University of Groningen

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

Bao, Te; Hommes, Cars; Makarewicz, Tomasz

Published in:
Economic Journal

DOI:
[10.1111/econj.12341](https://doi.org/10.1111/econj.12341)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2017

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Bao, T., Hommes, C., & Makarewicz, T. (2017). Bubble Formation and (In)Efficient Markets in Learning-to-forecast and optimise Experiments. *Economic Journal*, 127(605), F581-F609.
<https://doi.org/10.1111/econj.12341>

Copyright

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

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

BUBBLE FORMATION AND (IN)EFFICIENT MARKETS IN LEARNING-TO-FORECAST AND OPTIMISE EXPERIMENTS*

Te Bao, Cars Hommes and Tomasz Makarewicz

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 three times larger than the fundamental value, which were not seen in former experiments.

This article 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; and
- (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

* Corresponding author: Cars Hommes, Center for Nonlinear Dynamics in Economics and Finance (CeNDEF), Amsterdam School of Economics, University of Amsterdam, Roetersstraat 11, E-Building, 1018 WB Amsterdam, The Netherlands. Email: C.H.Hommes@uva.nl.

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), Integrated Macro-Financial Modelling for Robust Policy Design (MACFINROBODS, grant 612796), the Institute of New Economic Thinking (INET) grant project 'Heterogeneous Expectations and Financial Crises' (INO 1200026), National Natural Science Foundation of China grant (number 71301174), and a grant from the Netherlands Institute for Advanced Study in the Humanities and Social Sciences (NIAS).

experiment and should lead the subjects to an equilibrium with a constant fundamental price. In contrast, we find none of the experimental markets 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 in the modern economics and finance literature. One possible reason is that they contradict the standard theory of rational expectations (Muth, 1961; Lucas, 1972) and efficient markets (Fama, 1970). Recently, however, there has been 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 policymakers 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&P500 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 reproduce price bubbles and crashes of asset prices reliably in a laboratory setting. They let the subjects trade an asset that pays a dividend in each of 15 periods. The fundamental price in each period equals the sum of the remaining expected dividends and follows a decreasing step function. The authors find the price goes 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 finding in these papers is that the price boom and bust is robust, 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 2007/8 financial crisis.

The Smith *et al.* (1988) experiment is categorised as ‘learning to optimise’ (henceforth LtO) experiments; see Duffy (2008) for an extensive discussion. Besides

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

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 forecast, which is transformed into a quantity decision of buying/selling by a computer program based on optimisation 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:

- (i) the asset price fails to converge to the fundamental but oscillates and forms bubbles in several markets; and
- (ii) instead of having rational expectations, most subjects follow price trend extrapolation strategies (Bostian and Holt, 2009).

Heemeijer *et al.* (2009) and Bao *et al.* (2012) investigate whether the non-convergence result is driven by the positive expectation feedback nature of the experimental market. 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 overreact in the long run.

The subjects in Hommes *et al.* (2005) and other ‘learning to forecast’ experiments do not trade directly 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 artefact of the computerised trading in the LtF design?

In this article, we design an experiment, in which the fundamental price is constant over time (Hommes *et al.*, 2005) but the subjects are asked to indicate the amount of asset they want to buy/sell directly. 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 some very large and repeated price oscillations occur in the mixed, LtO and LtF designs, where the price peaks at more than three times the fundamental price. This was not observed in previous 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 observed differences in aggregate macrobehaviour 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 characterised 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 article 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, in which all markets converge to the fundamental price, and find differences in the speed of convergence across treatments.

The article 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 rational expectation equilibrium (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. The subject's task is either to predict

the asset price, or suggest the trading quantity, or both; 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 is explained in detail in subsection 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's 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:

1.1.1. *LtF: classical learning-to-forecast experiment*

The 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 B1 in online 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 REE, $\bar{p}_{t+1}^e \equiv 1/6 \sum_{i=1}^6 p_{i,t+1}^e$ 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 $p_{i,t+1}^e = 66$.

