



UvA-DARE (Digital Academic Repository)

Learning to forecast: Genetic algorithms and experiments

Makarewicz, T.A.

[Link to publication](#)

Citation for published version (APA):

Makarewicz, T. A. (2014). *Learning to forecast: Genetic algorithms and experiments*. Amsterdam: Tinbergen Institute.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 4

Bubble Formation and (In)Efficient Markets in Learning-to-Forecast and -Optimize Experiments

4.1 Introduction

This chapter investigates the price dynamics and bubble formation in an experimental asset pricing market with a price adjustment rule. The purpose of the study is to address a fundamental question about the origins of bubbles: do bubbles arise because agents fail to form rational expectations or because they fail to optimise their trading quantity given their expectations?

We design three experimental treatments: (1) subjects make a forecast only, and are paid according to forecasting accuracy; (2) subjects make a quantity decision only, and are paid according to the profitability of their decision; (3) subjects make both a forecast and a quantity decision, and are paid by their performance of either of the tasks with equal probability. Under perfect rationality and perfect competition, these three tasks are equivalent and should lead the subjects to an equilibrium with a constant fundamental price. In contrast, we find none of the experimental markets to show a reliable convergence to the fundamental outcome. The market price is relatively most stable, with an upward trend in the treatment where the subjects make a forecast only. There are recurring bubbles and crashes with high frequency and magnitude when the subjects submit both a price forecast and a trading quantity decision.

Asset bubbles can be traced back to the very beginning of financial market, but has not been investigated extensively by modern economics and finance literature. One possible reason is that it contradicts the standard theory of rational expectations

(Lucas Jr., 1972; Muth, 1961) and efficient markets (Fama, 1970). Recent finance literature however has shown growing interest in bounded rationality (Farmer and Lo, 1999; Shiller, 2003) and ‘abnormal’ market movement such as over- and under-reaction to changes in fundamentals (Bondt and Thaler, 2012) and excess volatility (Campbell and Shiller, 1989). The recent financial crisis and precedent boom and bust in the US housing market highlight the importance of understanding the mechanism of financial bubbles in order for the policies makers to design policies/institutions to enhance market stability.

It is usually difficult to identify bubbles using data from the field, since people may substantially disagree about the underlying fundamental price of the asset. Laboratory experiments have an advantage in investigating this question by taking full control over the underlying fundamental price. Smith et al. (1988) are among the first authors to reliably reproduce price bubbles and crashes of asset prices in a laboratory setting. They let the subjects trade an asset that pays a dividend in each of 15 periods. Therefore the fundamental price at each period equals the sum of the remaining expected dividends and follows a decreasing step function. The authors find the price to go substantially above the fundamental price after the initial periods before it crashes towards the end of the experiment. This approach has been followed in many studies *i.e.* Dufwenberg et al. (2005); Haruvy and Noussair (2006); Lei et al. (2001); Noussair et al. (2001).¹ A typical result of these papers is that the price boom and bust is a robust finding despite several major changes in the experimental environment.

Nevertheless, Huber and Kirchler (2012); Kirchler et al. (2012) argue that the non-fundamental outcomes in this type of experiments are due to misunderstanding: subjects may be simply confused by the declining fundamental price. They support their argument by showing that no bubble appears when the fundamental price is not declining or when the declining fundamental price is further illustrated by an example of ‘a depletable gold mine’. Another potential concern is that in these experiments, due to typically short horizon (15 periods), one cannot test whether financial crashes are likely to be followed by new bubbles. It is very important to study consequent boom-bust cycles in asset prices, for example to understand the evolution of the asset prices between the dot-com and the 2007 crises.

The Smith et al. (1988) experiment are categorised as ‘learning to optimise’ (henceforth LtO) experiments (see Duffy, 2008, for an extensive discussion). Besides this approach, there is ‘learning to forecast’ (henceforth LtF) experimental design introduced by Marimon et al. (1993) (see Hommes, 2011, for a comprehensive survey). Hommes

¹For survey of the literature, see Noussair and Tucker (2013); Sunder (1995).

et al. (2005) run an experiment where subjects act as professional advisers (forecasters) for a pension fund: they submit a price forecast, which is transformed into a quantity decision of buying/selling by a computer program based on optimization over a standard myopic mean-variance utility function. Subjects are paid according to their forecasting accuracy. The fundamental price is defined as the rational expectation equilibrium and remains constant over time. The result of this chapter is twofold: (1) the asset price fails to converge to the fundamental, but oscillates and forms bubbles in several markets; (2) instead of having rational expectations, most subjects coordinate on a price trend following strategies (*cf.* Bostian and Holt, 2009). Heemeijer et al. (2009) and Bao et al. (2012) investigate whether the non-convergence result is driven by the positive expectation feedback nature of the experimental market in Hommes et al. (2005). Positive/negative expectation feedback means that the realised market price increases/decreases when the average price expectation increases/decreases. The results show that while negative feedback markets converge quickly to the fundamental price, and adjust quickly to a new fundamental after a large shock, positive feedback markets usually fail to converge, but under-react to the shocks in the short run, and over-react in the long run.

The subjects in Hommes et al. (2005) and other ‘learning to forecast’ experiments do not directly trade, but are assisted by a computer program to translate their forecasts into optimal trading decisions. A natural question is what will happen if they submit explicit quantity decisions, i.e. if the experiment is based on the ‘learning to optimise’ design. Are the observed bubbles robust against the LtO design or are they just an artifact of the computerised trading in the LtF design?

In this chapter we design an experiment, in which the fundamental price is constant over time (as in Hommes et al., 2005), but the subjects are asked to directly indicate the amount of asset they want to buy/sell. Different from the double auction mechanism in the Smith et al. (1988) design, the price in our experiment is determined by a price adjustment rule based on excess supply/demand (Beja and Goldman, 1980; Campbell et al., 1997; LeBaron, 2006). Our experiment is helpful in testing financial theory based on such demand/supply market mechanisms. Furthermore, our design allows us to have a longer time span of the experimental sessions, which will enable a test for the recurrence of bubbles and crashes.

The main finding of our experiment is that the persistent deviation from the fundamental price in Hommes et al. (2005) is a robust finding against task design. Based on Relative Absolute Deviation (RAD) and Relative Deviation (RD) as defined by Stöckl et al. (2010), we find that the amplitude of the bubbles in treatment (2) and (3) is much higher than in treatment (1). We also find large heterogeneity in traded quan-

tities than individual price forecasts. These findings suggest that learning to optimise is even harder than learning to forecast, and therefore leads to even larger deviations from rationality and efficiency.

An important finding of our experiment is that in the mixed, LtO and LtF designs we find some repeated ‘*super bubbles*’, where the price increases to more than 3 times the fundamental price. This was not observed in the former experimental literature. Considering that bubbles in stock and housing prices reached similar levels (the housing price index increases by 300% in several local markets before it decreased by about 50% during the crisis), our experimental design may provide a potentially better test bed for policies that deal with large bubbles.

Another contribution is that, to our best knowledge, we are the first to perform a formal statistical test on individual heterogeneity in forecasting and trading strategies in an asset pricing experiment. In particular, in some trading markets we observe a large degree of heterogeneity in the quantity decision even when the price is rather stable. By examining treatment (3), we also find that many subjects fail to trade at the conditionally optimal quantity given their own forecast.

This chapter is related to Bao et al. (2013) who run an experiment to compare the LtF, LtO and Mixed designs in a cobweb economy. The main difference is that they consider a negative expectation feedback system, for which all markets converge to the RE fundamental price. This chapter is also related to a study by Haruvy et al. (2013) who follow the basic design of Smith et al. (1988), with an additional new issue or repurchase of stocks in order to increase/decrease the supply of stock shares on the market. Theoretically, since the fundamental price in this type of studies is based purely on the dividend process, and irrespective of the size of the share supply, the new issue and repurchase should generate no impact on the asset price. But the results suggest that the price level is actually negative related to the supply of asset. This outcome points in the same direction as the intuition behind the models based on excess supply/demand, which we used in our experiments. The difference is that we keep the asset supply constant in our experiment, and the price change is driven instead by the asset excess demand of the investors (played by subjects).

The chapter is organised as follows: Section 2 presents the experimental design, Section 3 states the research hypothesis of the experiment, Section 4 reports the experimental result, and finally, Section 5 concludes.

4.2 Experimental design

4.2.1 Experimental economy

The experiment is based on an asset market in Brock and Hommes (1998) where the price is determined by excess demand/supply. The agents are assumed to have a simple myopic mean-variance objective function. There are $I = 6$ agents, who allocate investment between a risky asset and a risk-free bond. Each agent i at time t has an objective function that is increasing in his wealth in the next period $W_{i,t+1}$, but decreasing with the perceived investment risk $V_{i,t}(W_{i,t+1})$. The wealth of agent i evolves according to

$$(4.1) \quad W_{i,t+1} = RW_{i,t} + z_{i,t}(p_{t+1} + y_{t+1} - Rp_t),$$

where $R = 1 + r$ is the gross interest rate of the risk-free bond (assumed constant over time), $z_{i,t}$ is the demand of risky asset by agent i in period t (positive sign for buying and negative sign for selling). p_t and p_{t+1} are the prices of the risky asset in periods t and $t + 1$ respectively, and y_{t+1} is the assets dividend paid at the beginning of period $t + 1$.

The agent solves the myopic optimisation problem:

$$(4.2) \quad \text{Max}_{z_{i,t}} E_{i,t} W_{i,t+1} - \frac{a}{2} E_{i,t} V(W_{i,t+1}),$$

where a is a parameter for risk aversion. $E_{i,t} V(W_{i,t+1})$ is the conditional expectation by the agent on the conditional variance of the wealth based on publicly available information. The conditional variance equals the $z_{i,t}^2$ times the conditional variance of the excess return per share, ρ_{t+1} defined by

$$(4.3) \quad \rho_{t+1} \equiv p_{t+1} + \bar{y} - Rp_t.$$

We assume that agents have homogeneous and constant expectations on the excess return, i.e. $E_{i,t} V(\rho_{t+1}) \equiv \sigma^2$.² This leads to $E_{i,t} V(W_{i,t+1}) = \sigma^2 z_{i,t}^2$. Let $p_{i,t+1}^e = E_{i,t} p_{t+1}$. For simplicity, we assume that the dividend follows an i.i.d. stochastic process,

²This assumption is made for analytical tractability. Brock and Hommes (1998) noted that the heterogeneity in the price expectations in the expected return of the objective function can lead to heterogeneity in the variance as well, but they ignore it as it is a second order effect. Nelson (1992) provides some justification that in such a model there is typically more disagreement on the mean than the variance of the return. Moreover, this implication assumption is particularly useful in a laboratory experiments, where it would be almost impossible for the subjects to solve the problem if heterogeneity in expected variance is introduced.

where the unconditional expected value $E(y_t) = \bar{y}$. The objective function of the agent can be rewritten as:

$$(4.4) \quad U_{i,t}(z_{i,t}) = RW_{i,t} + (p_{i,t+1}^e + \bar{y} - Rp_t)z_{i,t} - \frac{a\sigma^2 z_{i,t}^2}{2}$$

where agents need to choose $z_{i,t}$ optimally based on their prediction on p_{t+1} . The optimal solution of this quadratic function is shown by

$$(4.5) \quad z_{i,t}^* = \frac{\rho_{t+1}^e}{a\sigma^2} = \frac{p_{i,t+1}^e + \bar{y} - Rp_t}{a\sigma^2},$$

where ρ_{t+1}^e is the conditional expectation on the excess return in the next period.

The market price is set by a market maker using a simple price adjustment mechanism (Beja and Goldman, 1980),³ given by

$$(4.6) \quad p_{t+1} = p_t + \lambda (Z_t^D - Z_t^S) + \varepsilon_t,$$

where $\varepsilon_t \sim N(0, 1)$ is a small i.i.d. idiosyncratic shock, $\lambda > 0$ is a scaling factor, Z_t^S is the exogenous supply and Z_t^D is the total demand. This mechanism guarantees that excess demand/supply increases/decreases the price.

For simplicity, the exogenous supply Z_t^S is normalised to 0 in all periods. We take $R\lambda = 1$, specifically $R = 1 + r = 21/20$, $\lambda = 20/21$, $a\sigma_z^2 = 6$, and $\bar{y} = 3.3$. The price adjustment based on aggregate *individual demand* thus takes the form of

$$(4.7) \quad p_{t+1} = p_t + \frac{20}{21} \sum_{i=1}^6 z_{i,t} + \varepsilon_t.$$

For an optimising agent and the chosen parameters, the individual optimal demand (4.5) equals

$$(4.8) \quad z_{i,t}^* = \frac{\rho_{t+1}^e}{a\sigma^2} = \frac{p_{i,t+1}^e + 3.3 - 1.05p_t}{6},$$

Substituting it back into (4.7) gives

$$(4.9) \quad p_{t+1} = 66 + \frac{20}{21} (\bar{p}_{t+1}^e - 66) + \varepsilon_t,$$

where $\bar{p}_{t+1}^e = \frac{1}{6} \sum_{i=1}^6 p_{i,t+1}^e$ is the average prediction of price p_{t+1} .⁴ This price is the

³See *e.g.* Chiarella et al. (2009) for a survey on the abundant literature about the price adjustment market mechanisms.

