

TI 2015-107/II
Tinbergen Institute Discussion Paper



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

Te Bao^a

Cars Hommes^b

Tomasz Makarewicz^b

^a *University of Groningen, the Netherlands;*

^b *Faculty of Economics and Business, University of Amsterdam, and Tinbergen Institute, the Netherlands.*

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

More TI discussion papers can be downloaded at <http://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 1600

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

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

Te Bao^a Cars Hommes^b Tomasz Makarewicz^b

September 2015, *Economic Journal* forthcoming

^a *University of Groningen*

^b *University of Amsterdam and Tinbergen Institute*

Abstract. This experiment compares the price dynamics and bubble formation in an asset market with a price adjustment rule in three treatments where subjects (1) submit a price forecast only, (2) choose quantity to buy/sell and (3) perform both tasks. We find deviation of the market price from the fundamental price in all treatments, but to a larger degree in treatments (2) and (3). Mispricing is therefore a robust finding in markets with positive expectation feedback. Some very large, recurring bubbles arise, where the price is 3 times larger than the fundamental value, which were not seen in former experiments.

JEL Classification: C91, C92, D53, D83, D84

Keywords: Financial Bubbles, Experimental Finance, Rational Expectations, Learning to Forecast, Learning to Optimise

*The authors are grateful to Alan Kirman for stimulating discussions and to the Editor Andrea Galeotti and two anonymous referees for helpful and detailed comments. We also thank participants at seminars at New York University, University of Goettingen, University of Nijmegen, University of Namur, University of Groningen and conferences/workshops on “Computing in Economics and Finance” 2013, Vancouver, “Experimental Finance” 2014, Zurich, in particular our discussant David Schindler, “Expectations in Dynamic Macroeconomic Models” 2014, Bank of Finland, Helsinki and “Economic Dynamics and Quantitative Finance”, December 2014, Sydney. We gratefully acknowledge the financial support from NWO (Dutch Science Foundation) Project No. 40611142 “Learning to Forecast with Evolutionary Models” and other projects: the EUFP7 projects Complexity Research Initiative for Systemic Instabilities (CRISIS, grant 288501) and Integrated Macro-Financial Modeling for Robust Policy Design (MACFIN-ROBODS, grant 612796), and from the Institute of New Economic Thinking (INET) grant project “Heterogeneous Expectations and Financial Crises” (INO 1200026). Email addresses: T.Bao@rug.nl, C.H.Hommes@uva.nl, T.A.Makarewicz@uva.nl.

This paper 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? Our experiment indicates that both forces have a destabilising effect on the financial markets, which implies that both deviations from rationality deserve more attention in future theoretical or policy-oriented inquiries on bubble formation and market efficiency.

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. We design the payoff functions carefully so that under the assumptions of perfect rationality and price taking behaviour, these three tasks are equivalent in our experiment 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 a slow upward trend, in the treatment where the subjects make forecasts 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 markets, but have not been investigated extensively by modern economics and finance literature. One possible reason is that it contradicts the standard theory of rational expectations (Muth, 1961; Lucas Jr., 1972) 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 movements 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 the preceding boom and bust in the US housing market highlight the importance of understanding the mechanism of financial bubbles in order for policy 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 (see Hommes and in’t Veld, 2014, for a discussion about the S&P 500 example). 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. 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 *e.g.* Lei *et al.* (2001); Noussair *et al.* (2001); Dufwenberg *et al.* (2005); Haruvy and Noussair (2006); Akiyama *et al.* (2012); Haruvy *et al.* (2013); Füllbrunn *et al.* (2014).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 design.

Nevertheless, Kirchler *et al.* (2012); Huber and Kirchler (2012) argue that the non-fundamental outcomes in these 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 ‘stocks of a depletable gold mine’. Another potential concern about these experiments, due to a typically short horizon (15 periods), is that one cannot test whether financial crashes are likely to be followed by new bubbles. It is very important to study the recurrence of boom-bust cycles in asset prices, for example to understand the evolution of the asset prices between the dot-com bubble and crash and the 2007/2008 financial crisis.

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 a complementary ‘learning to forecast’ (henceforth LtF) experimental design introduced by Marimon *et al.* (1993) (see Hommes, 2011; Assenza *et al.*, 2014, for comprehensive surveys). Hommes *et al.* (2005) run an experiment where subjects act as professional advisers (forecasters) for a pension fund: they submit a price forecasts, 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 results are 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 follow price trend extrapolation strategies (*cf.* Bostian and Holt, 2009). Heemeijer *et al.* (2009) and Bao *et al.* (2012) investigate whether the non-convergence

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

result is driven by the positive expectation feedback nature of the experimental market. 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 happens 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 paper 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 mis-pricing in treatment (2) and (3) is much higher than in treatment (1). We also find larger heterogeneity in traded quantities 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 some very large and repeated price oscillations occur, where the price peaks at more than 3 times the fundamental price. This was not observed in the former experimental literature. Since 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 recurrent bubbles.

Another contribution is that we provide an *empirical micro foundation* of ob-

served differences in aggregate macro behaviour across treatments. We estimate individual forecasting and trading rules and find significant differences across treatments. In the LtF treatment individual forecasting behaviour is more cautious in the sense that subjects use a more conservative anchor (a weighted average of last observed price and last forecast) in their trend-following rules, while in the Mixed treatment almost all weight is given to the last observed price leading to a more aggressive trend-following forecasting rule. Individual trading behaviour of most subjects is characterized by extrapolation of past and/or expected returns. Moreover, in the LtO and Mixed treatments the return extrapolation coefficients are higher. These differences in individual behaviour explain the more unstable aggregate behaviour with recurring booms and busts in the LtO and Mixed treatments. We also perform a formal statistical test on individual heterogeneity in trading strategies under the Mixed treatment. In particular, in some trading markets we observe a large degree of heterogeneity in the quantity decision even when the price is rather stable. In the Mixed treatment most subjects fail to trade at the conditionally optimal quantity given their own forecast. Learning to trade optimally thus appears to be difficult.

Our paper 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 feedback system, for which all markets converge to the fundamental price, and find differences in the speed of convergence across treatments.

The paper is organised as follows: Section 1 presents the experimental design and formulates testable hypotheses. Section 2 summarises the experimental results and performs statistical tests of convergence to REE and for differences across treatments based on aggregate variables as well as individual decision rules. Finally, Section 3 concludes.

1 Experimental Design

In this section we explain the design of our experiment. We begin by defining the treatments, followed by a discussion of the information given to the subjects. Thereafter, we derive the micro foundations of the experimental economy, discuss the implementation of the experiment, and specify hypotheses that will be tested empirically.

1.1 Experimental Treatments

The experimental economy is based on a simple asset market with a constant fundamental price. There are $I = 6$ subjects in each market, and each subject plays an advisory role to a professional trading company. Subject task is either to predict the asset price, suggest trading quantity or both, and subjects are rewarded depending on their forecasting accuracy or trading profits. These decisions generate an excess demand that determines the market price for the asset. The experimental sessions last for $T = 50$ trading rounds. To present a quick overview of the treatment designs, we only show the reduced form law of motion of the price in each treatment in this section. The microfoundation of the experimental economy and choice of parameters will be explained in detail in section 1.3.

Based on the nature of the task and the payoff scheme, there are three treatments in the experiment. It is important to note here that the underlying market structure is the same regardless of the subject task. We carefully choose the parameters of the model and payoff function so that under rational expectations, these treatments are equivalent and lead to the same market equilibrium. The treatments are specified as follows:

LtF Classical Learning-to-Forecast experiment. Subjects act as forecasting advisers, namely they are asked for one-period ahead price predictions $p_{i,t+1}^e$. The subjects' reward depends only on the prediction accuracy, defined by (see also Table B.1 in Appendix B)

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

where $p_{i,t+1}^e$ denotes the forecast of price at period $t+1$ formulated by subject i and p_{t+1} is the realised asset price at period $t+1$.

The subject forecasts are automatically translated into excess demand for the asset, yielding the following law of motion for the LtF treatment economy:

$$p_{t+1} = 66 + \frac{20}{21} (\bar{p}_{t+1}^e - 66) + \varepsilon_t, \quad (2)$$

where $p^* = p^f = 66$ is the fundamental price of the asset as well as the unique Rational Expectations Equilibrium, $\bar{p}_{t+1}^e \equiv \frac{1}{6} \sum_{i=1}^6 p_{i,t+1}^e$ denotes the average price forecast of the six subjects and $\varepsilon_t \sim N(0, 1)$ is a small IID shock to price p_{t+1} .

For the price adjustment rule (2) the subjects' payoff is maximised when all predict the fundamental price, so that on average they make the smallest prediction errors. Hence, in the LtF treatment it is optimal for all subjects to predict $p_{i,t+1}^e = 66$.

LtO Classical Learning-to-Optimise experiment, where the subjects are asked to decide on the asset quantity $z_{i,t}$. Unlike the experiments in the spirit of Smith *et al.* (1988), subjects in this treatment *do not* accumulate the asset over periods. Instead, $z_{i,t}$ represents the final position of subject i in period t . This position can be short with $z_{i,t} < 0$ and is cleared once p_{t+1} is realised. Subjects earn payoff based on the realised return ρ_{t+1} , which is defined as a (constant) dividend $y = 3.3$ plus the capital gain over the constant gross interest rate $R = 1.05$ of a secure bond:

$$\rho_{t+1} \equiv p_{t+1} + y - Rp_t = p_{t+1} + 3.3 - 1.05p_t. \quad (3)$$

Subjects 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}^e for any price forecast p_{t+1}^e as in equation (3). Subjects are rewarded according to

$$\text{Payoff}_{i,t+1} = \max \{0, 800 + 40(z_{i,t}(p_{t+1} + 3.3 - 1.05p_t) - 3z_{i,t}^2)\}. \quad (4)$$

This payoff corresponds to a mean-variance utility function of the financial firms in the underlying economy, as explained below. Expected payoff can be computed by the subjects or read from a payoff table, depending on the chosen quantity and the expected excess return (see Table B.2 in Appendix B).

Under the assumption of price-taking behaviour, i.e., when the subjects ignore the impact of their own trading decisions on the realised market price, the optimal demand for asset given one's own price forecast $p_{i,t+1}^e$ is $z_{i,t}^* = \frac{p_{i,t+1}^e + 3.3 - 1.05p_t}{6} = \frac{\rho_{t+1}^e}{6}$.