1.1.2. *LtO classical learning-to-optimize experiment*

Here 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 a 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:

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

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 (3). Subjects are rewarded according to:

$$Payoff_{i,t+1} = \max \left\{ 0, 800 + 40[z_{i,t}(p_{t+1} + 3.3 - 1.05p_t) - 3z_{i,t}^2] \right\}. \quad (4)$$

This payoff corresponds to a mean-variance utility function of the financial firms in the underlying economy, as explained below. The 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 B2 in online 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:

$$z_{i,t}^* = \frac{p_{i,t+1}^e + 3.3 - 1.05p_t}{6} = \frac{p_{t+1}^e}{6}.$$

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 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.

1.1.3. *Mixed*

A detailed discussion is provided in online Appendix C. 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 the payoff in either (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 price forecasts submitted.

The points in each treatment are exchanged into euro at the end of the experiment with the conversion rate 3,000 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 online 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 realised 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 subjects 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 subsection 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 maximisers of next period's wealth, i.e. they solve the myopic optimisation problem:

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

where a is a parameter for risk aversion and ρ_{t+1} is the excess return as defined in (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 by:³

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

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

² Fixing the dividend permits 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.

³ The 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.

adjustment mechanism in response to excess demand (Beja and Goldman, 1980),⁴ given by:

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

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 coincides 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 $p_{i,t+1}^e$ is:

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

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 $\bar{p}_{t+1}^e = 1/6 \sum_{i=1}^6 p_{i,t+1}^e$ 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

⁴ See 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: newline (i) market maker is a stylised 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 three-dimensional payoff table instead of the two-dimensional payoff table in online 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 destabilising force in the financial markets – is likely to be robust in a similar two-period ahead LtO experiment with a market clearing design.

⁵ Heemeijer *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.

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.

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 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 an 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 14, 17, 18 and 20 December 2012 and 6 June 2014 at the CREED Laboratory, University of Amsterdam. 144 subjects were recruited. The experiment employs a group design with six subjects in each experimental market. There are 24 markets in total and eight for each treatment. No subject participates in more than one session. The duration of the experiment is typically about one 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 online 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 (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 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 3(a), (d) and (h)) 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 dashed lines. In the LtO treatment, however, there is a high level of heterogeneity in the quantity decisions shown by the 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

⁶ We 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.

Fig. 1. Price Dynamics in Learning to Forecast (LtF) Treatment

Notes. Groups 1–8 for the LtF treatment. Straight line shows the fundamental price $p^f = 66$, solid black line denotes the realised price, while grey dashed lines denote individual forecasts.

Fig. 1. (Continued).

Fig. 2. Price Dynamics in Learning to Optimize (LtO) Treatment

Notes. Groups 1–8 for the LtO 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 grey lines) and average trade (solid grey line). Notice the different y-axis scale for group 7 (g).

Fig. 2. (Continued).

Fig. 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 (grey dashed lines), while the lower panel displays individual trades (dashed lines) and average trade (solid grey line). Notice the different y-axis scale for groups 1, 4 and 8 ((a), (d) and (h) respectively).

Fig. 3. (Continued).

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. There may be two causes:

- (i) 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
- (ii) 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 discuss the statistical evidence for the hypotheses in detail.

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.

Table 1
RAD and RD

Treatment group no	LtF		LtO		Mixed	
	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. *** (**) denote 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 the $[-1.5\%, 1.5\%]$ interval on 5% significance level. LtF, learning to forecast; LtO, learning to optimise.

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 article, 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 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 (a high average $\overline{RAD}$ of 24.7%), but centred close to the fundamental price (average $\overline{RD}$ of 1.4%, compared to an average $\overline{RD}$ of −3% and 16.1% for the LtF and mixed treatments respectively). LtF groups on average are below the fundamental price and mixed groups typically overshoot it.

A simple t-test shows that for the LtO and mixed treatment, as well as for 6 out of 8 LtF groups (exceptions are markets 7 and 8), the means of the groups’ RAD measures

⁷ As 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).