⁴Heemeijer et al. (2009) used a very similar price adjustment rule, but the fundamental price is

temporary equilibrium with point-beliefs of prices. Equation (4.9) represents the price adjustment process as a function of the average *individual forecast*.

By plugging in the rational expectations condition, namely $\bar{p}_{t+1}^e = p^f = E(p_{t+1}) = E\left(\frac{20}{21}(\bar{p}_{t+1}^e - 66) + \varepsilon_t\right)$, $p^f = 66$ is the unique Rational Expectations Equilibrium (REE) of the system. It is also true that $p^f = \bar{y}/r$, namely, the fundamental price is equal to the discounted sum of dividend in infinite horizon. If all the agents have rational expectations, the realised price becomes $p_t = p^f + \varepsilon_t = 66 + \varepsilon_t$, *i.e.* the fundamental price plus a white noise, and, on average, the price forecasts are self-fulfilling.

4.2.2 Experimental treatments

Based on the nature of the task and the payoff structure, three treatments are set up:

LtF Classical Learning-to-Forecast experiment. Subjects are asked for one-period ahead price predictions $p_{i,t+1}^e$, based on which the realised price is generated according to the price adjustment rule (4.9). The subjects' reward depends only on the prediction accuracy, defined by (see also Table 4.2 in Appendix 4.B)

$$(4.10) \quad \text{Payoff}_{i,t} = \max \left\{ 0, \left(1300 - \frac{1300}{49} (p_{i,t+1}^e - p_{t+1})^2 \right) \right\}.$$

The law of motion of the treatment economy is given by (4.9).

LtO Classical Learning-to-Optimise experiment, where the subjects are asked to decide on the asset quantity $z_{i,t}$. They are not explicitly asked for a price prediction, but can use a built-in calculator in the experimental program to compute the expected asset return ρ_{t+1} for each price forecast p_{t+1}^e as in equation (4.3). Subjects are rewarded based on a linear transformation of the realised (profit) utility given by

$$(4.11) \quad U_{i,t} = \max \left\{ 0, 800 + 40(z_{i,t}(p_{t+1} + 3.3 - 1.05p_t) - 3z_{i,t}^2) \right\},$$

that is on how close their choice was to the optimal choice regardless of their individual prediction. The law of motion of the LtO treatment is given by (4.7).

We add a constant 800 in order to avoid negative payoff.

Mixed Each subject is asked first for his or her price forecast and second for the choice of the asset quantity. In order to avoid hedging, the reward for the whole

60 in a learning to forecast experiment that compares positive versus negative expectation feedback systems.

experiment is based on either Equation (4.10) or Equation (4.11) with equal probability. The law of motion of the treatment economy is given by (4.7), the same as in LtO and does not depend on the submitted price forecasts.

The same payoff scheme for the forecasting task in the LtF and mixed treatments is the same as in the previous LtF experiments. The points in each treatment are exchanged into Euro with the conversion rate 3000 points = 1 Euro.

Variable	Notation	Parameter
<i>Market parametrization</i>		
Subjects	I	6
Risk penalty	$a\sigma_z^2$	6
Expected value of dividend	y	3.3
Interest rate	r	0.05
Exogenous supply	Z^S	0
Price adjustment parameter	λ	$\frac{20}{21}$
<i>Stationary RE equilibrium</i>		
Price	p^f	66
Excess demand	$z^*(p^f)$	0
Points per 1 Euro		3000

Table 4.1: Parametrization of the experiment.

Finally, we would like to emphasise that the LtF and LtO treatments are equivalent under the assumption of perfect rationality and perfect competition, because the models of the economy in these two treatments, Equations (4.6) and (4.9) are equivalent.

4.2.3 Liquidity constraints

To limit the effect of extreme price forecasts or quantity decisions in the experiment, we impose the following liquidity constraints on the subjects. For the LtF treatment, price predictions such that $p_{i,t+1}^e > p_t + 30$ or $p_{i,t+1}^e < p_t - 30$ are treated as $p_{i,t+1}^e = p_t + 30$ and $p_{i,t+1}^e = p_t - 30$ respectively. For the LtO treatment, quantity decisions greater than 5 or smaller than -5 are treated as 5 and -5 respectively. These two liquidity constraints are roughly the same, since the optimal asset demand (4.8) is close to one sixth of the expected price difference. Nevertheless, the liquidity constraint in the LtF treatment was never binding, while under the LtO treatment subjects would sometimes trade at the edges of the allowed quantity interval.⁵

⁵We also imposed additional constraint that p_t has to be non-negative and not greater than 300. In the experiment, this constraint never had to be implemented.

4.2.4 Number of observations

Experimental instructions with the computer screen presented to the subjects are shown in Appendix 4.A. The experiment took place on December 14, 17, 18 and 20, 2012 and June 6, 2014 at the CREED Laboratory, University of Amsterdam. 144 subjects were recruited. The experiment employs a group design with 6 subjects in each experimental market. There are 24 markets in total and 8 for each treatment. No subject participates in more than one session. The duration of the experiment is typically about 1 hour for the LtF treatment, 1 hour and 15 minutes for the LtO treatment, and 1 hour 45 minutes for the Mixed treatment.

4.3 Testable Hypotheses

The RE benchmark suggests that the subjects should learn to play the rational expectations equilibrium and behave similarly in all treatments. In addition, a rational decision maker should be able to solve the optimal demand for the asset given his price forecast according to Equation (4.8) in the Mixed treatment. These theoretical predictions can be formulated to the following testable hypotheses:

Hypothesis 1: The asset prices converge to the rational expectation equilibrium in all markets;

Hypothesis 2: There is no systematic difference between the market prices across the treatments;

Hypothesis 3: Subjects' earnings efficiency (defined as the ratio of the experimental payoff divided by the hypothetical payoff when all subjects play the REE) are independent from the treatment;

Hypothesis 4: The quantity decision by the subjects are conditionally optimal to their price expectations in the Mixed treatment;

Hypothesis 5: There is no systematic difference between the decision rules used by the subjects for the same task across the treatments;

These hypotheses are further translated into rigorous statistical tests. More specifically, the distribution of Relative (Absolute) Deviation (Stöckl et al., 2010) measures price convergence and differences between the treatments (**Hypothesis 1** and **2**). Relative earnings can be compared with the Mann-Whitney-Wilcoxon rank-sum test (**Hypothesis 3**). Finally, we estimate individual behavioral rules for every subject: a simple restriction test will reveal whether **Hypothesis 4** is true, while the rank-sum test can again be used to test the rule differences between the treatments

(**Hypothesis 5**). Notice that **Hypothesis 1** is nested within **Hypothesis 2**, while **Hypothesis 4** is nested within **Hypothesis 5**.

4.4 Experimental results

4.4.1 Overview

The market prices in each treatment are shown in Figure 4.1 (LtF treatment), Figure 4.2 (LtO treatment) and Figure 4.3 (Mixed treatment). For most of the groups, the prices and predictions remained in the interval $[0, 100]$. The exceptions are the mixed treatment groups 1, 4 and 8 (Figures 4.3a, 4.3d and 4.3h). In the first two of these three groups, prices peaked at almost 150 (more than twice the fundamental price $p^f = 66$) and for the last group, the prices reached 225, almost 3.5 times the fundamental price.

As shown by the figures, the market price is the most stable in the LtF treatment, and the most unstable in the Mixed treatment. In the LtF treatment, there is little heterogeneity in the individual forecasts shown by the green dashed lines. In the LtO treatment, however, there is a high level of heterogeneity in the quantity decisions shown by the blue dashed lines. In the Mixed treatment, it is somewhat surprising that the low heterogeneity in price forecasts and the high heterogeneity in quantity decisions coexist.

The dynamics are diversified between the groups and treatments. In addition, convergence to the REE does not seem to occur in any of the treatments. This suggests the hypotheses based on the rational expectations benchmark are likely to be rejected. In the remainder of this section, we will discuss the statistical evidence in detail.

4.4.2 Quantifying the bubbles

We follow Stöckl et al. (2010) to evaluate the size of mispricing and the experimental asset bubbles, using the Relative Absolute Deviation (RAD) and Relative Deviation (RD). These two indices measure respectively the absolute and relative deviation from the fundamental in a specific period t and are given by

$$(4.12) \quad RAD_{g,t} \equiv \frac{|p_t^g - p^f|}{p^f} \times 100\%,$$

$$(4.13) \quad RD_{g,t} \equiv \frac{p_t^g - p^f}{p^f} \times 100\%,$$




Figure 4.1: Groups 1-8 for the Learning to Forecast treatment. Straight line shows the fundamental price $p^f = 66$, solid black line denotes the realised price, while green dashed lines denote individual forecasts.




Figure 4.2: Groups 1-8 for the Learning to Optimize treatment. Each group is presented in two panels. The upper panel displays the fundamental price $p^f = 66$ (straight line) and the realised prices (solid black line), while the lower panel displays individual trades (dashed blue lines) and average trade (solid red line). Notice the different y -axis scale for group 7 (picture g).

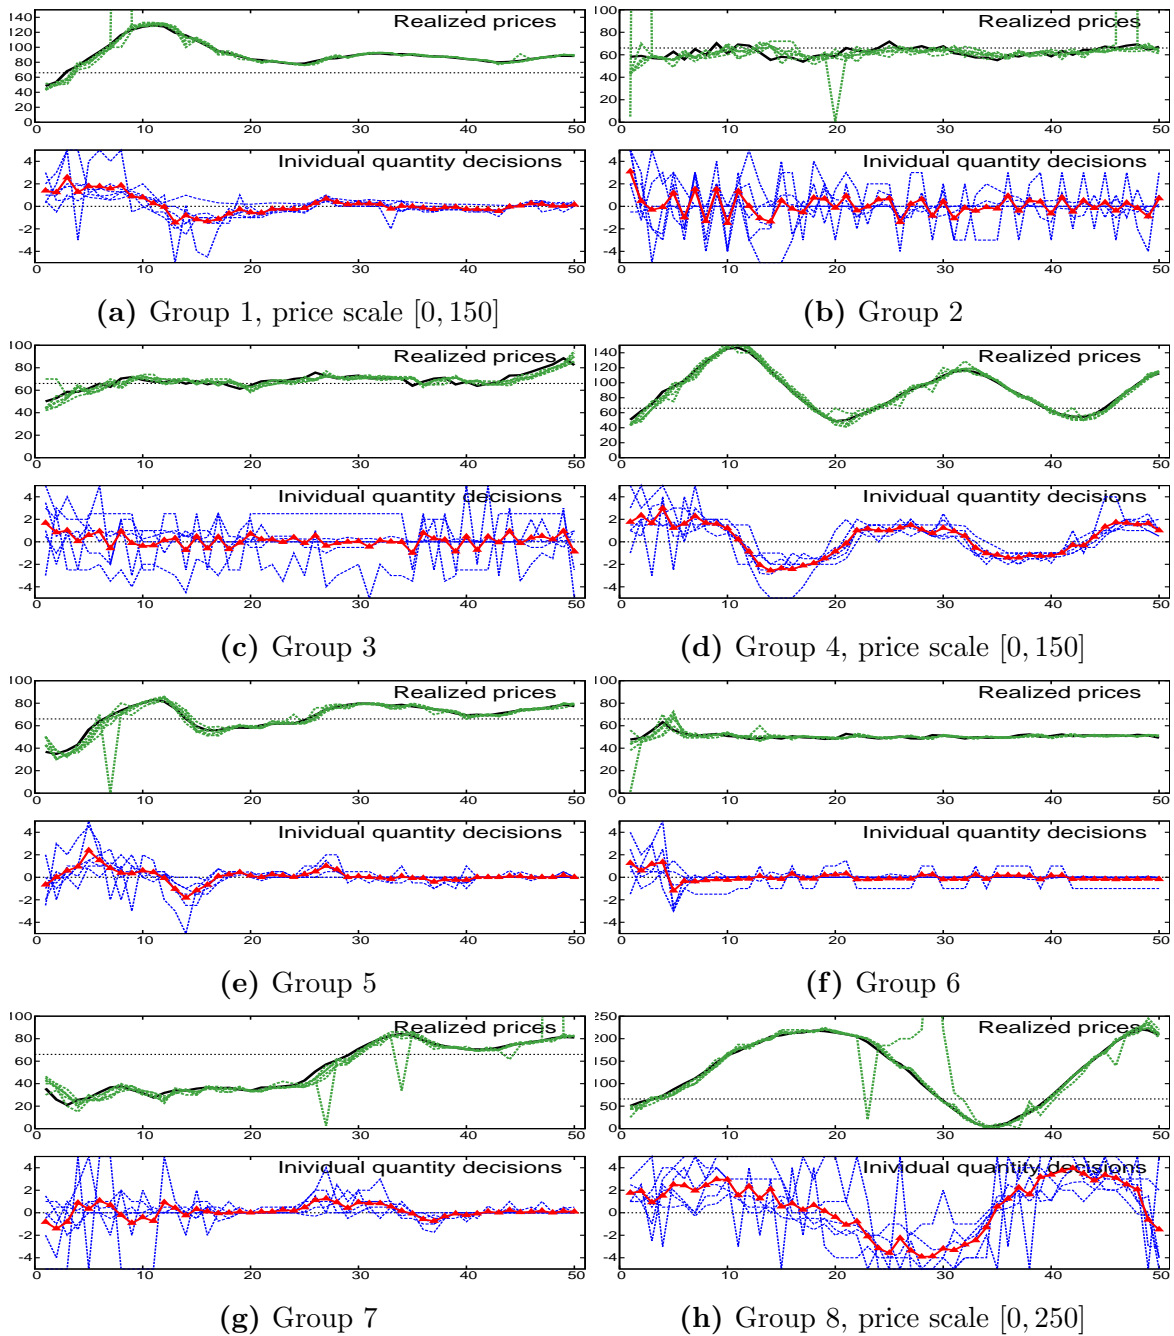


Figure 4.3: Groups 1-8 for the Mixed treatment with subject forecasting and trading. Each group is presented in a picture with two panels. The upper panel displays the fundamental price $p^f = 66$ (straight line), the realized prices (solid black line) and individual predictions (green dashed lines), while the lower panel displays individual trades (dashed blue lines) and average trade (solid red line). Notice the different y -axis scale for groups 1, 4 and 8 (pictures a, d and h respectively).