The law of motion of the LtO treatment is given by the price adjustment rule based on the aggregate excess demand

$$p_{t+1} = p_t + \frac{20}{21} \sum_{i=1}^6 z_{i,t} + \varepsilon_t \quad (5)$$

for the same set of IID shocks ε_t as in the LtF treatment. Under the assumption of price-taking behaviour, the Rational Expectation Equilibrium (REE) of the market is $p^* = p^f = 66$, and the associated optimal demand for the asset is $z^* = 0$ for each individual. Therefore, the optimal choices are equivalent in the LtF and LtO treatments. For other cases in which subjects deviate from price-taking behaviour, e.g. by taking their market power into account and playing collusive or non-cooperative Nash strategies, a detailed discussion is provided in Appendix C.

Mixed Each subject is asked first for his or her price forecast $p_{i,t+1}^e$ and second for the choice of the asset quantity $z_{i,t}$. In order to avoid hedging, the reward for the whole experiment is based on either the payoff in (1) or (4) with equal probability (flip of a coin at the end of the session). The law of motion of the Mixed treatment is given by (5), the same price adjustment rule as in LtO and does not depend on the submitted price forecasts.

The points in each treatment are exchanged into Euro at the end of the experiment with the conversion rate 3000 points = 1 Euro. We add a max function to the forecasting and trading payoffs to avoid negative rewards.

1.2 Information to the Subjects

At the beginning of the experimental sessions, subjects were informed about their task and payoff scheme, including the payoff functions (1) or (4) depending on the treatment. We supplemented the subjects with payoff tables (see Appendix B).

Subjects from the LtF treatment were told that the asset price depends positively on the average price forecast, while subjects in the two other treatments were informed that the price increases with the excess demand. In addition, in the Mixed treatment we made it clear that the subject payoffs may be related to the forecasting accuracy, but that the realized price itself depends exclusively on their trades. Regardless of the treatment, we provided the subjects only with *qualitative* information about the market, that is we did not explicate the respective laws of motion (2) or (5).

Throughout the experiment, the subject could observe past market prices and their individual decisions, in graphical and table form, but they could not see the decisions, or an average decision, of the other participants. We did not mention the fundamental price in the instructions at all, though we did provide the information about the interest rate and the asset dividend in all the three treatments, which could be used to compute the fundamental price $p^* = \bar{y}/r = 66$. Finally, the subjects know the specification of their payoff function, i.e., the payoff is higher if the prediction error (trading profit) is lower (higher) for the forecasters (traders).

1.3 Experimental Economy

This section provides some micro-foundations of our experimental economy. We build our experimental economy upon an asset market with heterogeneous beliefs as in Brock and Hommes (1998). There are $I = 6$ agents, who allocate investment between a risky asset that pays a fixed dividend y and a risk-free bond that pays

a fixed gross return $R = 1 + r$.² The wealth of agent i evolves according to

$$W_{i,t+1} = RW_{i,t} + z_{i,t}(p_{t+1} + y - Rp_t), \quad (6)$$

where $z_{i,t}$ is the demand (in the sense of the final position) for the risky asset by agent i in period t (positive sign for buying and negative sign for selling) and p_t and p_{t+1} are the prices of the risky asset in periods t and $t + 1$ respectively. Let $E_{i,t}$ and $V_{i,t}$ denote the beliefs or forecasts of agent i about the conditional expectation and the conditional variance based on publicly available information. The agents are assumed to be simple myopic mean-variance maximizers of next period's wealth, i.e. they solve the myopic optimisation problem:

$$\max_{z_{i,t}} \left\{ E_{i,t}W_{i,t+1} - \frac{a}{2}V_{i,t}(W_{i,t+1}) \right\} \equiv \max_{z_{i,t}} \left\{ z_{i,t}E_{i,t}\rho_{t+1} - \frac{a}{2}z_{i,t}^2V_{i,t}(\rho_{t+1}) \right\}, \quad (7)$$

where a is a parameter for risk aversion, and ρ_{t+1} is the excess return as defined in equation (3). In the experiment, we use an affine transformation of this utility function as in (4) as a payoff for the trading task.

Optimal demand of agent i is given by³

$$z_{i,t}^* = \frac{E_{i,t}(\rho_{t+1})}{aV_{i,t}(\rho_{t+1})} = \frac{p_{i,t+1}^e + y - Rp_t}{a\sigma^2}, \quad (8)$$

where $p_{i,t+1}^e = E_{i,t}p_{t+1}$ is the individual forecast by agent i of the price in period $t + 1$. The market price is set by a market maker using a simple price adjustment mechanism in response to excess demand (Beja and Goldman, 1980),⁴ given by

$$p_{t+1} = p_t + \lambda (Z_t^D - Z_t^S) + \varepsilon_t, \quad (9)$$

²Fixed dividend allows for a constant fundamental price throughout the experiment. In a more general model with the same demand functions and market equilibrium, y corresponds to the mean of an (exogenous) IID stochastic dividend process y_t ; see Brock and Hommes (1998) for a discussion.

³The last equality in (8) follows from a simplifying assumption made in Brock and Hommes (1998) that all agents have homogeneous and constant beliefs about the conditional variance, i.e. $V_{i,t}(\rho_{t+1}) = \sigma^2$. See Hommes (2013), Chapter 6, for a more detailed discussion.

⁴See *e.g.* Chiarella *et al.* (2009) for a survey on the abundant literature about the price adjustment market mechanisms. We decided to use (9) instead of a market clearing mechanism for two reasons: (i) market maker is a stylized description of a specialist driven market, a common case for financial markets (*e.g.* NASDAQ); and (ii) the current one-period ahead design is much simpler for the subjects than one based on a market clearing mechanism, which requires two-period ahead trading and forecasting. In particular, the two-period ahead trading/forecasting feature would lead to a 3-dimensional payoff table instead of the 2-dimensional payoff table in Appendix B.2. The two-period ahead market clearing design results in much more volatile price patterns in the LtF experiments (Hommes, 2011), which suggests that our main finding –that the boundedly rational trading can be a destabilizing force in the financial markets– is likely to be robust in a similar two-period ahead LtO experiment with a market clearing design.

where $\varepsilon_t \sim N(0, 1)$ is a small IID 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. In the experiment, we take $R\lambda = 1$, specifically $R = 1 + r = 21/20$, $\lambda = 20/21$, $a\sigma_z^2 = 6$, and $y = 3.3$. We chose these specific parameters mainly for simplicity of the law of motion of the price. For example, by imposing $a\sigma_z^2 = 6$, the total excess demand coincide with the average expected excess return, and when $R\lambda = 1$, this ensures that the final law of motion of asset price in the LtF treatment only depends on the average forecast \bar{p}_{t+1}^e , but does not contain p_t . The price adjustment based on aggregate *individual demand* thus takes the simple form

$$p_{t+1} = p_t + \frac{20}{21} \sum_{i=1}^6 z_{i,t} + \varepsilon_t, \quad (10)$$

which constitutes the law of motion (5) for the LtO and Mixed treatments, in which the subjects are asked to elicit their asset demands.

For an optimising agent and the chosen parameters, the individual optimal demand (8) *conditional* on a price forecast $p_{i,t+1}^e$ equals

$$z_{i,t}^* = \frac{\rho_{i,t+1}^e}{a\sigma^2} = \frac{p_{i,t+1}^e + 3.3 - 1.05p_t}{6}, \quad (11)$$

with $\rho_{i,t+1}^e$ the forecast of excess return in period $t + 1$ by agent i . Substituting it back into (5) gives

$$p_{t+1} = 66 + \frac{20}{21} (\bar{p}_{t+1}^e - 66) + \varepsilon_t, \quad (12)$$

where $\bar{p}_{t+1}^e = \frac{1}{6} \sum_{i=1}^6 p_{i,t+1}^e$ is the average prediction of the price p_{t+1} by six subjects.⁵ This price is the temporary equilibrium with point-beliefs about prices and represents the price adjustment process as a function of the average *individual forecast*. It constitutes the law of motion (2) for the LtF treatment, in which the subjects are asked to elicit their price expectations.

We note that from the optimal demand (11) it is clear that optimising the (quadratic) mean-variance utility function (7) is equivalent to minimising the quadratic penalty for forecasting errors as in the LtF payoff function (1). This implies that the trading and forecasting tasks in the experiment are equivalent under perfect rationality.

⁵Heemeijer et al. (2009) used a similar price adjustment rule in a learning to forecast experiment that compares positive versus negative expectation feedback, but their fundamental price is 60 instead of 66.

By imposing the rational expectations condition $\bar{p}_{t+1}^e = p^f = E_t(p_{t+1})$, a simple computation shows that $p^f = 66$ is the unique Rational Expectation Equilibrium (REE) of the system. This fundamental price equals the discounted sum of all expected future dividends, i.e., $p^f = y/r$. If all agents have rational expectations, the realised price becomes $p_t = p^f + \varepsilon_t = 66 + \varepsilon_t$, i.e. the fundamental price plus (small) white noise and, on average, the price forecasts are self-fulfilling. When the price is p^f , the (expected) excess return of the risky asset in (3) equals 0 and the optimal demand for the risky asset in (8) by each agent is also 0, that is excess demand is equal to 0.

1.4 *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 (11) 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. We 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.

1.5 *Number of Observations*

The experiment was conducted 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. Experimental instructions are shown in Appendix A.

1.6 Testable Hypotheses

The RE benchmark suggests that the subjects should learn to play the REE and behave similarly in all treatments. In addition, a rational decision maker should be able to find the optimal demand for the asset given his price forecast according to Equation (11) in the Mixed treatment. These theoretical predictions can be formulated into 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: In the Mixed treatment the quantity decisions by the subjects are optimal conditional on their price expectations;

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. To be specific, we will use Relative (Absolute) Deviation (Stöckl *et al.*, 2010) to measure price convergence, and test the difference of the distribution of this measure between the three treatments. (**HYPOTHESIS 1** and **2**). Relative earnings can be compared with the Mann-Whitney-Wilcoxon rank-sum test (**HYPOTHESIS 3**). Finally, we estimate individual behavioural 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**.

