	0	-0.885	-1.028	0	0	0.848	0.842	0.551	0.857
7	17	52.8	1.994	0	0	0.365	-0.505	-0.563	-0.275	0	0.845	0.883	0.108	0.091
7	18	139.6	1.420	0	-1.121	0	0	0	0	0.416	0.878	0.191	0.470	0.998
7	19	163.2	1.134	0	-0.813	0	0	-0.418	0	0.731	0.593	0.293	0.937	0.827
7	20	52.3	1.540	0.992	0	0.462	0	-1.357	-0.800	0	0.797	0.336	0.517	0.805
7	21	146.3	1.746	-1.069	0	0	0	0	0	0	0.752	0.731	0.701	0.620
7	22	91.2	1.397	0	0	0	0	-0.629	0	0	0.642	0.334	0.019	0.944
7	23	113.8	1.562	0	0.868	0	0	-0.814	-0.687	0	0.771	0.220	0.999	0.940
7	24	154.9	1.563	-0.874	0	0	0	0	0	0	0.758	0.555	0.501	0.531
7	25	184.2	1.544	-0.999	0	0	0	0	0	0	0.692	0.862	0.798	0.105
7	26	118.7	1.282	-1.086	0	0	0.335	0	0	0.232	0.741	0.571	0.053	0.103
7	27	172.0	0.864	0	0	0	0.632	-0.883	0	0	0.695	0.175	0.832	0.792
7	28	181.0	0.996	-0.734	0	0	0.361	0	0	0	0.452	0.467	0.777	0.306
7	29	163.2	1.752	-1.299	0	0	0	0	0	0	0.651	0.866	0.024	0.005
7	30	96.4	0	0	0	-0.203	0.869	0	0	0	0.752	0.531	0.000	0.168
7	31	30.8	1.464	0	-1.159	0.685	0	0	0	0	0.884	0.500	0.374	0.342
7	32	17.0	1.390	0	0.308	0	0	-0.670	0	0	0.904	0.557	0.139	0.233

Note: 'S' indicates the number of the session, 'P' indicates the number of the participant. A bold value of α indicates that the constant is significant at the 5% level. The columns 'BG', 'W' and 'RR' report p -values of the Breusch-Godfrey test, White test and Ramsey RESET test. Bold values indicate significance at the 5% level.

Table 11: First-order heuristics

S	P	α_1	α_2	$1 - \alpha_1 - \alpha_2$	β	anchor	adjustment
1	3	0.953	0	0.047	0.781	naive	trend-following
1	5	0	0.662	0.338	0.574	mixed	trend-following
1	6	0.897	0	0.103	0.682	naive	trend-following
1	8	1.694	-0.798	0.104	1.043	mixed	trend-following
1	11	0.516	0	0.484	1.025	mixed	trend-following
1	13	0	0.620	0.380	0.799	mixed	trend-following
1	14	0.742	0	0.258	1.050	mixed	trend-following
1	17	0.898	0	0.102	0.750	naive	trend-following
1	24	0	0.645	0.355	0.730	mixed	trend-following
1	25	0.424	0	0.576	1.070	mixed	trend-following
1	26	0.769	0	0.231	1.092	mixed	trend-following
S	P	α_1	α_2	$1 - \alpha_1 - \alpha_2$	β	anchor	adjustment
2	2	0.511	0	0.489	0.998	mixed	trend-following
2	5	0	0.597	0.403	0.496	mixed	trend-following
2	7	0	0.596	0.404	0.746	mixed	trend-following
2	11	0.517	0	0.483	0.673	mixed	trend-following
2	12	0.474	0	0.526	0.91	mixed	trend-following
2	15	0	0.591	0.409	0.598	mixed	trend-following
2	16	0.692	0	0.308	0.962	mixed	trend-following
2	21	0.551	0	0.449	0.776	mixed	trend-following
2	23	0	0	1	1.041	fundamental	trend-following
2	25	0.651	0	0.349	0.415	mixed	trend-following
S	P	α_1	α_2	$1 - \alpha_1 - \alpha_2$	β	anchor	adjustment
3	2	0	0.723	0.277	0.505	mixed	trend-following
3	7	1.730	-0.778	0.048	0.627	mixed	trend-following
3	19	1.781	-0.794	0.014	0.761	mixed	trend-following
S	P	α_1	α_2	$1 - \alpha_1 - \alpha_2$	β	anchor	adjustment
4	4	0.785	0	0.215	0.531	mixed	trend-following
4	6	0	0.473	0.527	0.589	mixed	trend-following
4	8	0	0.529	0.471	0.686	mixed	trend-following
4	9	0.793	0	0.207	0.516	mixed	trend-following
4	13	0.798	0	0.202	0.710	naive	trend-following
4	14	0	0.400	0.600	0.810	mixed	trend-following
4	18	0.592	0	0.408	0.794	mixed	trend-following
4	19	0.558	0	0.442	0.470	mixed	trend-following
4	23	0.663	0	0.337	0.536	mixed	trend-following
4	25	0	0.489	0.511	0.785	mixed	trend-following
4	27	0	0.433	0.567	0.433	mixed	trend-following
4	31	0	0.500	0.500	0.561	mixed	trend-following
S	P	α_1	α_2	$1 - \alpha_1 - \alpha_2$	β	anchor	adjustment
5	3	0	0.668	0.332	0.949	mixed	trend-following

Note: 'S' indicates the number of the session, 'P' indicates the number of the participant.

S	P	α_1	α_2	$1 - \alpha_1 - \alpha_2$	β	anchor	adjustment
6	1	0.912	0	0.088	0.892	naive	trend-following
6	2	0.846	0	0.154	0.703	mixed	trend-following
6	3	0.853	0	0.147	0.980	mixed	trend-following
6	10	0	0.584	0.416	0.874	mixed	trend-following
6	14	0.781	0	0.219	1.122	mixed	trend-following
6	15	0	0.774	0.226	0.743	mixed	trend-following
6	17	0.848	0	0.152	0.863	mixed	trend-following
6	18	0.891	0	0.109	0.852	naive	trend-following
6	19	1.536	-0.553	0.017	0.511	mixed	trend-following

S	P	α_1	α_2	$1 - \alpha_1 - \alpha_2$	β	anchor	adjustment
7	1	0.941	0	0.059	0.748	naive	trend-following
7	6	0.811	0	0.189	0	naive	none
7	7	0	0.741	0.259	0.611	mixed	trend-following
7	9	1.374	-0.493	0.118	0.899	mixed	trend-following
7	11	0.645	0	0.355	0.945	mixed	trend-following
7	12	0	0.533	0.467	0.739	mixed	trend-following
7	13	0.868	0	0.132	0.778	mixed	trend-following
7	21	0.688	0	0.312	1.072	mixed	trend-following
7	24	0.703	0	0.297	0.878	mixed	trend-following
7	25	0.554	0	0.446	1.002	mixed	trend-following
7	28	0	0.512	0.488	0.780	mixed	trend-following
7	29	0.446	0	0.554	1.297	mixed	trend-following

Note: ‘S’ indicates the number of the session, ‘P’ indicates the number of the participant.