where $p^f = 66$ is the fundamental price and p_t^g is the realised asset price at period t in the session of group g . The average RAD is defined as

$$(4.14) \quad \overline{RAD}_g = \frac{1}{50} \sum_{t=1}^{50} RAD_{g,t},$$

and it shows the average relative distance between the realised prices and the fundamental, while the average \overline{RD}_g (defined similarly as in (4.14)) focuses more on the sign of this relationship. Groups with average \overline{RD} close to zero could either converge to the fundamental (in which case the $RAD_{g,t}$ is also close to zero) or oscillate around the fundamental (with high $RAD_{g,t}$), while positive or negative average \overline{RD} signals that the group typically over- or underpriced the asset.

The results for average \overline{RAD} and \overline{RD} measures are presented in Table 4.4 in the appendix. They confirm that the LtF groups were the closest to, though still quite far from, the REE (with an average \overline{RAD} of about 9.5%), while Mixed groups exhibited largest bubbles with an average \overline{RAD} of 36%. Interestingly, LtO groups had significant oscillations (on average high \overline{RAD} of 24.6%), but centered close to the fundamental price (average \overline{RD} of 1.4%, compared to average \overline{RD} of -3% and 16.1% for the LtF and Mixed treatments respectively). LtF groups on average are below the fundamental price and Mixed groups typically overshoot it.

A simple t-test shows that for the LtO and Mixed treatment, as well as for 6 out of 8 LtF groups (exceptions are Markets 7 and 8), the means of the groups' RAD measures (disregarding the initial 10 periods to allow for learning) are significantly larger than 3%.⁶ Furthermore, for all groups in all three treatments, t-test on any meaningful significance level rejects null of the average price (for periods 11 – 50, we use the last 40 periods to allow for learning by the subjects in the first 10 periods) being equal to the fundamental value. This result shows negative evidence towards the **Hypothesis 1**: none of the treatments converges to the REE.

There is no significant difference between the treatments in terms of \overline{RD} according to Mann-Whitney-Wilcoxon test (henceforth MWWT; p -value > 0.1 for each pair of the treatments). However, the difference between the LtF treatment and each of the other treatments in terms of \overline{RAD} is significant at 5% according to MWWT (p -value = 0.002 and 0.003 respectively), while the difference between the LtO and Mixed is not significant (p -value = 0.753). This is strong evidence against **Hypothesis 2**, as it shows that trading and forecasting tasks yield different market dynamics.

⁶3% \overline{RAD} is equivalent to a typical price deviation of 2 in absolute terms, which corresponds to twice the standard deviations of the idiosyncratic supply shocks, *i.e.* 95% bounds of the REE.

Our results are comparable with the result in Stöckl et al. (2010) (see specifically their Table 4 for the $\overline{RAD}/\overline{RD}$ measures) in terms of the typical RAD values. Nevertheless, there are some important differences. First, group 8 from the mixed treatment (with \overline{RAD} equal to 120.7%) exhibits the largest relative price bubble among the experimental data. Second, the four experiments investigated by Stöckl et al. (2010) have shorter spans (with sessions of either 10 or 25 periods) and so typically witness one bubble. Our data shows that the mispricing in experimental asset markets is a robust finding. The crash of a bubble does not enforce the subjects to converge to the fundamental, but instead marks the beginning of a ‘crisis’ until the market turns around and a new bubble emerges. This succession of over- and under-pricing of the asset is reflected in our \overline{RD} measures, which are smaller than the typical ones reported by Stöckl et al. (2010), and can even be negative, despite high \overline{RAD} .

Result 1. *Among the three treatments, LtF incurs dynamics closest to the REE. Nevertheless, the average price is still far from the rational expectations equilibrium. Furthermore, in terms of aggregate dynamics LtF treatment is significantly different from the other two treatments, which are indistinguishable between themselves. We conclude that **Hypothesis 1** and **2** are rejected.*

4.4.3 Earnings efficiency

Subjects’ earnings in the experiment are compared to the hypothetical case where all subjects play according to the REE in all 50 periods. Subjects can earn 1300 points per period for the forecasting task when they play according to REE because they make no prediction errors, and 800 points for the trading task when they play according to the REE because the asset return is 0 and they should not buy or sell. We use the ratio of actual to hypothetical REE payoffs as a measure of payoff efficiency. This measure can be larger than 100% in treatments with the LtO and Mixed Treatments, because the subjects can profit if they buy and the price increases and vice versa. These earnings efficiency ratios, as reported in Table 4.5 in the appendix, are generally high (more than 75%).

The earnings efficiency for the forecasting task is higher in the LtF treatment than in the Mixed treatment (difference is significant at 5% level according to MWWT, p -value=0.001). At the same time, the earnings efficiency for the trading task is very similar in the LtO treatment and the Mixed treatment (difference is not significant at 5% level according to MWWT test, p -value=0.753).

Result 2. *Forecasting efficiency is significantly higher in the LtF than the Mixed treatment, while there is no significant difference in the trading efficiency in treatments LtO*

and Mixed. **Hypothesis 3** is partially rejected.

4.4.4 Conditional optimality of forecast and quantity decision in mixed treatment

In the Mixed treatment, each subject makes both a price forecast and a quantity decision. It is therefore possible to investigate whether these two are consistent, namely, whether the subjects' quantity choices are close to the optimal demand conditional on the price forecast as in Equation (4.8) ($1/6$ of the corresponding expected asset return). Figure 4.4 shows the scatter plot of the quantity decision against the implied predicted return (4.3), which we constructed based on the price predictions of each subject, for each period separately.⁷ If all individuals made consistent decisions, these points should lie on the (blue) line with slope $1/6$.




Figure 4.4: ML estimation for trading rule (4.15) in the Mixed treatment. Panel (a) is the scatter plot of the traded quantity (vertical axis) against the implied expected return (horizontal axis). Each point represents one decision of one subject in one period from one group. Panel (b) is the scatter plot of a trading rule (4.15) slope (reaction to expected return; horizontal axis) against constant (trading bias; vertical axis). Each point represents one subject from one group. Solid line (left panel)/triangle (right panel) denotes the optimal trade rule ($z_{i,t} = 1/6\rho_{i,t}^e$). Dashed line (left panel)/circle (right panel) denotes the estimated rule under restriction of homogeneity ($z_{i,t} = c + \phi\rho_{i,t}^e$).

Figure 4.4 brings two interesting observations. First, subjects have some degree of 'digit preference', in the sense that the trading quantities are typically round numbers or contain only one digit after the decimal. Second, the quantity choices are far from being consistent with the price expectations. In fact, the subjects sometimes sold

⁷Sometimes the subjects submit extremely high price predictions, which in many cases seem to be typos. We exclude these outliers, where the predicted returns on the asset greater than 60 in absolute terms.

(bought) the asset even though they believed its return will be substantially positive (negative).

To further evaluate this finding, we run a series of Maximum Likelihood (ML) regressions based on

$$(4.15) \quad z_{i,t} = c_i + \phi_i \rho_{i,t+1}^e + \eta_{i,t},$$

with $\eta_{i,t} \sim NID(0, \sigma_{\eta,i}^2)$. This model has a straightforward interpretation: it takes the quantity choice of subject i in period t as a linear function of the implied (by the price forecast) return on the asset. It has two important special cases: homogeneity and optimality (nested in homogeneity). To be specific, subject homogeneity (heterogeneity) corresponds to an insignificant (significant) variation in the slope $\phi_i = \phi_j$ ($\phi_i \neq \phi_j$) for any (some) two subjects i and j . Optimality of individual quantity decisions implies homogeneity with an additional restriction that $\phi_i = \phi_j = 1/6$. The constant c_i shows subject's i 'irrational' optimism/pessimism bias. Optimality thus corresponds to homogeneity with an additional condition $c_i = c_j = 0$ (no agent has a decision bias).

The assumptions of homogeneity and perfect optimisation are tested by estimation of equation (4.15) with the restrictions on the parameters c_i and ϕ_i .⁸ These regressions are compared with an unrestricted regression (with $\phi_i \neq \phi_j$ and $c_i \neq c_j$) via a Likelihood Ratio (LR) test. The detailed results can be found in Table 4.8 in Appendix 4.E, but they boil down to one observation: both the assumption of homogeneity and perfect optimisation are *rejected* by the data. Furthermore, we explicitly tested for $z_{i,t} = \rho_{i,t}^e/6$ when estimating individual rules. Estimations identified 11 subjects (23%) as consistent traders (see footnote 12 for a detailed discussion). We conclude that this is evidence for heterogeneity of individual sophistication. The majority of the subjects identifies the variables that are relevant to their economic decision, but only a minority is able to learn the (mathematically) optimal solution. Instead of optimising in the mathematical sense, most subjects follow simple rules of thumb.

This result has important implications for economic modelling. The RE hypothesis is built on homogeneous and model consistent expectations, which the agents in turn use to optimise their decisions. Many economists find the first element of RE unrealistic: it is difficult for the agents to form rational expectations due to limited

⁸We use ML since the optimality constraint does not exclude heterogeneity of the idiosyncratic shocks $\eta_{i,t}$ and so the model is non-linear. We exclude outliers defined as observations when a subject would predict the asset to have its return higher than 60 in absolute terms. To account for the initial learning, we exclude the first ten periods from the sample. We also drop subjects 4 and 5 from group 6, since they would always pass $z_{i,t} = 0$ for $t > 10$. Interestingly, these two subjects had non-constant price predictions, which suggests that they were not optimisers.

understanding of the structure of the economy. But the second part of RE is often taken as a good approximation: agents should make an optimal decision *conditional* on what they think about the economy, *even* if their forecast is wrong. Our subjects were endowed with as much information as possible, including an asset return calculator, a table for profits based on the predicted asset return and chosen quantity and the explicit formula for profits; and yet many failed to behave optimally. The design of the randomised payoff excludes risk hedging as a potential reason. The simplest explanation is that individuals in general lack the computation capacity to make perfect optimisations.