2 Experimental Results

2.1 Overview

Figure 1 (LtF treatment), Figure 2 (LtO treatment) and Figure 3 (Mixed treatment) show plots of the market prices in each treatment. For most of the groups, the prices and predictions remained in the interval $[0, 100]$. The exceptions are markets 1, 4 and 8 (Figures 3a, 3d and 3h) in the Mixed treatment. In the first

two of these three groups, prices peaked at almost 150 (more than twice the fundamental price $p^f = 66$) and in the last group, prices reached 225, almost 3.5 times the fundamental price. Moreover, markets 4 and 8 of the Mixed treatment show repeated booms and busts.

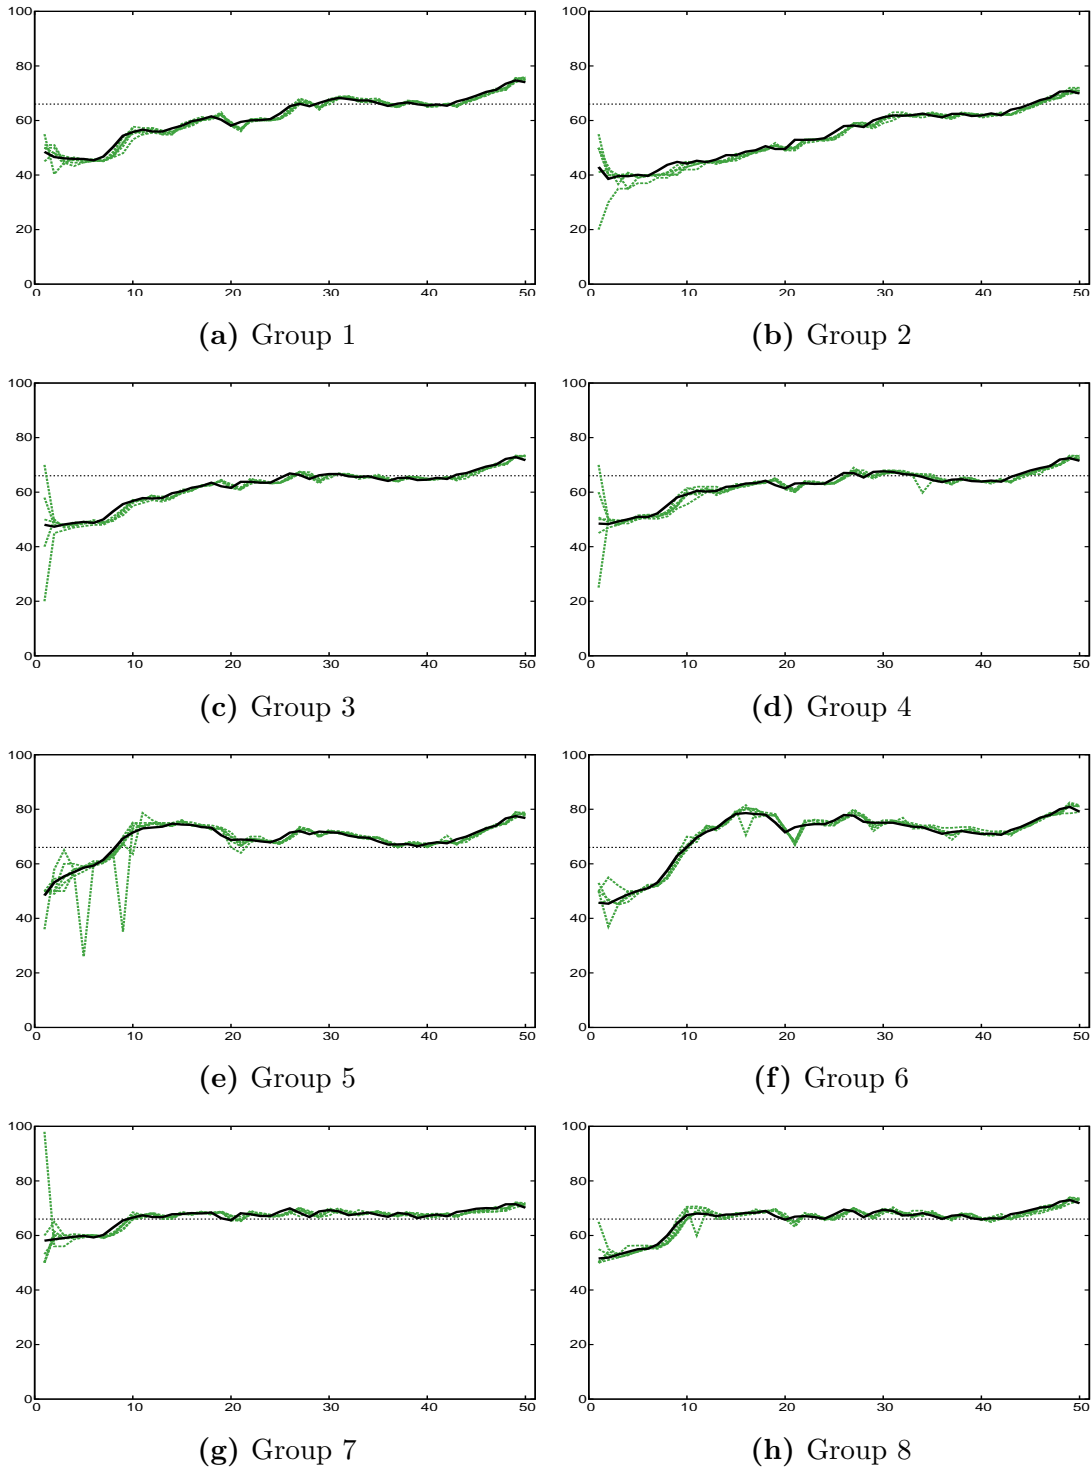
The figures suggest that the market price is the most stable in the LtF treatment, and the most unstable in the Mixed treatment. In the LtF treatment, there is little heterogeneity between the individual forecasts, shown by the green dashed lines. In the LtO treatment, however, there is a high level 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.⁶

It is noticeable that in two markets in the LtO and Mixed treatment, the market price stabilises after a few periods, but stays at a non REE level. Market 2 in the LtO treatment stabilises around price 40, and Market 6 in the Mixed treatment stabilises around price 50. In these two markets, the optimal demand by each individual as implied by (11) should be about 0.2 (0.15) when the price stabilises at 40 (50). However, the actual average demand in the experiment stays very close to 0 in both cases. This is an indication of sub-optimal behaviour by some subjects. It may be caused by two reasons: (1) the subjects mistakenly ignored the role of dividend in the return function, and thought that buying is not profitable unless the price change is strictly positive, or (2) some of them held a pessimistic view about the market, and kept submitting a lower demand than the optimal level as implied by their price forecast.

In general, convergence to the REE does not seem to occur in any of the treatments. This suggests that the hypotheses based on the rational expectations benchmark are likely to be rejected. Furthermore, the figures suggest clear differences between the treatments. In the remainder of this section, we will discuss the statistical evidence for the hypotheses in detail.

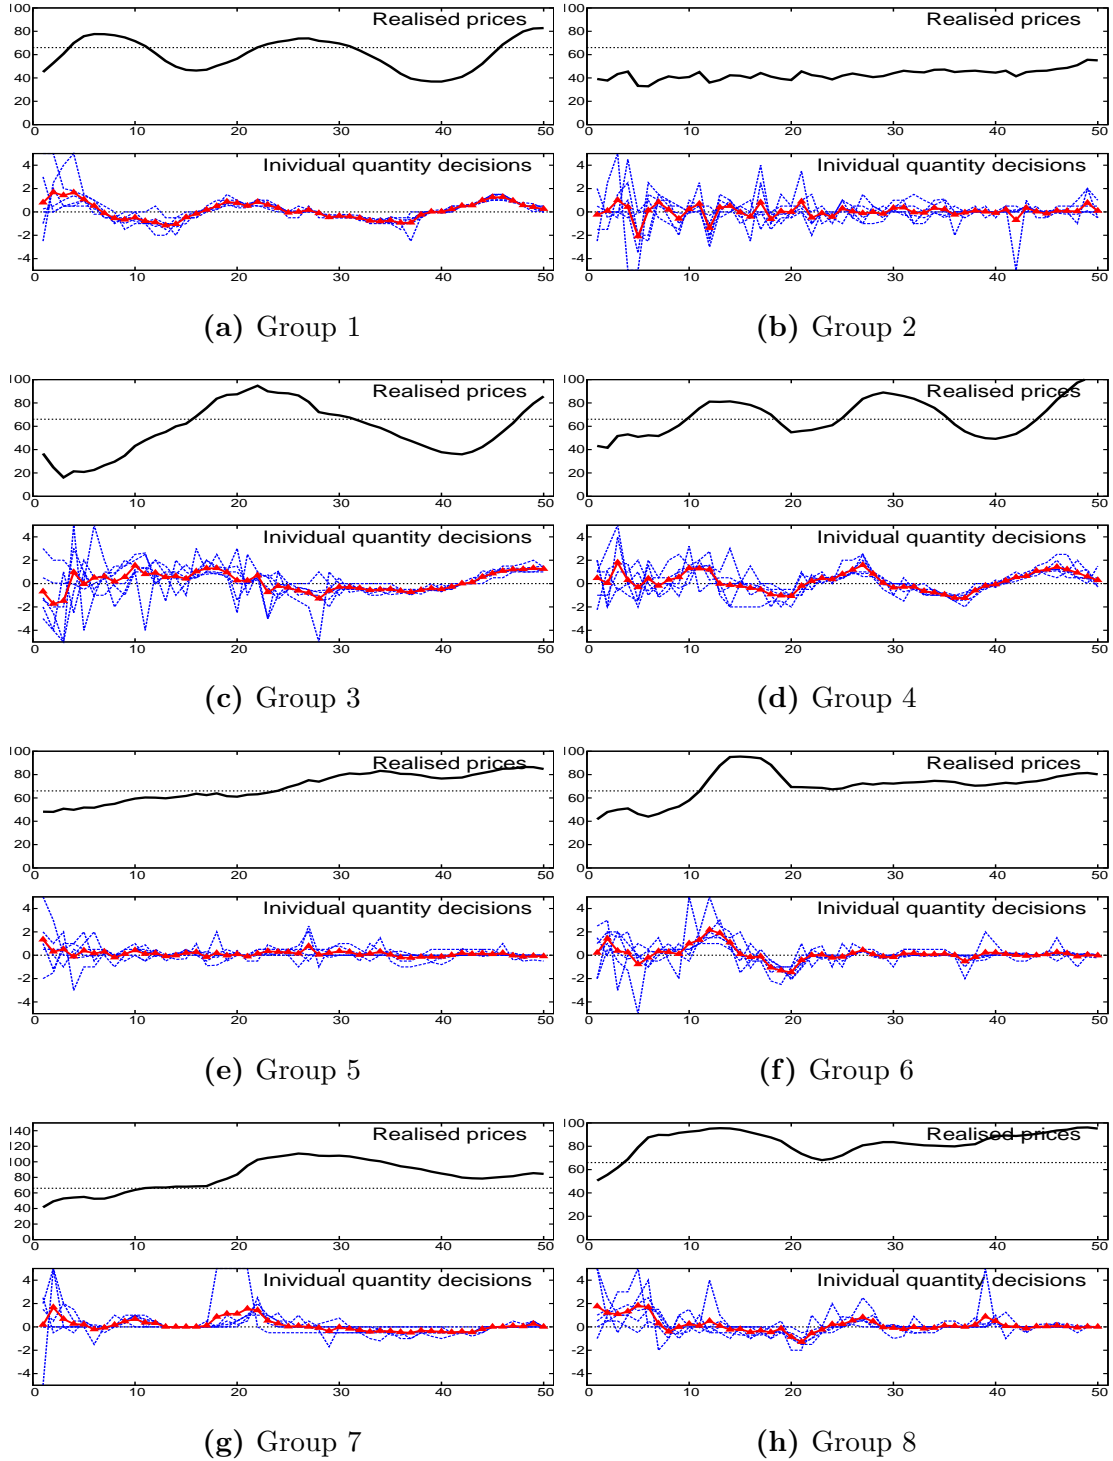
⁶We compare the dispersion of individual decisions using the standard deviation of the (implied) quantity decisions averaged over all periods in each market. A rank-sum test suggests that there is no difference between dispersion of quantity decisions in the LtO versus Mixed treatment (with p -value equal to 0.083 for dispersion over all periods and p -value equal to 0.161 for dispersion over last 40 periods). The dispersions of the quantity decisions in the LtO and Mixed treatments are indeed significantly larger than the dispersion of (implied) quantity decisions in the LtF treatment, with p -values equal to 0 for both all and last 40 periods.