(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 $\overline{RD}$ according to the Mann-Whitney–Wilcoxon rank-sum test ($p > 0.1$ 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 ($p = 0.002$ 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 ($p = 0.753$, 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 article 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 show that the mispricing in experimental asset markets is a robust finding. The collapse of a bubble does not force 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 underpricing 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 has dynamics closest to the REE. Nevertheless, the average price is still far from the rational expectations equilibrium. Furthermore, in terms of aggregate dynamics, the LtF treatment is significantly different from the other two treatments, which are indistinguishable from each other. 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 1,300 points per

⁸ 3% $\overline{RAD}$ 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.

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 D1, in online Appendix D, 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 $p = 0.001$). 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 $p = 0.753$).

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 (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 (solid grey) line with slope 1/6.

Figure 4(a) 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).

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 $\eta_{i,t} \sim \text{NID}[(0, \sigma_{\eta,i}^2)]$. The estimated coefficients for all subjects are shown in the scatter plot of Figure 4(b). 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

⁹ Sometimes 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.

Fig. 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 (top panel)/triangle (bottom panel) denotes the optimal trade rule ($z_{i,t} = \rho_{i,t}^e/6$). Dashed line (top panel)/square (bottom panel) denotes the estimated rule under restriction of homogeneity ($z_{i,t} = c + \theta \rho_{i,t}^e$).

$\theta_i = \theta_j$ ($\theta_i \neq \theta_j$) 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 (17) with the restrictions on the parameters c_i and θ_i .¹⁰ These regressions are compared with an unrestricted regression (with $\theta_i \neq \theta_j$ and $c_i \neq c_j$) using 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. Estimation identified 11 subjects (23%) as consistent optimal traders (see footnote 5 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; 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:

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

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:

¹⁰ We 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 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 report $z_{i,t} = 0$ for $t > 10$. Interestingly, these two subjects had non-constant price predictions, which suggests that they were not optimisers.

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

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, we estimate the behavioural heuristic for each subject, 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

¹¹ Adaptive expectations (20) impose a restriction $\alpha \in [0,1]$ (with $\alpha = 1 - \beta$), so here we follow 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 LR test against the likelihood of (18).

¹² See footnote 4.

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 (21–22) are tested based on reported or fitted price forecasts according to the Hausman test.¹³

The estimation results can be found in online Appendix E, in Tables E1, E2 and E3 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 (18) where $\gamma_i = 0$, but $\alpha_i + \beta_i \neq 1$. There were no subjects in the LtF treatment and only two 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 (22) 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

¹³ Whenever the estimation 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 as 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.

dominating rule (33 in the LtO and 32 in the mixed treatment). There are only 12 subjects using a significant ARI 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 realised and/or expected asset returns, which leads to relatively strong trend chasing behaviour. Nevertheless, there are 11 subjects in the LtO treatment and eight in the mixed treatment for whom we cannot identify a trading rule within this simple class. The average demand persistence was $\bar{\chi} = 0.07$ and $\bar{\chi} = 0.006$; the average trend extrapolation was $\bar{\phi} = 0.09$ and $\bar{\phi} + \bar{\zeta} = 0.06$ 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 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 with 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:

¹⁴ The 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.

¹⁵ The implied trading rule (8) however cannot be rewritten in exactly the form of (21a); it 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 cannot reject the hypothesis that the mean coefficient is 0 (p-value 0.15).

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.*

3. Conclusions

The origin of asset price bubbles is an important topic for both researchers and policy-makers. This article 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 previous empirical work quantifying forecast biases by households and finance professionals in real markets and theoretical work that starts to incorporate the stylised 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 cognitively 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 tests, 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, while 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 the 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 studies 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.

Nanyang Technological University and University of Groningen
University of Amsterdam and Tinbergen Institute
University of Amsterdam and Tinbergen Institute

Additional Supporting Information may be found in the online version of this article:

Appendix A. Experimental Instructions.

Appendix B. Payoff Tables.

Appendix C. Rational Strategic Behaviour.

Appendix D. Earnings Ratios.

Appendix E. Estimation of Individual Forecasting Rules.

Data S1.