Result 3. *The subjects' quantity decisions are not conditionally optimal to their price forecasts in the Mixed treatment. We conclude that **Hypothesis 4** is rejected.*

4.4.5 Estimation of individual behavioural rules

Prior experimental work (Heemeijer et al., 2009) suggests that in LtF experiments, subjects use heterogeneous forecasting rules which nevertheless can be typically described a linear First-Order Rule

$$(4.16) \quad p_{i,t}^e = \alpha_i p_{t-1} + \beta_i p_{i,t-1}^e + \gamma_i (p_{t-1} - p_{t-2}).$$

Two important special cases of (4.16) are pure **trend following** rule with $\alpha_i = 1$ and $\beta_i = 0$, yielding

$$(4.17) \quad p_{i,t}^e = p_{t-1} + \gamma_i (p_{t-1} - p_{t-2}),$$

and **adaptive expectations** with $\gamma_i = 0$ and $\alpha_i + \beta_i = 1$, namely

$$(4.18) \quad p_{i,t}^e = \alpha_i p_{t-1} + (1 - \alpha_i) p_{i,t-1}^e.$$

To explain the trading behaviour of the subjects from the LtO and Mixed treatments, we estimate a general trading strategy in the following specification:

$$(4.19) \quad z_{i,t} = \begin{cases} \text{constant}_i + \chi_i z_{i,t-1} + \phi_i \rho_{t-1}, & \text{(LtO)} \\ \text{constant}_i + \chi_i z_{i,t-1} + \phi_i \rho_{t-1} + \zeta_i \rho_{i,t+1}^e. & \text{(Mixed)} \end{cases}$$

This rule can capture the most relevant possible elements of the individual trading and has two interesting special cases. First, what we call **persistent demand** ($\phi_i = \zeta_i = 0$)

characterised by a simple AR process:

$$(4.20) \quad z_{i,t} = \text{constant}_i + \chi_i z_{i,t-1}.$$

A second special case is a **return extrapolation** rule (with $\chi_i = 0$):

$$(4.21) \quad z_{i,t} = \begin{cases} \text{constant}_i + \phi_i \rho_{t-1} & \text{(LtO)}, \\ \text{constant}_i + \phi_i \rho_{t-1} + \zeta_i \rho_{i,t}^e & \text{(Mixed)}. \end{cases}$$

For every subject from the LtF and LtO treatments, we estimate her behavioral heuristic starting with the general forecasting rule (4.16) or the general trading rule (4.19) respectively. To allow for learning, all the estimations are based on the last 40 periods. Testing for special cases of the estimated rules is straightforward: insignificant variables are dropped until all of the rule coefficients are significant at 5% level.⁹

The same estimation approach is used for the Mixed treatment (now also allowing for the expected return term ζ_i).¹⁰ The caveat is that the two rules (4.16) and (4.19) are closely linked. Their contemporary idiosyncratic errors are potentially correlated,¹¹ while the trade decision depends on the contemporary expected forecast (if $\zeta_i \neq 0$). Since the contemporary trade does not appear in the forecasting rule, the forecast based on the rule (4.16) is assumed to be exogenous and can be estimated independently from the trading rule (4.19). This leaves the potential endogeneity only in the trading heuristic (4.19), and we deal it with a simple instrumental variable approach. Estimated price expectations rule (4.16) (again testing for the special cases) yields *fitted* price forecasts of a subject. Next the trading rule (4.19) is estimated both with the *fitted* forecasts as instruments; and directly with the *reported* forecasts. We control for endogeneity by comparing the two estimators with the Hausman test. Finally, the special cases of (4.19) are tested based on *reported* or *fitted* price forecasts accordingly to the Hausman test.¹²

The estimation results can be found in Appendix 4.E, in Tables 4.6, 4.7 and 4.9 respectively for the LtF, LtO and Mixed treatments.

⁹Adaptive expectations (4.18) impose a restriction $\alpha \in [0, 1]$ (with $\alpha = 1 - \beta$), so we follow here a simple ML approach. If $\alpha_i > 1$ ($\alpha_i < 0$) maximises the likelihood for (4.18), we use the relevant corner solution $\alpha_i = 1$ ($\alpha_i = 0$) instead. We check the relevance of the two constrained models (trend and adaptive) with the Likelihood Ratio test against the likelihood of (4.16).

¹⁰See footnote 7.

¹¹This can happen *e.g.* when a subject makes the two decisions at the same time.

¹²Whenever the estimations indicated that a subject from the Mixed treatment used a return extrapolation rule (4.21) of a form $z_{i,t} = \zeta_i \rho_{i,t}^e$, that is a rule in which only the implied expected return was significant, we directly tested $\zeta_i = 1/6$. This restriction implies trading consistently with the price forecast, which we could not reject for 11 out of 48 subjects.

In order to quantify whether agents use different decision rules in different treatments, we test the differences of the coefficients in the decision rules with the rank sum test. The LtF treatment can be directly compared to the Mixed treatment according to coefficients in equation (4.16), and the LtO treatment can be compared to the Mixed treatment based on equation 4.19. Since the LtF design implies an optimal trade conditional on the price forecast, one can show that a forecasting rule (4.16) with coefficients $(\alpha_i, \beta_i, \gamma_i)$ is approximately equivalent to a trading rule with coefficients $\chi_i = \beta_i$ and $\phi_i = (\alpha_i + \gamma_i - R)/6$. Thus, the LtF and LtO treatment are also comparable in terms of adaptiveness or conservatism (former terms), and response to the asset return (latter terms).

First, when the forecasting rules in the LtF and Mixed treatments are compared, we observe rules with a trend extrapolation terms are popular in both treatments (respectively 39 in LtF and 25 in Mixed out of 48). Other subjects would rarely use a pure adaptive rule (4.18) (none and 3 subjects in the LtF and Mixed treatments respectively), but instead a general FOR (4.16) with insignificant $\gamma_i = 0$. There were none subjects in the LtF treatment, and only 2 in the Mixed treatment, for whom we could not identify a significant forecasting rule. The trend coefficients for both treatments are on average close to $\bar{\gamma} \approx 0.4$ (*i.e.* weak trend following, in line with the previous LtF experiments), and not significantly different in terms of distribution (with MWWT p -value of 0.736). The difference between the two treatments lies in the forecasts conservatism: whereas LtF subjects do prefer an adaptive rule with average coefficient $\bar{\beta} = 0.56$, Mixed treatment subjects almost never use their past predictions while forecasting ($\bar{\beta} = 0.06$ and the two variables have significantly different distribution with MWWT p -value close to zero).

Secondly, when the trading rules in the LtO and Mixed treatment are compared, we find that the rules with a term on past or expected return is the dominating rule in both treatments (33 in the LtO and 32 in the Mixed treatment). There are only 12 subjects using a significant AR1 coefficient χ_i in the LtO treatment, and 8 in the Mixed treatment. This shows that the majority of our subjects tried to extrapolate the asset return dynamics, which leads to a common behaviour of trend chasing. Nevertheless, no significant trading rule was found for 11 LtO treatment and 8 Mixed treatment subjects. The average demand persistence was equal to $\bar{\chi} = 0.07$ and $\bar{\chi} = 0.006$, and the average asset extrapolation was equal to $\bar{\phi} = 0.09$ and $\overline{\phi + \zeta} = 0.06$, for the LtO and Mixed treatment respectively.¹³ The distributions of the two rule coefficients is

¹³Notice that the Mixed treatment trading rule (4.19) is a function of both the past and the expected asset return, and the latter is both unobservable in the LtO treatment, while being popular among Mixed treatment subjects. In this sense the two rules are not equivalent. For the sake of comparability,

insignificant with MWWT p -values of 0.425 and 0.885 for χ_i and $\phi_i/\phi_i + \zeta_i$ respectively.

Finally, in order to evaluate the difference between the LtF and LtO treatments, we compare the implied trading rules in the LtF treatment with the trading rules in the LtO treatment as discussed above. The trading conservatism, with average $\bar{\beta} = 0.56$ and $\bar{\chi} = 0.07$ for the LtF and LtO treatments respectively, is significantly higher in the LtF treatment (MWWT p -value close to zero). This further means that the implied reaction to the asset return is weaker in the LtF treatment (average implied $\bar{\phi} = -0.03$) than in the LtO treatment (average $\bar{\phi} = 0.09$), and this difference is again confirmed by MWWT p -value close to zero. Hence, the results suggest that the LtF treatment is more stable than the other two treatments because agents use a stronger adaptive component in their forecasting rules.

We conclude that the estimated behavioral rules show that most subjects, regardless of the treatment, follow the observed price trend with a weak trend extrapolation type of rules. However, LtF subjects were much more adaptive in their behavior. This explains more stable dynamics under the LtF treatment. In addition, large individual heterogeneity persists within each of the treatment.

Result 4. *The subjects use similar trading strategies in the LtO and Mixed treatments. While the subjects from the LtF treatment prefer more adaptive type of forecasting strategy, which explains more stability of that treatment. We conclude by rejecting the Hypothesis 5.*

4.5 Conclusions

The origin of asset price bubbles is an important topic for both researchers and policy makers. This chapter investigates the price dynamics and bubble formation in an experimental asset pricing market with a price adjustment rule. A fundamental question about the origins of bubbles we address is: do bubbles arise because agents fail to learn to forecast accurately or because they fail to optimise their trading? We investigate the occurrence, the magnitude and the recurrence of bubbles in three treatments based on the tasks of the subjects: price forecasting, quantity trading and both. Under perfect rationality and perfect competition, these three tasks are equivalent and should lead the subjects to an equilibrium with a constant fundamental price. In contrast, we find none of the experimental markets to show a reliable convergence to the fundamental outcome, and recurring bubbles and crashes occur with the highest frequency

we look at what we interpret as an individual reaction to asset return dynamics: ϕ_i in LtO treatment and $\phi_i + \zeta_i$ in the Mixed treatment.

and magnitude when the subjects submit both a price forecast and a trading quantity decision.

This result shows that the deviation of market prices from the rational expectations equilibrium in former learning to forecast experiments (Hommes et al, 2005, 2008, and Heemeijer et al. 2009) is a robust phenomenon. Moreover, when the subjects act in a learning to optimise environment or submit both a forecast and a quantity, the deviation or asset bubbles become more severe. In contrast to the learning to forecast experiments, the coordination of individual decisions is lower in the trading treatments, which suggests that homogeneity of beliefs is not a necessary condition for mispricing or bubbles to occur. In particular, we provide a rigorous statistical test result on the individual heterogeneity in forecasting and trading strategies within the same treatment and across treatments. In the mixed treatment, in which we directly observe both the trading decisions and price expectations, there are only a quarter of the subjects being able to submit a trading quantity that is conditionally optimal to their price forecasts. We find that while there is no significant difference in the trading strategy in the Mixed and LtO treatment, the subjects in the LtF treatment use adaptive expectations with a higher frequency than their counterparts in the Mixed treatment, which leads to more stabilised price behaviour.

What is the behavioural foundation for the difference in the individual decisions and aggregate market outcomes in the learning to forecast and learning to optimise market? There are several candidate explanations: (1) the quantity decision task is more cognitive demanding than the forecasting task, in particular when the subjects in the LtF treatment are helped by a computer program. Following Rubinstein (2007), we use decision time as a proxy for cognitive load and compare the average decision time in each treatment. It turns out while subjects take significantly longer time in the Mixed treatment than the other two treatments according to MWWT, there is no significant difference between the LtF and LtO treatments. It helps to explain why the markets are particularly volatile in the Mixed treatment, but does not explain why the LtO treatment is more unstable than the LtF treatment. (2) In a LtF treatment, the subjects' goal is to find the accurate forecast. Only the size of the prediction error matters while the sign does not matter. Conversely, in a LtO market it is in a way more important for the subjects to predict the direction of the price movement right, and the size of the prediction error is important only to a secondary degree. (For example, if the subjects predict the return will be high and decided to buy, he can still make a profit if the price goes up far more than he expected, and his prediction error is large.) Therefore, the subjects may have a natural tendency to pay more attention to the trend of the price, which leads to a higher degree of trend extrapolation in the

forecasting rule.

Asset mispricing and financial bubbles can cause serious market inefficiencies, and may become a threat to the overall economic stability, as shown by the 2007 financial and economic crisis. It is therefore crucial to study the origins of assets' mispricing in order to design regulations on the financial market. Proponents of the rational expectations would often claim that the serious asset pricing bubbles cannot arise, because rational economic agents would efficiently arbitrage against it and quickly push the 'irrational' (non-fundamental) investors away from the market. Our experiment suggests otherwise: people exhibit heterogeneous and not necessarily optimal behaviour, but because they are trend-followers, their 'irrational' (non-fundamental) beliefs are correlated. This is reinforced by the *positive feedback* between expectations and realised prices on the asset pricing markets, as stressed e.g. in Hommes (2013). Therefore, price oscillations cannot be mitigated by more rational market investors, and trading heterogeneity persists. As a result, waves of optimism and pessimism can arise despite the fundamentals being relatively stable. A strong policy implication is that the financial authorities should remain skeptical about the moods of the investors: fast increase of asset prices should be considered as a warning signal, instead of a reassuring signal of growth of the economic fundamentals only.

The design of our experiment can be extended to study other topics related to financial bubbles, such as markets with financial derivatives and the housing market. The advantage of our framework is that we can define a constant fundamental with positive dividend process, and the price is easy to calculate, and the same for all participants in the market.¹⁴ However, the subjects in our experiment can short-sell the asset as much as they want in order to profit from the fall of asset price during the market crash, which may not be feasible in real markets. An interesting topic for future research are experimental markets where agents face short selling constraints (Anufriev and Tuinstra, 2013) or the role of financial derivatives in (de)stabilising markets.

Another possible extension is to impose a network structure among the traders, i.e. one trader can only trade with some, but not all the other traders; or traders need to pay a cost in order to be connected to other traders. This design can help us to examine the mechanism of bubble formation in financial networks (Gale and Kariv, 2007), and network games (Galeotti et al., 2010) in general. There has been a pioneering experimental literature by Gale and Kariv (2009) and Choi et al. (2013) that study how network structure influence market efficiency when subjects act as

¹⁴The asset price is usually defined for each transaction in a typical Smith et al. (1988) experiment, but it can also be the same for the whole market if the trading mechanism is a call market system, e.g. Akiyama et al. (2012)).

intermediaries between sellers and buyers. Our experimental setup can be extended to study how network structure influences market efficiency and stability when subjects act as traders of financial assets in the over the counter (OTC) market.