Figure 1: *Price Dynamics in LtF Treatment*



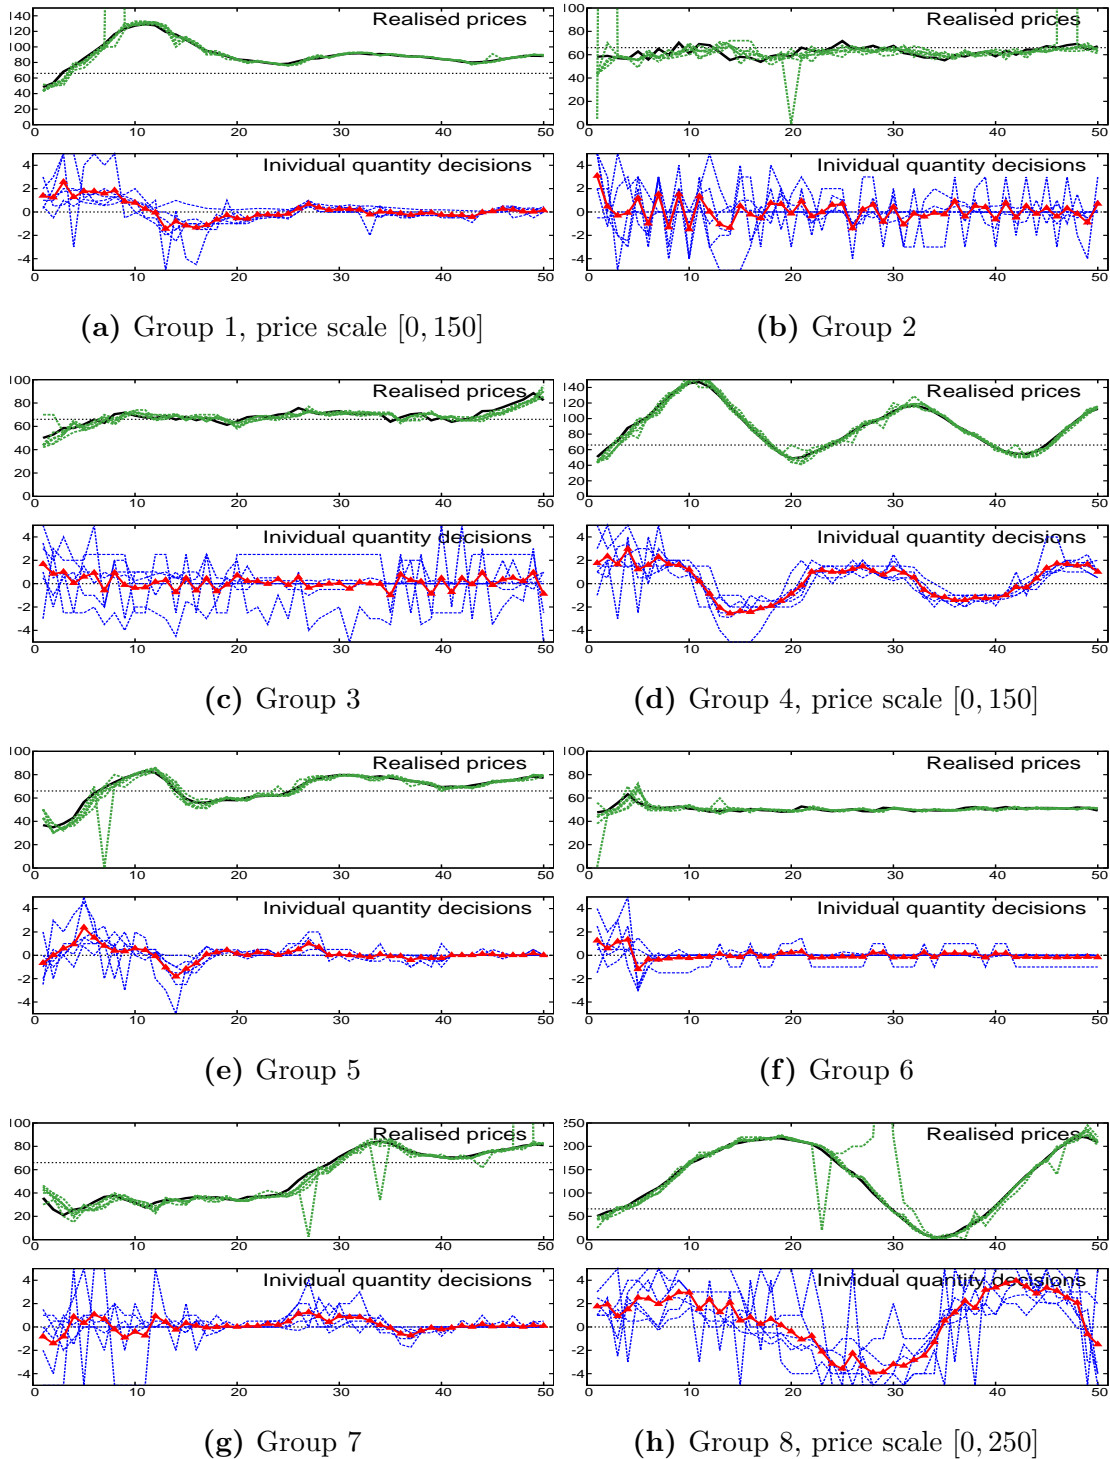
Notes. Groups 1-8 for the Learning to Forecast treatment. Straight line shows the fundamental price $p^f = 66$, solid black line denotes the realised price, while green dashed lines denote individual forecasts.

Figure 2: Price Dynamics in LtO Treatment



Notes. Groups 1-8 for the Learning to Optimise 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 3: Price Dynamics in Mixed Treatment



Notes. 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 realised 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).

2.2 Quantifying the Bubbles

The term ‘bubble’ is informally used in the literature to describe, loosely speaking, a prolonged spell of an asset price growth beyond its fundamental. In order to capture this notion with a rigid statistic, we follow Stöckl *et al.* (2010) and evaluate the experimental mispricing with 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

$$RAD_{g,t} \equiv \frac{|p_t^g - p^f|}{p^f} \times 100\%, \quad (13)$$

$$RD_{g,t} \equiv \frac{p_t^g - p^f}{p^f} \times 100\%, \quad (14)$$

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_g and RD_g are defined as

$$\overline{RAD}_g \equiv \frac{1}{50} \sum_{t=1}^{50} RAD_{g,t}, \quad (15)$$

$$\overline{RD}_g \equiv \frac{1}{50} \sum_{t=1}^{50} RD_{g,t}, \quad (16)$$

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

It is difficult to come up with a formal criterion for a bubble in terms of these measures. In particular, when bubbles are accompanied by a price plunge, or “negative bubbles”, the RD may be very close to 0. Therefore, in this paper we focus on the differences between the three treatments.⁷

The results for average \overline{RAD} and \overline{RD} measures for each treatment are presented in Table 1. They confirm that the LtF groups were the closest to, though

⁷As empirical benchmarks we computed these two measures for the US stock and housing markets. The RAD (RD) is 40% (20.2%) for S&P500, based on quarterly data 1950Q1 – 2012Q4 and the fundamental computed by a standard Gordon present discounted value model; for the same data set, using deviations from the Campbell-Cochrane consumption-habit fundamental model the RAD (RD) is 19% (3.9%) (Hommes and in’t Veld, 2014). For US housing market data in deviations from a benchmark fundamental based on housing rents the RAD (RD) are 7.7% (0.4%) for over 40 years of quarterly data 1970Q1 – 2013Q1 and 9.7% (2.2%) for 20 years of quarterly data 1993Q1 – 2013Q1 (Bolt *et al.*, 2014).

still quite far from, the REE (with an average \overline{RAD} of about 9.5%), while Mixed groups exhibited the largest price deviations with an average \overline{RAD} of 36%. Interestingly, LtO groups had significant oscillations (on average high \overline{RAD} of 24.7%), 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.

Table 1: *RAD and RD*

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

Notes. Relative Absolute Deviation (RAD) and Relative Deviation (RD) of the experimental prices for the three treatments, in percentages. *** (**) denotes groups for which the average RAD from the last 40 periods is larger than 3% on 1% (5%) significance level. * denotes groups for which the average RD from the last 40 periods is outside $[-1.5\%, 1.5\%]$ interval on 5% significance level.