References

- Akiyama, E., and Hanaki, N. and Ishikawa, R. (2012). ‘Effect of uncertainty about others’ rationality in experimental asset markets: an experimental analysis’, Available at: <http://ssrn.com/abstract=2178291> (last accessed: 9 November 2015).
- 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(8), pp. 1523–43.

- Assenza, T., Bao, T., Hommes, C. and Massaro, D. (2014). 'Experiments on expectations in macroeconomics and finance', in (J. Duffy, ed.), *Experiments in Macroeconomics*, vol. 17 of *Research in Experimental Economics*, pp. 11–70, Bingley: 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. 1101–20.
- Beja, A. and Goldman, M.B. (1980). 'On the dynamic behavior of prices in disequilibrium', *Journal of Finance*, vol. 35(2), pp. 235–48.
- 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', Working Paper No. 450, De Nederlandsche Bank.
- Bondt, W. and Thaler, R. (2012). 'Further evidence on investor overreaction and stock market seasonality', *Journal of Finance*, vol. 42(3), pp. 557–81.
- Bostian, A.A. and Holt, C.A. (2009). 'Price bubbles with discounting: a webbased 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. 1235–74.
- Campbell, J. and Shiller, R. (1989). 'Stock prices, earnings and expected dividends', Working Paper, NBER.
- Campbell, J.Y., Lo, A.W. and MacKinlay, A.C. (1997). *The Econometrics of Financial Markets*, Chichester: Princeton University Press.
- Chiarella, C., Dieci, R. and He, X.Z. (2009). 'Heterogeneity, market mechanisms, and asset price dynamics', in (T. Hens and K.R. Schenk-Hoppe, eds.), *Handbook of Financial Markets: Dynamics and Evolution*, vol. 231, pp. 277–344, Amsterdam: Elsevier North Holland.
- Choi, S., Galeotti, A. and Goyal, S. (2014). 'Trading in networks: theory and experiments', Working Paper No. 8, Cambridge-INET.
- Duffy, J. (2008). 'Macroeconomics: a survey of laboratory research', in (J. Kagel and A. Roth, eds.), *Handbook of Experimental Economics*, vol. 2, Princeton: Princeton University Press. Forthcoming.
- Dufwenberg, M., Lindqvist, T. and Moore, E. (2005). 'Bubbles and experience: an experiment', *American Economic Review*, vol. 95(5), pp. 1731–7.
- 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. 9991–2.
- Füllbrunn, S., Rau, H.A. and Weitzel, U. (2014). 'Does ambiguity aversion survive in experimental asset markets?', *Journal of Economic Behavior & Organization*, vol. 107(Part B), pp. 810–26.
- 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. 114–32.
- Galeotti, A., Goyal, S., Jackson, M.O., Vega-Redondo, F. and Yariv, L. (2010). 'Network games', *Review of Economic Studies*, vol. 77(1), pp. 218–44.
- 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. 1119–57.
- 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. 1052–72.
- 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: Cambridge University Press.
- Hommes, C. and in't Veld, D. (2014). 'Booms, busts and behavioural heterogeneity in stock prices', CeNDEF Working Paper No. 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. 955–80.
- 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. 865–83.
- LeBaron, B. (2006). 'Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics', in (L. Tesfatsion and K.L. Judd, eds.), Amsterdam: Elsevier Science B.V.
- 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. 831–59.

- Lucas Jr., R.E. (1972). 'Expectations and the neutrality of money', *Journal of Economic Theory*, vol. 4(2), pp. 103–24.
- 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. 315–35.
- Noussair, C.N. and Tucker, S. (2013). 'Experimental research on asset pricing', *Journal of Economic Surveys*, vol. 27(3), pp. 554–69.
- 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.
- Rubinstein, A. (2007). 'Instinctive and cognitive reasoning: a study of response times', *Economic Journal*, vol. 117(523), pp. 1243–59.
- 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. 1119–51.
- Stöckl, T., Huber, J. and Kirchler, M. (2010). 'Bubble measures in experimental asset markets', *Experimental Economics*, vol. 13(3), pp. 284–98.
- Sunder, S. (1995). 'Experimental asset markets: a survey', in (J. Kage 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. 1124–31.