Appendix 4.A Instructions and computer screen

4.A.1 LtF treatment

General information

In this experiment you participate in a market. Your role in the market is a professional **Forecaster** for a large firm, and the firm is a major trading company of an asset in the market. In each period the firm asks you to make a prediction of the market price of the asset. The price should be predicted one period ahead. Based on your prediction, your firm makes a decision about the quantity of the asset the firm should buy or sell in this market. Your forecast is the only information the firm has on the future market price. The more accurate your prediction is, the better the quality of your firm's decision will be. You will get a payoff based on the accuracy of your prediction. You are going to advise the firm for 50 successive time periods.

About the price determination

The price is determined by the following price adjustment rule: when there is more demand (firm's willingness to buy) of the asset, the price goes up; when there is more supply (firm's willingness to sell), the price will go down.

There are several large trading companies on this market and each of them is advised by a forecaster like you. Usually, higher price predictions make a firm to buy more or sell less, which increases the demand and vice versa. Total demand and supply is largely determined by the sum of the individual demand of these firms.

About your job

Your only task in this experiment is to predict the market price in each time period as accurately as possible. **Your prediction in period 1 should lie between 0 and 100.** At the beginning of the experiment you are asked to give a prediction for the price in period 1. When all forecasters have submitted their predictions for the first period, the firms will determine the quantity to demand, and the market price for period 1 will be determined and made public to all forecasters. Based on the accuracy of your prediction in period 1, your earnings will be calculated.

Subsequently, you are asked to enter your prediction for period 2. When all participants have submitted their prediction and demand decisions for the second period, the market price for that period, will be made public and your earnings will be calculated, and so on, for all 50 consecutive periods. The information you can refer to at period t consists of all past prices, your predictions and earnings.

Please note that due to liquidity constraint, your firm can only buy and sell up to

a maximum amount of assets in each period. This means although you can submit any prediction for period 2 and all periods after period 2, if the price in last period is p_{t-1} , and you prediction is p_t^e : the firm's trading decision is constrained by $p_t^e \in [p_{t-1} - 30, p_{t-1} + 30]$. More precisely, **the firm will trade as if** $p_t^e = p_{t-1} + 30$ if $p_t^e > p_{t-1} + 30$, **and trade as if** $p_t^e = p_{t-1} - 30$ if $p_t^e < p_{t-1} - 30$.

About your payoff

Your earnings depend only on the accuracy of your predictions. The earnings shown on the computer screen will be in terms of points. If your prediction is p_t^e and the price turns out to be p_t in period t , your earnings are determined by the following equation:

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

The maximum possible points you can earn for each period (if you make no prediction error) is 1300, and the larger your prediction error is, the fewer points you can make. You will earn 0 points if your prediction error is larger than 7. There is a **Payoff Table** on your table, which shows the points you can earn for different prediction errors. We will pay you in cash at the end of the experiment based on the points you earned. You earn 1 euro for each 2600 points you make.

4.A.2 LtO treatment

General information

In this experiment you participate in a market. Your role in the market is a Trader of a large firm, and the firm is a major trading company of an asset. In each period the firm asks you to make a trading decision on the quantity D_t your firm will BUY to the market. (You can also decide to sell, in that case you just submit a negative quantity.) You are going to play this role for 50 successive time periods. The better the quality of your decision is, the better your firm can achieve her target. The target of your firm is to maximize the expected asset value minus the variance of the asset value, which is also the measure by the firm concerning your performance:

$$(1) \quad \pi_t = W_t - \frac{1}{2} Var(W_t)^2$$

The total asset value W_t equals the return of the per unit asset multiplied by the number of unit you buy D_t . **The return of the asset** is $p_t + y - Rp_{t-1}$, where R is the gross interest rate which equals 1.05, p_t is the asset price at period t , therefore

$p_t - Rp_{t-1}$ is the capital gain of the asset, and $y = 3.3$ is the dividend paid by the asset. We assume the variance of the price of a unit of the asset is $\sigma^2 = 6$, therefore the expected variance of the asset value is $6D_t^2$. Therefore we can rewrite the performance measure in the following way

$$(2) \quad \pi_t = (p_t + y - Rp_{t-1})D_t - 3D_t^2$$

The asset price in the next period p_{t+1} is not observable in the current period. You can make a forecast p_t^e on it. **There is an asset return calculator in the experimental interface** that gives the asset return for each price forecast p_t^e you make. Your own payoff is a function of the value of target function of the firm:

$$(3) \quad \text{Payoff}_t = 800 + 40 * \pi_t$$

This function means you get 800 points (experimental currency) as basic salary, and 40 points for each 1 unit of performance (target function of the firm) you make. If your trades will be unsuccessful, you may lose points and earn less than your basic salary, down to 0. Based on the asset return, you can look up your payoff for each quantity decision you make in the **payoff table**.

You can of course also calculate your payoff for each given forecast and quantity using equation (2) and (3) directly. In that situation you can ask us for a calculator.

About the price determination

The price is determined by the following price adjustment rule: when there is more demand than supply of the asset (namely, more traders want to buy), the price will go up; and when there is more supply than demand of the asset (namely, more people want to sell), the price will go down.

About your job

Your only task in this experiment is to decide the quantity the firm will buy/sell. At the beginning of period 1 you determine the quantity to buy or sell (submitting a positive number means you want to buy, and negative number means you want to sell) for period 1. After all traders submit their quantity decisions, the market price for period 1 will be determined and made public to all traders. Based on the value of the target function of your firm in period 1, your earnings in the first period will be calculated. Subsequently, you make trading decisions for the second period, the market price for that period will be made public and your earnings will be calculated, and so on, for

all 50 consecutive periods. The information you can refer to at period t consists of all previous prices, your quantity decisions and earnings.

Please notice that due to the liquidity constraint of the firm, the amount of asset you buy or sell cannot be more than 5 units. Which means your quantity decision should be between -5 and 5 . The numbers on the payoff table are just examples. You can use any other number such as 0.01 , -1.3 , 2.15 etc., as long as they are within $[-5, 5]$. If when you want to submit numbers with a decimal point, please write a “.”, NOT a “,”.

About your payoff

In each period you are paid according to equation (3). The earnings shown on the computer screen will be in terms of points. We will pay you in cash at the end of the experiment based on the points you earned. You earn 1 euro for each 2600 points you make.

4.A.3 Mixed treatment

General information

In this experiment you participate in a market. Your role in the market is a Trader of a large firm, and the firm is a major trading company of an asset. In each period the firm asks you to make a trading decision on the quantity D_t your firm will BUY to the market. (You can also decide to sell, in that case you just submit a negative quantity.) You are going to play this role for 50 successive time periods. The better the quality of your decision is, the better your firm can achieve her target. The target of your firm is to maximize the expected asset value minus the variance of the asset value, which is also the measure by the firm concerning your performance:

$$(1) \quad \pi_t = W_t - \frac{1}{2} \text{Var}(W_t)^2$$

The total asset value W_t equals the return of the per unit asset multiplied by the number of unit you buy D_t . **The return of the asset** is $p_t + y - Rp_{t-1}$, where R is the gross interest rate which equals 1.05 , p_t is the asset price at period t , therefore $p_t - Rp_{t-1}$ is the capital gain of the asset, and $y = 3.3$ is the dividend paid by the asset. We assume the variance of the price of a unit of the asset is $\sigma^2 = 6$, therefore the expected variance of the asset value is $6D_t^2$. Therefore we can rewrite the performance measure in the following way

$$(2) \quad \pi_t = (p_t + y - Rp_{t-1})D_t - 3D_t^2$$

The asset price in the next period p_{t+1} is not observable in the current period. You can make a forecast p_t^e on it. **There is an asset return calculator in the experimental interface** that gives the asset return for each price forecast p_t^e you make. Your own payoff is a function of the value of target function of the firm:

$$(3) \quad \text{Payoff}_t = 800 + 40 * \pi_t$$

This function means you get 800 points (experimental currency) as basic salary, and 40 points for each 1 unit of performance (target function of the firm) you make. If your trades will be unsuccessful, you may lose points and earn less than your basic salary, down to 0. Based on the asset return, you can look up your payoff for each quantity decision you make in the **payoff table**.

You can of course also calculate your payoff for each given forecast and quantity using equation (2) and (3) directly. In that situation you can ask us for a calculator.

The payoff for the forecasting task is simply a decreasing function of your forecasting error (the distance between your forecast and the realized price). When your forecasting error is larger than 7, you earn 0 points.

$$(4) \quad \text{Payoff}_{forecasting} = \max \left[1300 - \frac{1300}{49} (p_t^e - p_t)^2, 0 \right]$$

About the price determination

The price is determined by the following price adjustment rule: when there is more demand than supply of the asset (namely, more traders want to buy), the price will go up; and when there is more supply than demand of the asset (namely, more people want to sell), the price will go down.

About your job

Your task in this experiment consists of two parts: (1) to make a price forecast; (2) to decide the quantity the firm will buy/sell. At the **beginning of period 1 you submit your price forecast between 0 and 100**, and then determine the quantity to buy or sell (submitting a positive number means you want to buy, and negative number means you want to sell) for period 1, and the market price for period 1 will be determined and made public to all traders. Based on your forecasting error and performance measure for the trading task, in period 1, your earnings in the first period will be calculated.

Subsequently, you make forecasting and trading decisions for the second period, the market price for that period will be made public and your earnings will be calculated,

and so on, for all 50 consecutive periods. The information you can refer to at period t consists of all previous prices, your past forecasts, quantity decisions and earnings. Please notice that due to the liquidity constraint of the firm, the amount of asset you buy or sell cannot be more than 5 units. Which means your quantity decision should always be **between -5 and 5**. The numbers on the payoff table are just examples. You can use **any other numbers** such as 0.01, -1.3, 2.15 etc. as long as they are within $[-5, 5]$.

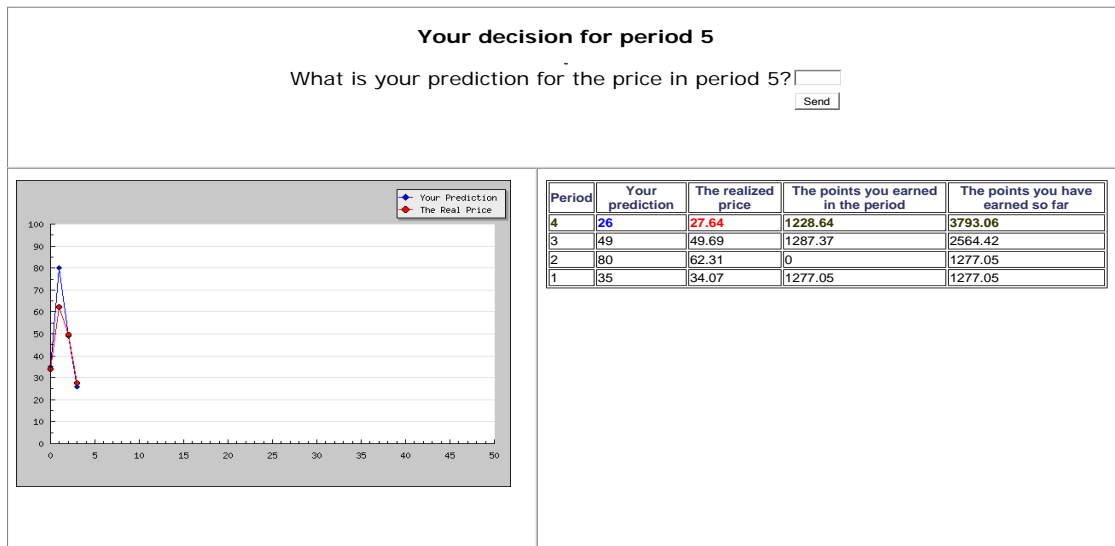
About your payoff

In each period you are paid for the forecasting task according to equation (4) and trading task according to equation (3). The earnings shown on the computer screen will be in terms of points. We will pay you in cash at the end of the experiment based on the points you earned for **either** the forecasting task or the trading task. Which task will be paid will be determined randomly (we will invite one of the participants to toss a coin). **That is, depending on the coin toss, your earnings will be calculated either based on equation (3) or equation (4)**. You earn 1 euro for each 2600 points you make.