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%.⁸ Furthermore, for all groups in all three treatments, t-test on any meaningful significance level rejects the null of the average price (for periods 11–50, i.e. the last 40 periods to allow for learning by the subjects) being equal to the fundamental value. This result shows negative evidence on **HYPOTHESIS 1**: none of the treatments converges to the REE.

There is no significant difference between the treatments in terms of \overline{RD} ac-

⁸3% \overline{RAD} is approximately 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% confidence bounds of the REE.

according to the Mann-Whitney-Wilcoxon rank-sum test (p -value > 0.1 for each pair of the treatments, z -statistic is -0.735 , -0.735 and -0.420 for LtF, LtO and Mixed respectively. The unit of observation is per market, i.e. 8 for each treatment). However, the difference between the LtF treatment and each of the other treatments in terms of \overline{RAD} is significant at 5% according to the rank-sum test (p -value = 0.002 and 0.003, and z -statistic is -3.151 and -2.205 for the LtO and Mixed respectively, number of observations: 8 for each treatment), while the difference between the LtO and Mixed is not significant (p -value = 0.753, z -statistic = -0.135 number of observations: 8 for each treatment). This is strong evidence against **HYPOTHESIS 2**, as it shows that trading and forecasting tasks yield different market dynamics.

The RAD values in our paper are similar to those in Stöckl *et al.* (2010) (see specifically their Table 4 for the $\overline{RAD}/\overline{RD}$ measures). Nevertheless, there are some important differences. First, group 8 from the Mixed treatment (with \overline{RAD} equal to 120.7%) exhibits the largest price bubble in the experiment. 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} .

In addition, our experiment yields measures resembling the above mentioned benchmark stock and housing markets (see footnote 7). Indeed, the LtF, LtO and Mixed experimental treatments yields boom/boost cycles of a realistic magnitude, comparable to what has been observed in recent stock and housing market bubbles and crashes.

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.*

2.3 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 against 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 D.1 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 (rank-sum test for difference in distributions with 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 (rank-sum test with p -value=0.753).

RESULT 2. *Forecasting efficiency is significantly higher in the LtF than in the Mixed treatment, while there is no significant difference in the trading efficiency in treatments LtO and Mixed. **HYPOTHESIS 3** is partially rejected.*

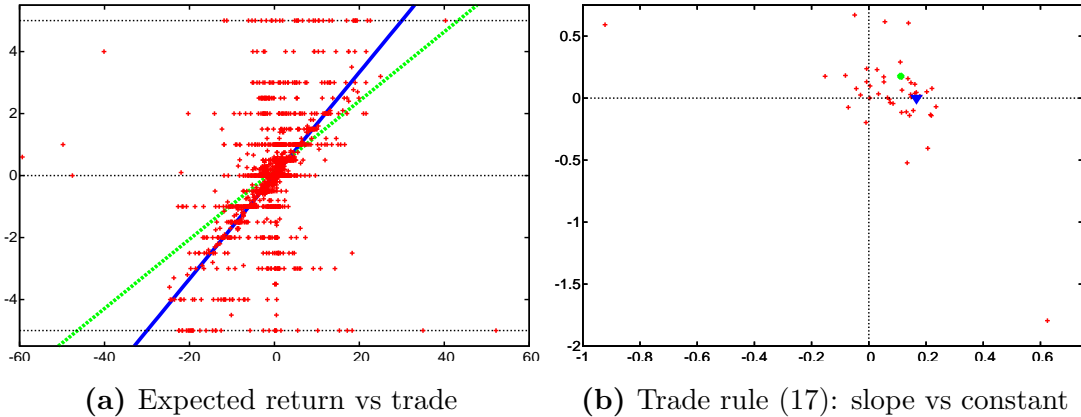
2.4 *Conditional Optimality of Forecast and Quantity Decision in Mixed Treatment*

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 Eq. (11) (the optimal quantity is $1/6$ of the corresponding expected asset return). Figure 4 shows the scatter plot of the quantity decision against the implied predicted return $\rho_{i,t+1}^e = p_{i,t+1}^e + 3.3 - 1.05p_t$, for each subject and each period separately.⁹ If all individuals made consistent decisions, these points should lie on the (blue) line with slope $1/6$.

Figure 4a illustrates 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 (bought) the asset even though they believed its return will be substantially positive (negative).

⁹Sometimes the subjects submit extremely high price predictions, which in most cases seem to be typos. The scatter plot excludes these outliers, by restricting the horizontal scale of predicted returns on the asset between -60 and 60 .

Figure 4: *Conditional Optimality of Quantity Decisions*



Notes. ML estimation for trading rule (17) 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 the estimated trading rules (17) 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 ($z_{i,t} = \rho_{i,t}^e/6$). Dashed line (left panel)/circle (right panel) denotes the estimated rule under restriction of homogeneity ($z_{i,t} = c + \theta\rho_{i,t}^e$).

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

$$z_{i,t} = c_i + \theta_i\rho_{i,t+1}^e + \eta_{i,t}, \quad (17)$$

with $\eta_{i,t} \sim NID(0, \sigma_{\eta,i}^2)$. The estimated coefficients for all subjects are shown in the scatter plot of Figure 4b. 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) expected return on the asset. It has two important special cases: homogeneity and optimality (nested within homogeneity). To be specific, subject homogeneity (heterogeneity) corresponds to an insignificant (significant) variation in the slope $\theta_i = \theta_j$ ($\theta_i \neq \theta_j$) for any (some) pair of subjects i and j . The constant c_i shows the ‘irrational’ optimism/pessimism bias of subject i . Optimality of individual quantity decisions implies homogeneity with the additional restrictions that $\theta_i = \theta_j = 1/6$ and $c_i = c_j = 0$ (no agent has a decision bias).

The assumptions of homogeneity and perfect optimisation are tested by estimation of equation (17) with the restrictions on the parameters c_i and θ_i .¹⁰ These

¹⁰We use ML since the optimality constraint does not exclude heterogeneity of the idiosyncratic shocks $\eta_{i,t}$. We exclude outliers defined as observations when a subject predicts an asset return

regressions are compared with an unrestricted regression (with $\theta_i \neq \theta_j$ and $c_i \neq c_j$) *via* a Likelihood Ratio (LR) test. The result of the LR test shows that both the assumption of homogeneity and perfect optimisation are *rejected* (with p -values below 0.001). Furthermore, we explicitly tested for $z_{i,t} = \rho_{i,t}^e/6$ when estimating individual rules. Estimations identified 11 subjects (23%) as consistent optimal traders (see footnote 13 for a detailed discussion). In sum, we find evidence for heterogeneity of individual trading rules. The majority of the subjects are unable to learn the optimal solution.

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 understanding of the structure of the economy. But the second part of RE is often taken as a good approximation: agents are assumed to 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 in forecasting as well as choosing quantities. The simplest explanation is that individuals in general lack the computational capacity to make perfect mathematical optimisations.

RESULT 3. *The subjects' quantity decisions are not conditionally optimal given their price forecasts in the Mixed treatment. We conclude that **HYPOTHESIS 4** is rejected for 77% (37 out of 48) of the subjects.*

2.5 Estimation of Individual Behavioural Rules

In this subsection we estimate individual forecasting and trading rules and investigate whether there are significant differences between treatments. Prior experimental work (Heemeijer *et al.*, 2009) suggests that in LtF experiments, subjects use heterogeneous forecasting rules which nevertheless typically are well described by a simple linear First-Order Rule

$$p_{i,t}^e = \alpha_i p_{t-1} + \beta_i p_{i,t-1}^e + \gamma_i (p_{t-1} - p_{t-2}). \quad (18)$$

higher than 60 in absolute terms. To account for an initial learning phase, 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.

This rule may be viewed as an *anchor and adjustment rule* (Tversky and Kahneman, 1974), as it extrapolates a price change (the last term) from an anchor (the first two terms). Two important special cases of (18) are the pure **trend following** rule with $\alpha_i = 1$ and $\beta_i = 0$, yielding

$$p_{i,t}^e = p_{t-1} + \gamma_i(p_{t-1} - p_{t-2}), \quad (19)$$

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

$$p_{i,t}^e = \alpha_i p_{t-1} + (1 - \alpha_i) p_{i,t-1}^e. \quad (20)$$

The pure trend-following rule (19) uses an anchor giving all weight to the last observed price (p_{t-1}), while in the general rule (18) the anchor gives weight to the last observed price (p_{t-1}) as well as the last forecast ($p_{i,t-1}^e$). In this sense the general rule (18) is more cautious and extrapolates the trend from a more gradually evolving anchor, while the pure trend-following rule is more aggressive extrapolating the trend from the last price observation.

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

$$z_{i,t} = c_i + \chi_i z_{i,t-1} + \phi_i \rho_t, \quad (\text{LtO}) \quad (21a)$$

$$z_{i,t} = c_i + \chi_i z_{i,t-1} + \phi_i \rho_t + \zeta_i \rho_{i,t+1}^e. \quad (\text{Mixed}) \quad (21b)$$

This rule captures the most relevant and most recent possible elements of individual trading. Notice however, that the trading rule (21a) in the LtO treatment only contains a past return (ρ_t) term, while the trading rule (21b) in the Mixed treatment contains an additional term for expected excess return ($\rho_{i,t+1}^e$), which is not observable in the LtO treatment because subjects did not give price forecasts. Both trading rules have two interesting special cases. First, what we call **persistent demand** ($\phi_i = \zeta_i = 0$) is characterised by a simple AR(1) process:

$$z_{i,t} = c_i + \chi_i z_{i,t-1}. \quad (22)$$

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

$$z_{i,t} = c_i + \phi_i \rho_t \quad (\text{LtO}), \quad (23a)$$

$$z_{i,t} = c_i + \phi_i \rho_t + \zeta_i \rho_{i,t+1}^e \quad (\text{Mixed}). \quad (23b)$$

For the LtF and LtO treatments, for each subject we estimate her behavioural heuristic starting with the general forecasting rule (18) or the general trading

rule (21a) respectively. To allow for learning, all estimations are based on the last 40 periods. Testing for special cases of the estimated rules is straightforward: insignificant variables are dropped until all the remaining coefficients are significant at 5% level.¹¹

A similar approach is used for the Mixed treatment (now also allowing for the expected return coefficient ζ_i).¹² Equations (18) and (21b) are estimated simultaneously. One potential concern for the estimation is that the contemporary idiosyncratic errors in these two equations are correlated, given that 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 (18) can be estimated independently in the first step. The potential endogeneity only affects the trading heuristic (23b), and can be solved with a simple instrumental variable approach. The first step is to estimate the forecasting rule (18), which yields *fitted* price forecasts of each subject. In the second step, the trading rule (23b) is estimated with both the *fitted* forecasts as instruments, and directly with the *reported* forecasts. Endogeneity can be tested by comparing the two estimators using the Hausman test. Finally, the special cases of (21a–21b) are tested based on *reported* or *fitted* price forecasts according to the Hausman test.¹³

The estimation results can be found in Appendix E, in Tables E.1, E.2 and E.3 respectively for the LtF, LtO and Mixed treatments. 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.