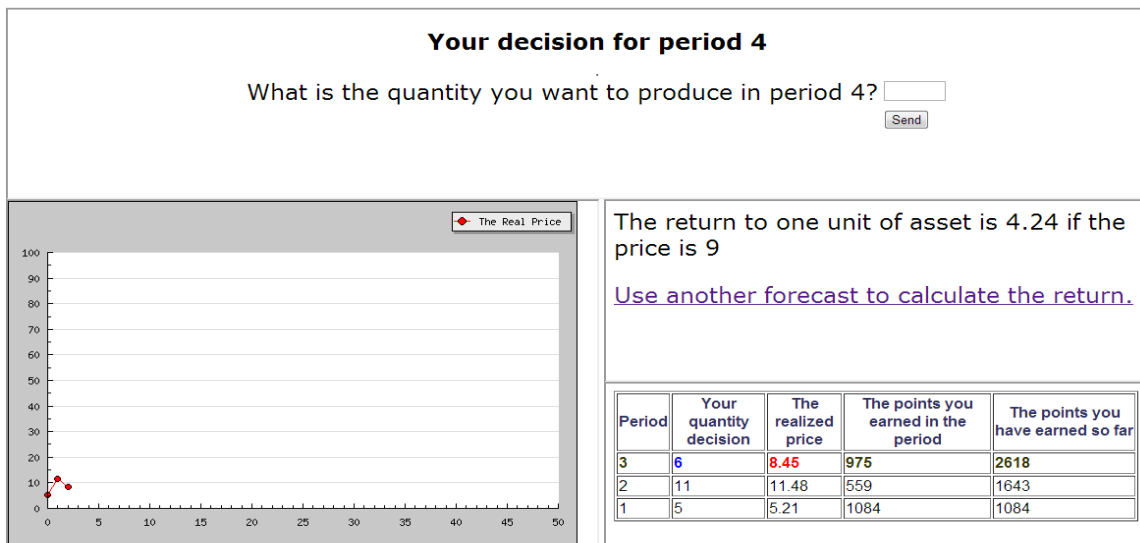
4.A.4 Computer screen

An illustration of the computer screens seen by the subjects is shown on Figure 4.5. The screen was divided into 3 mini pages. In the top mini page, subjects were prompted to submit their decisions, *i.e.*, their price forecast or the amount they want to trade. After submitting their decisions, they go to a waiting page until all the subjects have made their decisions for this period, and then the price and payoff of this period is calculated, the program goes to next period and the screen is reloaded to show the updated information. In the bottom left mini page there was a graph plotting past market prices (the “Real Price”) and, if they were a forecaster, they also saw their own past price forecast history (“Your Prediction”). Finally, in the bottom right mini page they saw a table reporting the history of realized prices, as well as their own prior decisions and cumulative payoffs. If the subject was a quantity decision maker, he/she was also helped by an imbedded calculator. In each period, the subjects could type in their price forecast and press “calculate”, and the calculator will tell them the asset return for this forecast in this period.

Subjects in LtF/LtO treatment saw the screen for a forecaster/trader only. In a Mixed treatment, the subjects first see the screen of the forecast, and then go to the trading page.



(a) Screen for a forecaster.



(b) Screen for a quantity decision maker.

Figure 4.5: Computer screen for subjects in LtF treatment (upper panel) and LtO panel (lower panel).

Appendix 4.B Payoff tables

Payoff Table for Forecasting Task							
Your Payoff= $\max[1300 - \frac{1300}{49}(\text{Your Prediction Error})^2, 0]$							
3000 points equal 1 euro							
error	points	error	points	error	points	error	points
0	1300	1.85	1209	3.7	937	5.55	483
0.05	1300	1.9	1204	3.75	927	5.6	468
0.1	1300	1.95	1199	3.8	917	5.65	453
0.15	1299	2	1194	3.85	907	5.7	438
0.2	1299	2.05	1189	3.9	896	5.75	423
0.25	1298	2.1	1183	3.95	886	5.8	408
0.3	1298	2.15	1177	4	876	5.85	392
0.35	1297	2.2	1172	4.05	865	5.9	376
0.4	1296	2.25	1166	4.1	854	5.95	361
0.45	1295	2.3	1160	4.15	843	6	345
0.5	1293	2.35	1153	4.2	832	6.05	329
0.55	1292	2.4	1147	4.25	821	6.1	313
0.6	1290	2.45	1141	4.3	809	6.15	297
0.65	1289	2.5	1134	4.35	798	6.2	280
0.7	1287	2.55	1127	4.4	786	6.25	264
0.75	1285	2.6	1121	4.45	775	6.3	247
0.8	1283	2.65	1114	4.5	763	6.35	230
0.85	1281	2.7	1107	4.55	751	6.4	213
0.9	1279	2.75	1099	4.6	739	6.45	196
0.95	1276	2.8	1092	4.65	726	6.5	179
1	1273	2.85	1085	4.7	714	6.55	162
1.05	1271	2.9	1077	4.75	701	6.6	144
1.1	1268	2.95	1069	4.8	689	6.65	127
1.15	1265	3	1061	4.85	676	6.7	109
1.2	1262	3.05	1053	4.9	663	6.75	91
1.25	1259	3.1	1045	4.95	650	6.8	73
1.3	1255	3.15	1037	5	637	6.85	55
1.35	1252	3.2	1028	5.05	623	6.9	37
1.4	1248	3.25	1020	5.1	610	6.95	19
1.45	1244	3.3	1011	5.15	596	<i>error</i> ≥ 0	
1.5	1240	3.35	1002	5.2	583		
1.55	1236	3.4	993	5.25	569		
1.6	1232	3.45	984	5.3	555		
1.65	1228	3.5	975	5.35	541		
1.7	1223	3.55	966	5.4	526		
1.75	1219	3.6	956	5.45	512		
1.8	1214	3.65	947	5.5	497		

Table 4.2: Payoff table for forecasters.

		Your profit																					
		Asset quantity: positive number means to buy, negative to sell																					
		-5	-4.5	-4	-3.5	-3	-2.5	-2	-1.5	-1	-0.5	0	0.5	1	1.5	2	2.5	3	3.5	4	4.5	5	
A s s e t r e t u r n	-15	800	1070	1280	1430	1520	1550	1520	1430	1280	1070	800	470	80	0	0	0	0	0	0	0	0	0
	-14	600	890	1120	1290	1400	1450	1440	1370	1240	1050	800	490	120	0	0	0	0	0	0	0	0	0
	-13	400	710	960	1150	1280	1350	1360	1310	1200	1030	800	510	160	0	0	0	0	0	0	0	0	0
	-12	200	530	800	1010	1160	1250	1280	1250	1160	1010	800	530	200	0	0	0	0	0	0	0	0	0
	-11	0	350	640	870	1040	1150	1200	1190	1120	990	800	550	240	0	0	0	0	0	0	0	0	0
	-10	0	170	480	730	920	1050	1120	1130	1080	970	800	570	280	0	0	0	0	0	0	0	0	0
	-9	0	0	320	590	800	950	1040	1070	1040	950	800	590	320	0	0	0	0	0	0	0	0	0
	-8	0	0	160	450	680	850	960	1010	1000	930	800	610	360	50	0	0	0	0	0	0	0	0
	-7	0	0	0	310	560	750	880	950	960	910	800	630	400	110	0	0	0	0	0	0	0	0
	-6	0	0	0	170	440	650	800	890	920	890	800	650	440	170	0	0	0	0	0	0	0	0
	-5	0	0	0	30	320	550	720	830	880	870	800	670	480	230	0	0	0	0	0	0	0	0
	-4	0	0	0	0	200	450	640	770	840	850	800	690	520	290	0	0	0	0	0	0	0	0
	-3	0	0	0	0	80	350	560	710	800	830	800	710	560	350	80	0	0	0	0	0	0	0
	-2	0	0	0	0	0	250	480	650	760	810	800	730	600	410	160	0	0	0	0	0	0	0
	-1	0	0	0	0	0	150	400	590	720	790	800	750	640	470	240	0	0	0	0	0	0	0
0	0	0	0	0	0	50	320	530	680	770	800	770	680	530	320	50	0	0	0	0	0	0	
1	0	0	0	0	0	0	240	470	640	750	800	790	720	590	400	150	0	0	0	0	0	0	
2	0	0	0	0	0	0	160	410	600	730	800	810	760	650	480	250	0	0	0	0	0	0	
3	0	0	0	0	0	0	80	350	560	710	800	830	800	710	560	350	80	0	0	0	0	0	
4	0	0	0	0	0	0	0	290	520	690	800	850	840	770	640	450	200	0	0	0	0	0	
5	0	0	0	0	0	0	0	230	480	670	800	870	880	830	720	550	320	30	0	0	0	0	
6	0	0	0	0	0	0	0	170	440	650	800	890	920	890	800	650	440	170	0	0	0	0	
7	0	0	0	0	0	0	0	110	400	630	800	910	960	950	880	750	560	310	0	0	0	0	
8	0	0	0	0	0	0	0	50	360	610	800	930	1000	1010	960	850	680	450	160	0	0	0	
9	0	0	0	0	0	0	0	0	320	590	800	950	1040	1070	1040	950	800	590	320	0	0	0	
10	0	0	0	0	0	0	0	0	280	570	800	970	1080	1130	1120	1050	920	730	480	170	0	0	
11	0	0	0	0	0	0	0	0	240	550	800	990	1120	1190	1200	1150	1040	870	640	350	0	0	
12	0	0	0	0	0	0	0	0	200	530	800	1010	1160	1250	1280	1250	1160	1010	800	530	200	0	
13	0	0	0	0	0	0	0	0	160	510	800	1030	1200	1310	1360	1350	1280	1150	960	710	400	0	
14	0	0	0	0	0	0	0	0	120	490	800	1050	1240	1370	1440	1450	1400	1290	1120	890	600	0	
15	0	0	0	0	0	0	0	0	80	470	800	1070	1280	1430	1520	1550	1520	1430	1280	1070	800	0	

Note that 3000 points of your profit corresponds to €1.

Table 4.3: Payoff table for traders.

Appendix 4.C RADs and RDs

Treatment	LtF		LtO		Mixed	
Group	RAD	RD	RAD	RD	RAD	RD
#1	10.03	-7.011	18.26	-8.148	38.65	36.84
#2	17.98	-16.94	34.52	-34.52	7.27	-5.657
#3	8.019	-6.048	30.2	-12.95	8.025	4.014
#4	7.285	-5.196	20.63	3.844	42.86	35.46
#5	8.366	4.152	16.55	5.256	14.98	3.341
#6	14.52	6.503	17.51	7.056	23.08	-23.08
#7	4.222	1.104	31.22	23.82	32.14	-18.71
#8	5.365	-0.2539	28.48	26.65	120.7	96.5
Average	9.473	-2.961	24.67	1.376	35.97	16.09

Table 4.4: Relative Absolute Deviation (RAD) and Relative Deviation (RD) of the experimental prices for the three treatments, in percentages.

Appendix 4.D Earnings Ratios

Treatment	Payoff	Payoff REE	Earnings efficiency
LtF			
Market 1	25.69	26.67	96.35%
Market 2	25.19	26.67	94.47%
Market 3	25.61	26.67	96.03%
Market 4	25.65	26.67	96.18%
Market 5	25.38	26.67	95.15%
Market 6	25.09	26.67	94.06%
Market 7	25.65	26.67	96.18%
Market 8	25.75	26.67	96.54%
Average	25.50	26.67	95.62%
LtO			
Market 1	18.80	18.33	102.54%
Market 2	17.46	18.33	95.25%
Market 3	18.00	18.33	98.21%
Market 4	18.41	18.33	100.43%
Market 5	17.85	18.33	97.39%
Market 6	18.27	18.33	99.64%
Market 7	18.07	18.33	98.58%
Market 8	18.04	18.33	98.41%
Average	18.11	18.33	98.81%
Mixed Forecasting			
Market 1	23.36	26.67	87.62%
Market 2	17.94	26.67	67.27%
Market 3	20.17	26.67	75.63%
Market 4	20.64	26.67	77.41%
Market 5	23.22	26.67	87.07%
Market 6	24.52	26.67	91.94%
Market 7	21.65	26.67	81.20%
Market 8	16.21	26.67	60.80%
Average	20.96	26.67	78.62%
Mixed Trading			
Market 1	18.50	18.33	100.89%
Market 2	16.01	18.33	87.33%
Market 3	14.60	18.33	79.61%
Market 4	21.02	18.33	114.63%
Market 5	18.15	18.33	99.03%
Market 6	17.83	18.33	97.24%
Market 7	17.33	18.33	94.55%
Market 8	24.20	18.33	132.01%
Average	18.45	18.33	100.66%

Table 4.5: Average earnings (in Euro) and earnings efficiency for each market.

Appendix 4.E Estimation of individual forecasting rules

Subject	Rule coefficients				R ²	Type
	cons.	Past price	AR(1)	Past trend		
Group 1						
1		0.288	0.756	0.680	0.995	
2	-1.952		1.090	0.448	0.996	
3		1.000		0.744	0.734	TRE
4	-1.349		0.982	0.427	0.998	
5	-2.080	0.307	0.725	0.362	0.997	
6		1.000		0.770	0.648	TRE
Group 2						
1			1.014		0.998	
2		0.626	0.347	0.519	0.998	
3	-2.110	0.346	0.697		0.996	
4			1.013		0.992	
5			1.013		0.997	
6	-1.857	0.475	0.561	0.391	0.996	
Group 3						
1		0.463	0.522	0.707	0.993	
2		0.513	0.495	0.655	0.994	
3		0.476	0.660	0.395	0.993	
4		1.000		0.302	0.310	TRE
5		1.000		0.364	0.390	TRE
6		0.471	0.544	0.579	0.998	
Group 4						
1		0.596	0.568	0.482	0.988	
2		1.000		0.679	0.320	TRE
3		1.000		0.161	0.025	TRE
4	-2.553	0.418	0.621	0.405	0.992	
5		0.389	0.608	0.539	0.996	

Table 4.6: Estimated individual rules for the LtF treatment.

4.E. Estimation of individual forecasting rules

Subject	Rule coefficients				R ²	Type
	cons.	Past price	AR(1)	Past trend		
6		1.000		0.341	0.385	TRE
Group 5						
1		0.260	0.715	0.729	0.990	
2			1.021		0.895	
3	-53.068	-0.369	2.125	-1.591	0.655	
4		0.178	0.902	0.836	0.980	
5		0.452	0.587	0.791	0.993	
6		0.281	0.719	1.245	0.985	
Group 6						
1			0.993		0.880	
2		1.000		0.921	0.507	TRE
3		1.000		0.712	0.761	TRE
4		1.000		0.827	0.804	TRE
5		0.452	0.411	0.977	0.986	
6		1.000		0.804	0.809	TRE
Group 7						
1	6.914		0.902		0.910	
2			1.010		0.998	
3			0.926		0.924	
4		0.359	0.590	0.399	0.966	
5			0.990		0.973	
6		0.308	0.536	0.545	0.960	
Group 8						
1		1.000		0.451	0.293	TRE
2		1.000		0.370	0.502	TRE
3	2.778		0.822	0.470	0.984	
4	7.958		0.884	0.783	0.911	
5		0.316	0.701	0.471	0.992	
6		1.000		0.342	0.081	TRE

Table 4.6: (continued) Estimated individual rules for the LtF treatment.

Subject	Rule coefficients			R ²	rule	stability
	cons.	AR(1)	past return			
Group 1						
1		-0.447	0.203	0.904	mixed	S
2			0.175	0.819	return	U
3			0.167	0.804	return	U
4			0.111	0.856	return	S
5	-0.125		0.168	0.833	return	U
6			0.159	0.854	return	S
Group 2						
1				0.0451	random	S
2				0.168	random	S
3				0.00997	random	S
4				0.106	random	S
5		0.478	-0.0473	0.24	mixed	U
6				0.0473	random	S
Group 3						
1	-0.188	-0.291	0.221	0.836	mixed	U
2			0.16	0.272	return	S
3	-0.26		0.16	0.645	return	S
4			0.0781	0.124	return	S
5		0.283	0.105	0.676	mixed	S
6			0.152	0.879	return	S
Group 4						
1		0.811		0.677	AR(1)	N
2			0.174	0.549	return	U
3			0.113	0.69	return	S
4			0.14	0.824	return	S
5			0.174	0.798	return	U
6			0.119	0.346	return	S

Table 4.7: Estimated individual rules for the LtO treatment. S, N and U denote respectively stable, neutrally stable and unstable rule if all six subjects would use this rule.

4.E. Estimation of individual forecasting rules

Subject	Rule coefficients			R^2	rule	stability
	cons.	AR(1)	past return			
Group 5						
1				0.0975	random	S
2				0.0695	random	S
3		0.579		0.333	AR(1)	N
4				0.00356	random	S
5				0.0238	random	S
6			0.0487	0.183	return	S
Group 6						
1				0.0496	random	S
2			0.135	0.588	return	S
3			0.125	0.854	return	S
4		0.566		0.663	AR(1)	N
5			0.108	0.468	return	S
6			0.148	0.595	return	S
Group 7						
1		0.29	0.0795	0.741	mixed	S
2		0.743		0.551	AR(1)	N
3		-0.3	0.177	0.759	mixed	S
4		0.44	0.0893	0.675	mixed	S
5	0.136	0.269	0.0521	0.59	mixed	S
6			0.156	0.884	return	S
Group 8						
1			0.2	0.258	return	U
2			0.118	0.439	return	S
3	0.118		0.207	0.757	return	U
4	0.0522		0.0482	0.546	return	S
5				0.131	random	S
6			0.143	0.703	return	S

Table 4.7: (continued) Estimated individual rules for the LtO treatment.

Type of bias	Parameters	Sample restriction		
		None	$\rho_{i,t}^e \leq 60$	$\rho_{i,t}^e \leq 30$
LogLikelihood				
Heterogeneous return				
Heterogeneous (unrestricted)	138	-2843.05	-2787.24	-2737.28
Common ($c_i = c$)	93	-2926.99 (0.00000)	-2872.52 (0.00000)	-2825.44 (0.00000)
No ($c_i = 0$)	92	-2927.43 (0.00000)	-2873.08 (0.00000)	-2825.99 (0.00000)
Common return				
Heterogeneous ($\phi_i = \phi$)	93	-3460.58 (0.00000)	-3169.44 (0.00000)	-3075.8 (0.00000)
Common ($\phi_i = \phi, c_i = c$)	48	-3496.66 (0.00000)	-3262.95 (0.00000)	-3177.11 (0.00000)
No ($\phi_i = \phi, c_i = 0$)	47	-3497.7 (0.00000)	-3265.54 (0.00000)	-3179.06 (0.00000)
Perfect return				
Heterogeneous ($\phi_i = 1/6$)	92	-3666.97 (0.00000)	-3242.27 (0.00000)	-3129.37 (0.00000)
Common ($\phi_i = 1/6, c_i = c$)	47	-3795.54 (0.00000)	-3371.24 (0.00000)	-3259.28 (0.00000)
No ($\phi_i = 1/6, c_i = 0$)	46	-3796.43 (0.00000)	-3371.48 (0.00000)	-3259.48 (0.00000)
No return				
Heterogeneous ($\phi_i = 0$)	92	-3461 (0.00000)	-3443.13 (0.00000)	-3415.29 (0.00000)
Common ($\phi_i = 0, c_i = c$)	47	-3496.99 (0.00000)	-3479.39 (0.00000)	-3451.83 (0.00000)
No ($\phi_i = 0, c_i = 0$)	46	-3498.1 (0.00000)	-3480.5 (0.00000)	-3452.93 (0.00000)
<i>Observations</i>		1840	1840	1840

Table 4.8: Mixed treatment: quantities chosen by individuals explained by their contemporary expected asset returns. Log-likelihood measures for models with various restrictions on the parameters and parameter heterogeneity. In parenthesis, likelihood ratio test p-values for the restrictions imposed in the estimation on the unrestricted model (reported in first row). Estimation for 46 individuals, unrestricted sample and sample restricted for observations with expected asset return above 60 and 30.

Subject	Quantity rule coefficients			Price prediction rule coefficients				Type	Stability
	cons.	AR(1)	Exp. return	Past return	cons.	AR(1)	Past price		
Group 1									
1	0.112		0.161						U
2[†]			0.167		0.347	0.650	0.946		S
3		-0.485		0.373	5.654	0.936	1.210		U
4[†]			0.167		6.335	0.929	0.842		S
5	0.092	0.714				1.000	0.884	<i>TRE</i>	N
6				0.129		1.000	1.089	<i>TRE</i>	S
Group 2									
1			0.037		48.425	0.216			S
2		0.580			14.590	0.759	-0.285		N
3[†]			0.167		22.056	0.649			S
4[†]			0.167		17.683	0.712	-0.649		S
5						0.251	0.749	<i>ADA</i>	N
6		0.327							N
Group 3									
1			-0.807			0.212	0.788	<i>ADA</i>	U
2					-13.408	0.535	0.661		N
3					9.092		0.867		N
4				-0.319			1.004		U
5[†]			0.167		-12.399	0.302	0.878		U
6		0.635				0.402	0.598	<i>ADA</i>	N

Table 4.9: Estimated individual rules for the mixed treatment (system of quantity and predicted price rules). S, N and U denote respectively stable, neutrally stable and unstable system of rules if all six subjects would use this system. ADA and TRE denote a price prediction rule of a subject that could be classified as adaptive or trend extrapolation expectations respectively.

Subject	Quantity rule coefficients			Price prediction rule coefficients				Stability
	cons.	AR(1)	Exp. return	cons.	AR(1)	Past price	Past trend	
Group 4								
1 [†]			0.167			0.931		S
2	0.256		0.115	9.748		0.887		S
3 [†]			0.167	3.696	-0.421	1.381	0.994	S
4 [†]			0.167			1.000	0.669	TRE
5	0.870		0.112			1.000	1.022	TRE
6	0.916					1.000	0.996	TRE
Group 5								
1					-0.831	1.795		N
2 [†]			0.167	6.572		1.257	0.814	U
3						1.000	0.860	TRE
4	0.093		0.065	0.096		1.000	1.158	TRE
5			0.149	7.634		0.886	1.231	U
6	0.564		0.220	6.024		1.103		U
Group 6								
1	-0.129		0.216	23.173		0.539		S
2						1.119		S
3	0.657							U
4*				15.897	0.685		0.618	N
5*				28.577		0.431		N
6						1.044		N

Table 4.9: Estimated individual rules for the mixed treatment (system of quantity and predicted price rules). S, N and U denote respectively stable, neutrally stable and unstable system of rules if all six subjects would use this system. ADA and TRE denote a price prediction rule of a subject that could be classified as adaptive or trend extrapolation expectations respectively.

Subject cons.	Quantity rule coefficients			Price prediction rule coefficients				Type	Stability
	AR(1)	Exp. return	Past return	cons.	AR(1)	Past price	Past trend		
Group 7									
1		0.048				1.000			S
2[†]		0.167				1.000	0.543	<i>TRE</i>	S
3			0.085			0.926			S
4						1.054			N
5[†]		0.167				1.000	0.471	<i>TRE</i>	S
6		-0.202	0.203		0.585	0.424	0.599		U
Group 8									
1			0.135			1.014			S
2		0.131		5.329	0.953		1.057		S
3		0.195				1.000	0.898	<i>TRE</i>	U
4	0.816			2.947		0.983	0.902		N
5	-0.572		0.281			1.000	0.801	<i>TRE</i>	U
6	0.815					1.000	0.944	<i>TRE</i>	N

Table 4.9: Estimated individual rules for the mixed treatment (system of quantity and predicted price rules). S, N and U denote respectively stable, neutrally stable and unstable system of rules if all six subjects would use this system. ADA and TRE denote a price prediction rule of a subject that could be classified as adaptive or trend extrapolation expectations respectively.

Appendix 4.F Rational strategic behaviour

Our experimental results are clearly different from the predictions of the rational expectation equilibrium (REE). Previous sections discussed some evidence that non-fundamental prices and oscillations are caused by bounded rationality and simple individual heuristics of our subjects. However, we also found that they would typically earn high payoffs, implying some sort of successful profit seeking behaviour.

In this appendix, we discuss whether rational strategic behaviour can explain our experimental results. Indeed, different types of rational equilibria may exist depending on agents' perception of the economy. Three cases are discussed: (1) agents are price takers; (2) agents know their market power and coordinate on monopolistic behaviour; (3) agents know their market power but play a non-cooperative game. We show that in the price-taking case, the LtF and LtO treatments are equivalent, with the same rational fundamental solution. If the subjects behave strategically or try to collude, the economy can have alternative rational equilibria, where the subjects collectively 'ride a bubble', or jump around the fundamental price. Nevertheless, these rational equilibria predict different outcomes than the individual and aggregate behaviour observed in the experiment.

Without loss of generality, in the case of the non-cooperative game we focus on the one-shot game version of the experimental market to derive our results. More precisely, we look at the optimal decisions that the agents in period t (knowing prices and individual traded quantity until and including period t) have to formulate *only* for the next period $t + 1$. We follow this approach for two reasons. First, by definition agents are myopic and their payoff in $t + 1$ depends only on the realised profit from that period, and not on the stream of future profits from period $t + 2$ onward. Second, the experiment is a repeated game with finitely many repetitions, and subjects knew it would end after 50 periods. Using the standard backward induction reasoning, one can easily show that a sequence of one-period game equilibria forms a rational equilibrium of the finitely repeated game as well.

4.F.1 Price takers

Realised utility of investors in the LtO treatment is given by (4.4) and is equivalent to the following form:

$$(4.22) \quad U_{i,t}(z_{i,t}) = z_{i,t}(p_{t+1} + y - Rp_t) - \frac{a\sigma_z^2}{2} z_{i,t}^2,$$

where $z_{i,t}$ is the traded quantity and $U_{i,t}$ is a quadratic function of the traded quantity. As shown in Section 2 discussed that, assuming the agent is a price taker, the optimal traded quantity conditional on the expected price $p_{i,t+1}^e$ is given by

$$(4.23) \quad z_{i,t}^{PT} = \arg \max_{z_{i,t}} U_{i,t} = \frac{p_{i,t+1}^e + y - Rp_t}{a\sigma_z^2}.$$

Note that this result relies on the assumption that the subjects do not know the price determination function. We argue that the subjects also have an incentive to minimise their forecasting error. To see that, suppose that the realised market price in the next period is p_{t+1} , and the subject makes a prediction error of ϵ , *i.e.* her prediction is $p_{i,t+1}^e = p_{t+1} + \epsilon$. The payoff function can be rewritten as:

$$(4.24) \quad \begin{aligned} U_{i,t}(z_{i,t}) &= z_{i,t}(p_{t+1} + y - Rp_t) - \frac{a\sigma_z^2}{2}z_{i,t}^2 \\ &= \frac{(p_{t+1} + \epsilon + y - Rp_t)(p_{t+1} + y - Rp_t)}{a\sigma_z^2} - \frac{(p_{t+1} + \epsilon + y - Rp_t)^2}{2a\sigma_z^2} \\ &= \frac{(p_{t+1} + y - Rp_t)^2}{2a\sigma_z^2} - \frac{\epsilon^2}{2a\sigma_z^2}. \end{aligned}$$

This shows that utility is maximised when $\epsilon = 0$. Assuming perfect rationality and price taking behaviour (perfect competition), the task of finding the optimal trade *coincides* with the task of minimizing the forecast error. Thus when all agents have rational expectations and are price takers, the market price equals the REE regardless of the task.