2.5.1 Forecasting rules in LtF versus Mixed

The LtF treatment can be directly compared to the Mixed treatment by comparing the estimated forecasting rules (18). We observe that rules with a trend extrapolation term γ_i are popular in both treatments (respectively 39 in LtF and 25 in Mixed out of 48). A few other subjects use a pure adaptive rule (20) (none in LtF and 3 in Mixed treatments respectively). A few others use a rule defined by

¹¹Adaptive expectations (20) 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 (20), 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 (18).

¹²See footnote 9.

¹³Whenever the estimations indicated that a subject from the Mixed treatment used a return extrapolation rule of the form $z_{i,t} = \zeta_i \rho_{i,t+1}^e$, that is a rule in which only the implied expected return was significant, we directly tested $\zeta_i = 1/6$. This restriction implies optimal trading consistently with the price forecast, which we could not reject for 11 out of 48 subjects.

(18) where $\gamma_i = 0$, but $\alpha_i + \beta_i \neq 1$. There were no subjects in the LtF treatment and only 2 in the Mixed treatment, for whom we could not identify a significant forecasting rule. The average trend coefficients in both treatments are close to $\bar{\gamma} \approx 0.4$, and not significantly different in terms of distribution (with p -value of the rank-sum test equal to 0.736). The difference between the two treatments lies in the anchor of the forecasting rule. For the LtF treatment the average coefficients are $\bar{\alpha} = 0.45$ and $\bar{\beta} = 0.56$, while in the Mixed treatment these are $\bar{\alpha} = 0.84$ and $\bar{\beta} = 0.06$ (the differences are significant according to the rank-sum test, with both p -values close to zero). This suggests that subjects in the LtF treatment are more cautious in revising their expectations, with a gradually evolving anchor that puts equal weight on past price and their previous forecast. In contrast, in the Mixed treatments subjects use an anchor that puts almost all weight on the last price observation and are thus closer to using a pure trend-following rule extrapolating a trend from the last price observation.

2.5.2 Trading rules in LtO versus Mixed

The LtO and Mixed treatments can be compared by the estimated trading rules. Recall however, that the trading rule (21a) in the LtO treatment only contains a past return term (ρ_t) with coefficient ϕ_i , while the trading rule (21b) in the Mixed treatment contains an additional term for expected excess return (ρ_{t+1}^e) with a coefficient ζ_i . In both treatments we find that the rules with a term on past or expected return is the dominating rule (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 in both the LtO and Mixed treatments the majority of subjects tried to extrapolate realized and/or expected asset returns, which leads to relatively strong trend chasing behaviour. Nevertheless, there are 11 subjects in the LtO treatment and 8 in the Mixed treatment for whom we can not identify a trading rule within this simple class. The average demand persistence was $\bar{\chi} = 0.07$ and $\bar{\chi} = 0.006$, and the average trend extrapolation was $\bar{\phi} = 0.09$ and $\overline{\phi + \zeta} = 0.06$ in the LtO and Mixed treatment respectively.¹⁴ The distributions of the two coefficients are not significantly different across the treatments according to the rank-sum test, with p -values of 0.425 and

¹⁴The trading rules (21a) and (21b) are not directly comparable, since (21b) is a function of both the past and the expected asset return, and the latter is unobservable in the LtO treatment. For the sake of comparability, 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. As a robustness check, we also estimated the simplest trading rule (21a) for both the LtO and Mixed treatments (ignoring expected asset returns) and found no significant difference between treatments.

0.885 for χ_i and $\phi_i/\phi_i + \zeta_i$ respectively. Hence, based upon individual trading rules we do not find significant differences between the LtO and Mixed treatments.

2.5.3 Implied trading rules in LtF versus LtO

It is more difficult to compare the LtF and LtO treatments based upon individual decision rules, since there was no trading in the LtF and no forecasting in the LtO treatment. We can however use the estimated individual forecasting rules to obtain the implied optimal trading rules (8) in the LtF treatment and compare these to the general trading rule (21a) in the LtO treatment. A straightforward computation shows that for a forecasting rule (18) with coefficients $(\alpha_i, \beta_i, \gamma_i)$, the implied optimal trading rule has coefficients $\chi_i = \beta_i$ and $\phi_i = (\alpha_i + \gamma_i - R)/6$.¹⁵ Hence, for the LtF and LtO treatments we can compare the coefficients for the adaptive term, i.e. the weight given to the last trade, and the return extrapolation coefficients. The averages of the first coefficient are $\bar{\beta} = 0.56$ and $\bar{\chi} = 0.07$ for the LtF and LtO treatments respectively, and it is significantly higher in the LtF treatment (rank-sum test p -value close to zero). Moreover the second coefficient, the implied reaction to the past 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 significant (rank-sum test p -value close to zero). Hence, these results on the individual (implied) trading rules show differences between the LtO and LtF treatments. The LtO treatment is more unstable than the LtF treatment because subjects are less cautious in the sense that they give less weight to their previous trade and they give more weight to extrapolating past returns.

We summarise the results on estimated individual behavioural rules as follows:

RESULT 4. *Most subjects, regardless of the treatment, follow an anchor and adjustment rule. In forecasting, LtF subjects were more cautious, using an anchor that puts more weight on their previous forecast, while the Mixed treatments subjects use an anchor with almost all weight on recent prices. In trading, most subjects extrapolate past returns and/or expected returns. In the LtO subjects give more weight to past return extrapolation compared to the implied trading behaviour in the LtF. These individual rules explain more unstable aggregate dynamics in the LtO and Mixed treatments. We conclude that **HYPOTHESIS 5** is rejected.*

¹⁵The implied trading rule (8) however cannot exactly be rewritten in the form (21a), but has one additional term p_{t-1} with coefficient $[R(\beta_i + \alpha_i + \gamma_i - R) - \gamma_i]/(a\sigma^2)$. This coefficient typically is small however, since γ_i is small and $\alpha_i + \beta_i$ is close to 1. The mean estimated coefficient over 48 subjects is very close to zero (-0.00229), and with a simple t-test we can not reject the hypothesis that the mean coefficient is 0 (p -value 0.15).

3 Conclusions

The origin of asset price bubbles is an important topic for both researchers and policy makers. This paper investigates the price dynamics and bubble formation in an experimental asset pricing market with a price adjustment rule. We find that the mispricing is largest in the treatment where subjects do both forecasting and trading, and smallest when subjects only make a prediction. Our result suggests that price instability is the result of both inaccurate forecasting and imperfect optimisation. There has been empirical work quantifying forecast biases by households and finance professionals in real markets, and theoretical works that start to incorporate the stylized facts into modelling of expectations in macroeconomics. Our result suggests it may be equally important to collect evidence on failure in making optimal decisions conditional on one's own belief by market participants, and incorporate this behavioural bias into modelling of simple heuristics as an alternative to perfectly optimal individual decisions.

Which behavioural biases can explain the differences in the individual decisions and aggregate market outcomes in the learning to forecast and learning to optimise markets? A first possibility is that the quantity decision task is more cognitive demanding than the forecasting task, when the subjects in the LtF treatment are assisted 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 that while subjects take significantly longer time in the Mixed treatment than the other two treatments according to Mann-Whitney-Wilcoxon test, 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. Second, in the LtF treatment, the subjects' goal is to find an 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 a subject predicts 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 price changes or follow the "wisdom of crowd", which leads to assigning more weight to past or expected returns. Furthermore, for price forecasting past individual behaviour can be directly compared to observed market prices. If an individual forecasting strategy fits well with observed price behaviour, more weight may be given to past own

individual behaviour. In contrast, individual trading decisions cannot be directly compared or anchored against trading volume or other aggregate information. It then becomes more natural to evaluate and anchor individual trading by giving more weight to recent past prices and/or recent past returns.

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-economic crisis. Proponents of rational expectations often claim that serious asset bubbles cannot arise, because rational economic agents would efficiently arbitrage against it and quickly push the ‘irrational’ (non-fundamental) investors out of the market. Our experiment suggests otherwise: people exhibit heterogeneous and not necessarily optimal behaviour. Because they are trend-followers, their non-fundamental beliefs are correlated. This is reinforced by the *positive feedback* between expectations and realised prices in asset markets, as stressed e.g. in Hommes (2013). Therefore, price oscillations cannot be mitigated by more rational market investors. As a result, waves of optimism and pessimism can arise despite the fundamentals being relatively stable.

Our experiment can be extended in several ways. For example, the subjects in our experiment can short-sell the assets, which may not be feasible in real markets. An interesting topic for future research is the case where agents face short selling constraints (Anufriev and Tuinstra, 2013). 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.* (2014) that study how network structure influences market efficiency when subjects act as 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.