Finding 7. *When the subjects act as price takers, the utility function in the Learning to Optimise treatment is a quadratic function of the prediction error, the same (up to a monotonic transformation) as in the Learning to Forecast treatment. Hence, LtF and LtO treatments have equivalent tasks under Rational Expectations.*

4.F.2 Collusive outcome

Suppose now that the agents realise how their trading quantities influence the price and are able to coordinate on a common strategy. This results in a collusive market, similar to a producers' oligopoly. We assume that in the collusive solution, all agents behave as a monopoly that maximises joint (unweighted) utility; thus the solution is symmetric, that is for each agent i , $z_{i,t} = z_t$. In our experiment the price determination

function is:

$$(4.25) \quad p_{t+1} = p_t + 6\lambda z_t,$$

and so the monopoly under perfect rationality maximises

$$(4.26) \quad \begin{aligned} U_t &= \sum_{i=1}^6 U_{i,t}(z_t) = 6 \left[z_t(p_{t+1} + y - Rp_t) - \frac{a\sigma_z^2}{2} z_t^2 \right] \\ &= 6 \left[z_t^2 \left(6\lambda - \frac{a\sigma_z^2}{2} \right) + z_t(y - rp_t) \right]. \end{aligned}$$

Notice that when $\lambda = 20/21, a\sigma_z^2 = 6$, as in the experiment, the coefficient before z_t^2 is positive, $6\lambda - \frac{a\sigma_z^2}{2} = \frac{19}{7} > 0$, and thus the profit function is U shaped, instead of inversely U shaped.¹⁵ This means that a finite global maximum does not exist (utility goes to $+\infty$ when z_t goes to either $+\infty$ or $-\infty$). The global *minimum* is obtained when $z_{i,t} = \frac{7}{38}(rp_t - y) = \frac{7r}{38}(p_t - p^f)$.

In our experiment, the subjects are constrained to choose a quantity from $[-5, +5]$ and the price is bound to the interval $[0, 300]$. Collusive equilibrium in the one-shot game implies that the subjects coordinate on $z_{i,t} = 5$ or $z_{i,t} = -5$, depending on which is further away from $\frac{7(rp_t - y)}{38}$ (as (4.26) is a symmetric parabola). Since $\frac{7(rp_t - y)}{38} > 0$ when the price is above the fundamental ($p^f = y/r$), we can see that the agents coordinate on -5 if the price is higher than the REE ($p_t > y/r$). Similarly, rational agents coordinate on $+5$ if the price is lower than the REE ($p_t < y/r$). If the price is exactly at the fundamental, rational agents are indifferent between -5 and 5 . Notice that in such a case trading the REE quantity ($z_{i,t} = 0$) gives the global *minimum* for the monopoly.

As a consequence, the collusive outcome predicts that the subjects will ‘jump up and down’ around the fundamental. When the price is just below the fundamental, rational agents will buy the asset, which brings the price above the fundamental, and hence the agents in the next period will sell the asset, and so forth. Notice that if the initial price is far below (above) the fundamental, the monopoly will buy (sell) the asset until the price overshoots (undershoots) the fundamental. Then the subjects start to ‘jump up and down’ as described before.

¹⁵If $6\lambda - \frac{a\sigma_z^2}{2} < 0$, this objective function is inversely U shaped. The maximum point is achieved when $z_{i,t} = \frac{rp_t - y}{12\lambda - a\sigma_z^2}$. This means when $p_t = \bar{y}/r$, namely when the price is at the REE, the optimal quantity under collusive equilibrium is still 0. When the price is higher or lower than the REE, the optimal quantity increases with the difference between the price and the REE. This means there is a continuum of equilibria when the economy does not start at the REE.

Finding 8. *When the subjects know the price determination function and are able to form a coalition, the collusive profit function in the LtO treatment is U shaped. Subjects would buy under-priced and sell an over-priced asset. In the long run rational collusive subjects will alternate their trading quantities between -5 and 5 and so the price will alternate around the equilibrium.*

Such alternating dynamics would resemble coordination on contrarian type of behaviour, but has not been observed in any of the experimental groups. Instead, our subjects coordinated on trend-following trading rules, which resulted in smooth, gradual price oscillations.

Notice also that the demand at the edge of the liquidity constraint ($z_{i,t} = \pm 5$) would generate rapid price changes, namely $p_{t+1} = p_t \pm (20/21)/(6 * 5) \approx p_t \pm 28.57$, that is the price would change *in every period* by around 28.57 in absolute terms. This has not been observed even in the super-bubble group 8 from the Mixed treatment. Indeed, quantity decisions equal to 5 or -5 happened only 7 times in the LtO and 44 times in the Mixed treatment (*i.e.* 0.39% and 1.83% of observations respectively). Typical subject behaviour was much more conservative: 97% and 91% traded quantities in the LtO and Mixed treatments respectively were confined in the interval $[-2.5, 2.5]$. A good example is group 4 from the Mixed treatment, in which the price reached 150, but the individual trades were rarely outside the interval $[-3, 3]$ (see Fig. 4.3).

4.F.3 Perfect information non-cooperative game

Consider a scenario, in which the subjects realise the experimental price determination mechanism, but cannot coordinate their actions and play a symmetric Nash equilibrium (NE) instead of the collusive one. There is a positive externality of the subjects' decisions: when one subject buys the asset, it pushes up the price and also the benefits of all the other subjects. The collusive equilibrium internalises this externality, but could the same happen if the subjects in the experiment could not coordinate? In other words, would they have an incentive to 'free-ride' on the demand of others, and would that push the price back to the REE?

In the case of a non-cooperative one-shot game, we again focus on a symmetric solution. Consider agent i , who optimises her quantity choice believing that all other agents will choose z_t^o . This means that the price at $t + 1$ becomes

$$(4.27) \quad p_{t+1} = p_t + 5\lambda z_t^o + \lambda z_{i,t}.$$

Agent i maximises therefore

$$\begin{aligned}
 U_{i,t} &= z_{i,t} (\lambda z_{i,t} + 5\lambda z_t^o + y - rp_t) - \frac{a\sigma_z^2}{2} z_{i,t}^2 \\
 (4.28) \quad &= z_{i,t}^2 \frac{2\lambda - a\sigma_z^2}{2} + z_{i,t} (5\lambda z_t^o + y - rp_t).
 \end{aligned}$$

Notice that $2\lambda - a\sigma_z^2 = -86/21 < 0$. This is an inversely U shaped parabola with the unique maximum given by the reaction function

$$(4.29) \quad z_{i,t}^*(z_t^o) = \frac{5\lambda z_t^o + y - rp_t}{a\sigma_z^2 - 2\lambda}.$$

A symmetric solution requires $z_{i,t}^*(z_t^o) = z_t^o$, which implies

$$(4.30) \quad z_t^* = \frac{rp_t - y}{7\lambda - a\sigma_z^2} = \frac{3}{2}(rp_t - y).$$

Furthermore the reaction function $z_{i,t}^*(z_t^o)$ is linear with respect to z_t^o , with a slope $\frac{5\lambda}{a\sigma_z^2 - 2\lambda} = \frac{100}{86} > 1$. Thus, $z_t^o > z_t^*$ ($<$ and $=$) implies $z_{i,t}^* > z_t^o$ ($<$ and $=$), or in words, if agent i believes that the other players will buy (sell) the asset, she has an incentive to buy (sell) *even more*. Then as a best response, the other agents should further increase/decrease their demand, and this is limited only by the liquidity constraints. The strategy (4.30) thus defines the threshold point between the two corner strategies, *i.e.* the full NE strategy is defined as

$$(4.31) \quad z_{i,t}^{NE} = \begin{cases} 5 & \text{if } z_t^o > z_t^* \\ z_t^* & \text{if } z_t^o = z_t^* \\ -5 & \text{if } z_t^o < z_t^*. \end{cases}$$

The boundary strategies can be infeasible if the previous price is too close to zero or 300.¹⁶ To sum up, as long as the price p_t is sufficiently far from the edges of the allowed interval $[0, 300]$, there are *three* NE of the one-shot non-cooperative game, which are defined by all players playing $z_{i,t} = -5$, $z_{i,t} = z_t^*$ and $z_{i,t} = +5$ for all $i \in \{1, \dots, 6\}$.

If the agents coordinate on the strategy $z_{i,t} = z_t^*$, the price evolves according to the following law of motion:

$$(4.32) \quad p_{t+1} = \frac{10p_t - 60y}{7}.$$

¹⁶Notice that we can interpret z_t^o as the average quantity traded by all other agents, besides agent i , and the reasoning for NE strategy (4.31) remains intact. This implies that NE has to be symmetric.

In contrast to the collusive game, in the non-cooperative game the fundamental price is therefore a possible steady state, but *only* if it is an outcome in the initial period. Additional equilibrium refinements may further exclude it as a rational outcome, since it is the least profitable one. Recall that the subjects earn 0 when they play z_t^* with price at the fundamental (because there is no trade). On the other hand, they may earn a positive profit by coordinating on -5 or 5 . For example, when all of them buy 5 units of asset, the utility for each of them will be $(p_t + y + 6\lambda z_{i,t} - (1+r)p_t)z_{i,t} - \frac{\alpha\sigma^2}{2}z_{i,t}^2 = (33.3 - 0.05p_t) * 5 - 75$. This equals 76.5 when $p_t = 60$, 16.5 when $p_t = 300$ and 75 when the previous price is equal to the fundamental, $p_t = 66$. This explains why the payoff efficiency (average experimental payoff divided by payoff under REE) is larger than 100% in some markets in the LtO or Mixed treatments where prices have large oscillations.

Notice that the linear equation (4.32) is unstable, so the NE of the one-shot game leads to unstable price dynamics in the repeated game even if the agents coordinate on $z_{i,t} = z_t^*$, as long as the initial price is different from the fundamental price. Indeed, if the initial price is 67 or 65 (fundamental price plus or minus one), the price hits the upper cap of 300 or the lower cap of 0 in 16 and 12 periods respectively, and rational non-cooperative agents would be forced to use appropriate corner strategies (-5 and 5 respectively). Furthermore the agents can switch *at any moment* between the three one-shot game NE defined by (4.31). This implies that in the repeated non-cooperative game, many rational price paths are possible. This includes many price paths where agents will often coordinate on 5 or -5 strategies. Furthermore, notice that the up and down alternating price behaviour around the fundamental, which was the solution for the collusive equilibrium, is a NE as well, and hence this is the Pareto efficient equilibrium for this game.

Finding 9. *In the non-cooperative game with perfect information, there are two possible types of NE. The fundamental outcome is a possible outcome only if the initial price is equal to the fundamental price. Otherwise, the agents will coordinate on unstable, possibly oscillatory price dynamics, with traded quantities of -5 or 5 . When they coordinate on a non-zero quantity, their payoff can be higher than their payoff under the REE under the price-taking beliefs.*

Altogether, the perfectly rational agents can coordinate on price boom-bust cycles and earn positive profit. However, this would require even stronger assumptions than the fundamental equilibrium, namely that the agents perfectly understand the underlying price determination mechanism. Furthermore, the cycle of bubbles and crashes is suboptimal in comparison with the ‘jumping up and down around the fundamen-

tal' equilibrium. If the agents were rational enough to coordinate, then it remains a mystery why they would coordinate on the less efficient path of bubbles and crashes.

Furthermore, such rational equilibria with price oscillations predict that the subjects to coordinate on *homogeneous* trades at the edge of the liquidity constraints (traded quantities should often, or even always, be either 5 or -5). The subjects from the LtO and Mixed treatments behaved differently. Their traded quantities were highly heterogeneous (which implied the observed heterogeneity of the estimated trading and forecasting rules), and rarely reached the liquidity constraints, as discussed above.

Therefore, the rational solutions, in particular the ones from the perfect information, non-cooperative games provide some useful insights on why subjects “ride the bubbles” in the LtO and Mixed treatment. However, since the rational solution cannot explain the heterogeneity of the individual decisions and the fact that the subjects shy away from the boundary solutions, the bubbles and crashes we see from the data is probably a result of the joint forces of rational (profit seeking) and boundedly rational behaviour with some coordination on trend-following buy and hold and short sell strategies.