University of Amsterdam and Tinbergen Institute
University of Groningen

References

- Akiyama, E., Hanaki, N. and Ishikawa, R. (2012). ‘Effect of uncertainty about others’ rationality in experimental asset markets: An experimental analysis’, *Available at SSRN 2178291*.
- Anufriev, M. and Tuinstra, J. (2013). ‘The impact of short-selling constraints on financial market stability in a heterogeneous agents model’, *Journal of Economic Dynamics and Control*, vol. 31, pp. 1523–1543.
- Assenza, T., Bao, T., Hommes, C. and Massaro, D. (2014). ‘Experiments on expectations in macroeconomics and finance’, in (J. Duffy, ed.), *Experiments in Macroeconomics*, pp. 11–70, vol. 17 of *Research in Experimental Economics*, Emerald Publishing.
- Bao, T., Duffy, J. and Hommes, C. (2013). ‘Learning, forecasting and optimizing: An experimental study’, *European Economic Review*, vol. 61, pp. 186–204.
- Bao, T., Hommes, C., Sonnemans, J. and Tuinstra, J. (2012). ‘Individual expectations, limited rationality and aggregate outcomes’, *Journal of Economic Dynamics and Control*, vol. 36(8), pp. 1101–1120.
- Beja, A. and Goldman, M.B. (1980). ‘On the dynamic behavior of prices in disequilibrium’, *Journal of Finance*, vol. 35(2), pp. 235–248.
- Bolt, W., Demertzis, M., Diks, C.G.H., Hommes, C.H. and van der Leij, M. (2014). ‘Identifying booms and busts in house prices under heterogeneous expectations’, De Nederlandsche Bank Working Paper No. 450.
- Bondt, W. and Thaler, R. (2012). ‘Further evidence on investor overreaction and stock market seasonality’, *Journal of Finance*, vol. 42(3), pp. 557–581.
- Bostian, A.A. and Holt, C.A. (2009). ‘Price bubbles with discounting: A web-based classroom experiment’, *Journal of Economic Education*, vol. 40(1), pp. 27–37.
- Brock, W.A. and Hommes, C.H. (1998). ‘Heterogeneous beliefs and routes to chaos in a simple asset pricing model’, *Journal of Economic Dynamics and Control*, vol. 22(8-9), pp. 1235–1274.
- Campbell, J. and Shiller, R. (1989). ‘Stock prices, earnings and expected dividends’, NBER working paper.

- Campbell, J.Y., Lo, A.W. and MacKinlay, A.C. (1997). *The Econometrics of Financial Markets*, Princeton University Press.
- Chiarella, C., Dieci, R. and He, X.Z. (2009). ‘Heterogeneity, market mechanisms, and asset price dynamics’, *Handbook of financial markets: Dynamics and evolution*, vol. 231, pp. 277–344.
- Choi, S., Galeotti, A. and Goyal, S. (2014). ‘Trading in networks: theory and experiments’, Cambridge-INET Working Paper 8.
- Duffy, J. (2008). ‘Macroeconomics: a survey of laboratory research’, *Handbook of Experimental Economics*, vol. 2.
- Dufwenberg, M., Lindqvist, T. and Moore, E. (2005). ‘Bubbles and experience: An experiment’, *American Economic Review*, vol. 95(5), pp. 1731–1737.
- Fama, E.F. (1970). ‘Efficient capital markets: A review of theory and empirical work*’, *Journal of Finance*, vol. 25(2), pp. 383–417.
- Farmer, J. and Lo, A. (1999). ‘Frontiers of finance: Evolution and efficient markets’, *Proceedings of the National Academy of Sciences*, vol. 96(18), pp. 9991–9992.
- Füllbrunn, S., Rau, H.A. and Weitzel, U. (2014). ‘Does ambiguity aversion survive in experimental asset markets?’, *Journal of Economic Behavior & Organization*, vol. 107, pp. 810–826.
- Gale, D.M. and Kariv, S. (2007). ‘Financial networks’, *American Economic Review*, vol. 97(2), pp. 99–103.
- Gale, D.M. and Kariv, S. (2009). ‘Trading in networks: A normal form game experiment’, *American Economic Journal: Microeconomics*, vol. 1(2), pp. 114–132.
- Galeotti, A., Goyal, S., Jackson, M.O., Vega-Redondo, F. and Yariv, L. (2010). ‘Network games’, *Review of Economic Studies*, vol. 77(1), pp. 218–244.
- Haruvy, E. and Noussair, C. (2006). ‘The effect of short selling on bubbles and crashes in experimental spot asset markets’, *Journal of Finance*, vol. 61(3), pp. 1119–1157.
- Haruvy, E., Noussair, C. and Powell, O. (2013). ‘The impact of asset repurchases and issues in an experimental market’, *Review of Finance*, vol. 18(2), pp. 1–33.

- Heemeijer, P., Hommes, C., Sonnemans, J. and Tuinstra, J. (2009). ‘Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation’, *Journal of Economic Dynamics and Control*, vol. 33(5), pp. 1052–1072.
- Hommes, C. (2011). ‘The heterogeneous expectations hypothesis: Some evidence from the lab’, *Journal of Economic Dynamics and Control*, vol. 35(1), pp. 1–24.
- Hommes, C. (2013). *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems*, Cambridge University Press.
- Hommes, C. and in’t Veld, D. (2014). ‘Booms, busts and behavioural heterogeneity in stock prices’, CeNDEF Working paper 14-14 University of Amsterdam.
- Hommes, C., Sonnemans, J., Tuinstra, J. and Van de Velden, H. (2005). ‘Coordination of expectations in asset pricing experiments’, *Review of Financial Studies*, vol. 18(3), pp. 955–980.
- Huber, J. and Kirchler, M. (2012). ‘The impact of instructions and procedure on reducing confusion and bubbles in experimental asset markets’, *Experimental Economics*, vol. 15(1), pp. 89–105.
- Kirchler, M., Huber, J. and Stöckl, T. (2012). ‘Thar she bursts: Reducing confusion reduces bubbles’, *American Economic Review*, vol. 102(2), pp. 865–883.
- LeBaron, B. (2006). ‘Agent-based computational finance’, *Handbook of Computational Economics*, vol. 2, pp. 1187–1233.
- Lei, V., Noussair, C.N. and Plott, C.R. (2001). ‘Nonspeculative bubbles in experimental asset markets: Lack of common knowledge of rationality vs. actual irrationality’, *Econometrica*, vol. 69(4), pp. 831–859.
- Lucas Jr., R.E. (1972). ‘Expectations and the neutrality of money’, *Journal of Economic Theory*, vol. 4(2), pp. 103–124.
- Marimon, R., Spear, S.E. and Sunder, S. (1993). ‘Expectationally driven market volatility: an experimental study’, *Journal of Economic Theory*, vol. 61(1), pp. 74–103.
- Muth, J.F. (1961). ‘Rational expectations and the theory of price movements’, *Econometrica*, vol. 29(3), pp. 315–335.

- Noussair, C., Robin, S. and Ruffieux, B. (2001). ‘Price bubbles in laboratory asset markets with constant fundamental values’, *Experimental Economics*, vol. 4(1), pp. 87–105.
- Noussair, C.N. and Tucker, S. (2013). ‘Experimental research on asset pricing’, *Journal of Economic Surveys*, vol. 27(3), pp. 554–569.
- Shiller, R. (2003). ‘From efficient markets theory to behavioral finance’, *Journal of Economic Perspectives*, vol. 17(1), pp. 83–104.
- Smith, V., Suchanek, G. and Williams, A. (1988). ‘Bubbles, crashes, and endogenous expectations in experimental spot asset markets’, *Econometrica: Journal of the Econometric Society*, vol. 56(5), pp. 1119–1151.
- Stöckl, T., Huber, J. and Kirchler, M. (2010). ‘Bubble measures in experimental asset markets’, *Experimental Economics*, vol. 13(3), pp. 284–298.
- Sunder, S. (1995). ‘Experimental asset markets: A survey.’, in (J. Kagel and A. Roth, eds.), *Handbook of Experimental Economics*, pp. 445–500, Princeton: NJ: Princeton University Press.
- Tversky, A. and Kahneman, D. (1974). ‘Judgment under uncertainty: Heuristics and biases’, *Science*, vol. 185(4157), pp. 1124–1131.

A Experimental Instructions

(Not For Publication)

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.

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 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 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.

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:

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

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 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.

B Payoff Tables

(Not For Publication)

Table B.1: *Payoff Table for Forecasting Task*

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> \geq 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 B.2: *Payoff Table for Trading Task*

		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.

C Rational Strategic Behaviour

(Not For Publication)

Our experimental results are clearly different from the predictions of the rational expectation equilibrium (REE). However, we also found that subjects typically earn high payoffs, implying some sort of profit seeking behaviour.

In this appendix, we discuss whether rational strategic behaviour can explain our experimental results, in particular in the LtO and Mixed treatments. Based on the assumption on the subjects' perception of the game and information structure, 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 non-cooperatively. We show that under price-taking behaviour, the LtF and LtO treatments are equivalent. If the subjects behave strategically or try to collude, the economy can have alternative 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, we focus on the one-shot game version of the experimental market to derive our results. More precisely, we look at the optimal quantity decisions $z_{i,t}$ that the agents in period t (knowing prices and individual traded quantity until and including period t) have to formulate to maximise the expected utility in period $t + 1$. This is supported by two observations. 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.

C.1 Price takers

Realised utility of investors in the LtO treatment is given by (4) and is equivalent to the following form:

$$U_{i,t}(z_{i,t}) = z_{i,t}(p_{t+1} + y - Rp_t) - \frac{a\sigma^2}{2}z_{i,t}^2, \quad (\text{C.1})$$

where $z_{i,t}$ is the traded quantity and $U_{i,t}$ is a quadratic function of the traded quantity. As discussed in Section 2, assuming the agent is a price taker, the

optimal traded quantity conditional on the expected price $p_{i,t+1}^e$ is given by

$$z_{i,t}^* = \arg \max_{z_{i,t}} U_{i,t} = \frac{p_{i,t+1}^e + y - Rp_t}{a\sigma^2}. \quad (\text{C.2})$$

Note that this result relies on the assumption that the subjects do not know the price generating function. We argue that the subjects also have an incentive to minimise their forecasting error when they choose the quantity and are paid according to the risk adjusted profit. 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:

$$\begin{aligned} U_{i,t}(z_{i,t}) &= z_{i,t}(p_{t+1} + y - Rp_t) - \frac{a\sigma^2}{2} z_{i,t}^2 \\ &= \frac{(p_{t+1} + \epsilon + y - Rp_t)(p_{t+1} + y - Rp_t)}{a\sigma^2} - \frac{(p_{t+1} + \epsilon + y - Rp_t)^2}{2a\sigma^2} \\ &= \frac{(p_{t+1} + y - Rp_t)^2}{2a\sigma^2} - \frac{\epsilon^2}{2a\sigma^2}. \end{aligned} \quad (\text{C.3})$$

This shows that utility is maximised when $\epsilon = 0$, namely, when all subjects have correct belief. 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. Subjects have incentives to search for and play the REE also when they choose the quantity. We summarise this finding below.

FINDING 1. *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. The subjects' payoff is maximised when they play the Rational Expectation Equilibrium regardless of the design: the REE of LtF and LtO treatments are equivalent.*

C.2 Collusive outcome

Consider now the case when agents realise how their predictions/trading quantities influence the price and are able to coordinate on a common strategy. This resembles a collusive (oligopoly) market, *e.g.* similar to a cobweb economy in which the sellers can coordinate their production.

In the collusive case, 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:

$$p_{t+1} = p_t + 6\lambda z_t, \quad (\text{C.4})$$

and so the monopoly maximises

$$\begin{aligned}
U_t &= \sum_{i=1}^6 U_{i,t}(z_t) = 6 \left[z_t(p_{t+1} + y - Rp_t) - \frac{a\sigma^2}{2} z_t^2 \right] \\
&= 6 \left[z_t^2 \left(6\lambda - \frac{a\sigma^2}{2} \right) + z_t(y - rp_t) \right]. \tag{C.5}
\end{aligned}$$

Here we assume that a rational agent has perfect knowledge about the pricing function (C.4). Notice that when $\lambda = 20/21$, $a\sigma^2 = 6$, as in the experiment, the coefficient before z_t^2 is positive, $6\lambda - \frac{a\sigma^2}{2} = \frac{19}{7} > 0$, and thus the profit function is U shaped, instead of inversely U shaped.¹⁶ 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 (C.5) 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 initial price is below (above) the fundamental, the monopoly will buy (sell) the asset until the price overshoots (undershoots) the fundamental, and so forth. Then the subjects start to ‘jump up and down’ as described before.

FINDING 2. *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.*

¹⁶If $6\lambda - \frac{a\sigma^2}{2} < 0$, this objective function is inversely U shaped. The maximum point is achieved when $z_{i,t} = \frac{rp_t - y}{12\lambda - a\sigma^2}$. This means when $p_t = y/r$, 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.

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. Moreover, quantity decisions equal to 5 or -5 happened rarely in the experiment (7 times in the LtO and 44 times in the Mixed treatment). Typical subject behaviour was much more conservative: 97% and 91% traded quantities in the LtO and Mixed treatments respectively were confined in the interval $[-2.5, 2.5]$.

C.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, while the non-cooperative NE does not. What will rational subjects do in this situation?

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

$$p_{t+1} = p_t + 5\lambda z_t^o + \lambda z_{i,t}. \quad (\text{C.6})$$

Agent i maximises therefore

$$\begin{aligned} U_{i,t} &= z_{i,t} (\lambda z_{i,t} + 5\lambda z_t^o + y - rp_t) - \frac{a\sigma^2}{2} z_{i,t}^2 \\ &= z_{i,t}^2 \frac{2\lambda - a\sigma^2}{2} + z_{i,t} (5\lambda z_t^o + y - rp_t). \end{aligned} \quad (\text{C.7})$$

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

$$z_{i,t}^*(z_t^o) = \frac{5\lambda z_t^o + y - rp_t}{a\sigma^2 - 2\lambda}. \quad (\text{C.8})$$

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

$$z_t^* = \frac{rp_t - y}{7\lambda - a\sigma^2} = \frac{3}{2}(rp_t - y). \quad (\text{C.9})$$

Furthermore the reaction function $z_{i,t}^*(z_t^o)$ is linear with respect to z_t^o , with a slope $\frac{5\lambda}{a\sigma^2 - 2\lambda} = \frac{100}{86} > 1$. 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 (C.9) thus defines the threshold point between the two corner strategies, *i.e.* the full NE best response strategy is defined as

$$z_{i,t}^{NE} = \begin{cases} 5 & \text{if } z_t^o > z_t^* \\ z_t^* & \text{if } z_t^o = z_t^* \\ -5 & \text{if } z_t^o < z_t^*. \end{cases} \quad (\text{C.10})$$

The boundary strategies can be infeasible if the previous price is too close to zero or 300.¹⁷ 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 as fixed points of (C.10), namely all players playing $z_{i,t} = -5$, $z_{i,t} = z_t^*$ and $z_{i,t} = +5$ for all $i \in \{1, \dots, 6\}$.

A simple interpretation is that, given the parametrization, our model is an example of a (Nash) coordination game. As long as $\frac{5\lambda}{a\sigma^2 - 2\lambda} > 1$, the best response (C.8) is to amplify the average trade of the other players. This is not a surprising result, as it merely exhibits the strength of the positive feedback present in this market¹⁸.

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

$$p_{t+1} = \frac{10p_t - 60y}{7}. \quad (\text{C.11})$$

In contrast to the collusive game, in the non-cooperative game the fundamental price is 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 $(p_{t-1} + y + 6\lambda z_{i,t} - (1+r)p_{t-1})z_{i,t} - \frac{\alpha\sigma^2}{2}z_{i,t}^2 = (33.3 - 0.05p_{t-1}) * 5 - 75$. This equals 76.5 when $p_{t-1} = 60$, 16.5 when $p_{t-1} = 300$ and 75 when the previous price is equal to the fundamental, $p_{t-1} = 66$. This explains why the payoff efficiency (average

¹⁷Notice 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 (C.10) remains intact. This implies that NE has to be symmetric.

¹⁸In practice, such an equilibrium could not be sustained in the long run, since then the market maker would incur accumulating losses every period.

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 (C.11) 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 will go to the upper cap of 300 or the lower cap of 0 respectively. Furthermore the agents can switch *at any moment* between the three one-shot game NE defined by (C.10). This implies that in the repeated non-cooperative game, many rational price paths are possible. This includes many price paths where agents coordinate on 5 or -5 , including the alternating collusive equilibrium discussed in the last section.

FINDING 3. *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.*

C.4 Summary

To conclude, the perfectly rational agents can coordinate on price boom-bust cycles and earn positive profit¹⁹. However, this would require even stronger assumptions than the fundamental equilibrium, namely that the agents perfectly understand the underlying price determination mechanism.

Furthermore, such rational equilibria with price oscillations predict that the subjects coordinate on *homogeneous* trades at the edge of the liquidity constraints. The subjects from the LtO and Mixed treatments behaved differently. Their traded quantities were highly heterogeneous, and rarely reached the liquidity constraints.

Therefore, the alternative rational equilibrium 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

¹⁹Note that the subjects earn more in collusive and non-cooperative Nash setting because we pay them according to the book value of the asset, and the tâtonnement process ensures the price movement is relatively smooth. In real life, people may not be able to realize the full book value of their asset holdings because the asset price will fall when a large fraction of them start to sell, and without the marker maker in the tâtonnements process absorbing all these losses, they may suffer huge losses when the asset price declines sharply.

cannot explain the heterogeneity of the individual decisions and non-boundary trading quantities, the mispricing in the experimental data is more likely 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.

D Earnings Ratios (Not For Publication)

Table D.1: *Earnings Efficiency*

Treatment	LtF	LtO	Mixed Forecasting	Mixed Trading
Market 1	96.35%	102.54%	87.62%	100.89%
Market 2	94.47%	95.25%	67.27%	87.33%
Market 3	96.03%	98.21%	75.63%	79.61%
Market 4	96.18%	100.43%	77.41%	114.63%
Market 5	95.15%	97.39%	87.07%	99.03%
Market 6	94.06%	99.64%	91.94%	97.24%
Market 7	96.18%	98.58%	81.20%	94.55%
Market 8	96.54%	98.41%	60.80%	132.01%
Average	95.62%	98.81%	78.62%	100.66%

Notes. Earnings efficiency for each market. The efficiency is defined as the average experimental payoff divided by the payoff under REE, which is 26.67 euro for the forecasting task, and 18.33 euro for the trading task.

E Estimation Of Individual Forecasting Rules (Not For Publication)

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	
6		1.000		0.341	0.385	TRE

Table E.1: Estimated individual rules for the LtF treatment.

Subject	Rule coefficients				R ²	Type
	cons.	Past price	AR(1)	Past trend		
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 E.1: (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 E.2: *Estimated individual rules for the LtO treatment.*

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 E.2: (continued) Estimated individual rules for the LtO treatment.

Subject	Quantity rule coefficients				Price prediction rule coefficients				Type	Stability
	cons.	AR(1)	Exp. return	Past return	cons.	AR(1)	Past price	Past trend		
11	0.112		0.161		-2.344		1.214	0.694		U
†12			0.167			0.347	0.650	0.946		S
13		-0.485		0.373	5.654		0.936	1.210		U
†14			0.167		6.335		0.929	0.842		S
15	0.092	0.714					1.000	0.884	TRE	N
16				0.129			1.000	1.089	TRE	S
21			0.037		48.425		0.216			S
22		0.580			14.590		0.759	-0.285		N
†23			0.167		22.056		0.649			S
†24			0.167		17.683		0.712	-0.649		S
25						0.251	0.749		ADA	N
26		0.327								N
31			-0.807			0.212	0.788		ADA	U
32					-13.408	0.535	0.661			N
33					9.092		0.867			N
34				-0.319			1.004			U
†35			0.167		-12.399	0.302	0.878			U
36		0.635				0.402	0.598		ADA	N
†41			0.167				0.931			S
42		0.256			9.748		0.887			S
†43			0.167		3.696	-0.421	1.381	0.994		S
†44			0.167				1.000	0.669	TRE	S
45		0.870					1.000	1.022	TRE	U
46		0.916					1.000	0.996	TRE	N

Table E.3: Estimated individual rules for the mixed treatment

Subject	Quantity rule coefficients				Price prediction rule coefficients				Stability	
	cons.	AR(1)	Exp. return	Past return	cons.	AR(1)	Past price	Past trend		Type
51						-0.831	1.795			N
†52			0.167		6.572		1.257	0.814		U
53			0.065	0.081			1.000	0.860	TRE	S
54	0.093		0.149	0.096			1.000	1.158	TRE	U
55			0.220		7.634		0.886	1.231		U
56		0.564	0.216		6.024		1.103			U
61	-0.129				23.173		0.539			S
62				0.005			1.119			S
63		0.657		-0.447						U
*64					15.897	0.685		0.618		N
*65					28.577		0.431			N
66							1.044			N
71			0.048				1.000			S
†72			0.167				1.000	0.543	TRE	S
73				0.085			0.926			S
74							1.054			N
†75			0.167				1.000	0.471	TRE	S
76			-0.202	0.203		0.585	0.424	0.599		U
81				0.135			1.014			S
82			0.131		5.329	0.953		1.057		S
83			0.195				1.000	0.898	TRE	U
84		0.816			2.947		0.983	0.902		N
85		-0.572		0.281			1.000	0.801	TRE	U
86		0.815					1.000	0.944	TRE	N

Table E.3: Estimated individual rules for the mixed treatment