ARTICLE IN PRESS

Journal of Economic Dynamics & Control ∎ (∎∎∎∎) ∎∎∎–∎∎∎

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation

Peter Heemeijer, Cars Hommes, Joep Sonnemans, Jan Tuinstra *

Center for Nonlinear Dynamics in Economics and Finance (CeNDEF), University of Amsterdam, School of Economics, Roetersstraat 11, 1018WB Amsterdam, The Netherlands

ARTICLE INFO

Article history:
Received 30 January 2008
Accepted 19 September 2008

JEL classification:
D02
G12
C92

Keywords:
Market behavior
Coordination
Expectations feedback
Experimental economics

ABSTRACT

The evolution of many economic variables is affected by expectations that economic agents have with respect to the future development of these variables. We show, by means of laboratory experiments, that market behavior depends to a large extent on whether realized market prices respond positively or negatively to average price expectations. In the case of negative expectations feedback, as in commodity markets, prices converge quickly to their equilibrium value, confirming the rational expectations hypothesis. In the case of positive expectations feedback, as is typical for speculative asset markets, large fluctuations in realized prices and persistent deviations from the benchmark fundamental price are likely. We estimate individual forecasting rules and investigate how these explain the differences in aggregate market outcomes.

1. Introduction

A key difference between natural and social sciences is that in social systems individual expectations or beliefs can affect the aggregate outcome. An investor buys a stock that he expects to go up in the future, a chip-manufacturer builds a new production facility because she expects that demand and therefore prices will be high after goods have been produced. Expectations determine individual behavior of economic agents and the actual market outcome, i.e. prices and traded quantities, are determined as an aggregation of individual behavior. Simultaneously, economic agents form their expectations on the basis of market history. Therefore, a market, like other social environments, may be viewed *as an expectations feedback system*: past market behavior determines individual expectations which, in turn, determine current market behavior and so on. Two important types of feedback can be distinguished: positive and negative. A market has positive (negative) expectations feedback, when a higher average price forecast yields a higher (lower) realized market price. This paper may be viewed as an experimental testing of whether, and if so how expectations feedback affects aggregate market outcome and individual forecasting behavior.

In many markets both types of feedback play a role. Positive feedback however, seems to be particularly relevant for speculative asset markets. If many agents expect the price of an asset to rise and therefore start buying the asset, aggregate demand will increase, and so, by the law of supply and demand, will the asset price. High price expectations thus become self-confirming and lead to high realized asset prices. Similarly, when a majority of investors expects markets to go down, this belief will be self-fulfilling and the market will go down. In markets where the role of speculative demand is less

* Corresponding author. Department of Quantitative Economics and CeNDEF, University of Amsterdam, Roetersstraat 11, 1018WB Amsterdam, The Netherlands. Tel.: +3120 525 4227; fax: +3120 525 4349.
E-mail address: J.Tuinstra@uva.nl (J. Tuinstra).

0165-1889/$ - see front matter © 2009 Published by Elsevier B.V.
doi:10.1016/j.jedc.2008.09.009

important, e.g. in markets for non-storable commodities, negative feedback may play a more prominent role. Consider e.g. a supply-driven commodity market. If many producers expect future prices to be high they will increase production which, according to the law of supply and demand, will lead to a low realized market price. In contrast, if firms expect a low price, they will produce little and, as a consequence, the realized market price will be high.[1]

To investigate how the expectations feedback structure affects aggregate market outcomes and individual forecasting behavior we designed experimental market environments that only differ in the sign of the expectations feedback, but are equivalent along all other dimensions. We compare these markets with respect to the coordination of individual expectations and (speed of) convergence to the fundamental market equilibrium price.

The distinction between positive and negative expectation feedback is related to the concept of, respectively, strategic complements and strategic substitutes. Strategic substitutability (complementarity) prevails if an increase in the action of individual *i* generates an incentive for *j* to decrease (increase) his action. Haltiwanger and Waldman (1985) argue that when actions are strategic complements, agents have an incentive to imitate other agents. This is the case in an asset market, where predicting a price close to the predictions of the other participants turns out to be most profitable. Coordination of predictions enhances the impact of the irrational participants upon the realized prices and convergence to the rational equilibrium price becomes unlikely. When actions are strategic substitutes, agents have an incentive to deviate from what other agents are doing. This is the case in negative feedback markets, where agents have an incentive to predict high (low) prices when the majority predicts prices below (above) the equilibrium price. The impact of irrational individuals will be limited and convergence to the equilibrium price is more likely. Coordination of predictions will only take place after convergence. Guesnerie (2005), on the other hand, argues that eductive stability is typically easier to obtain in an economy with strategic complements than in an economy with strategic substitutes. Finally, Haltiwanger and Waldman (1989) argue that, if not all agents are rational and provided agents' actions are strategic complements, a one-time shock will lead to disproportionally more persistence than is to be expected given the fraction of adaptive agents.

In recent experiments Fehr and Tyran (2008) study the impact of different strategic environments (strategic complementarity versus strategic substitutability) on individual rationality and aggregate outcomes. Based upon the ideas of Haltiwanger and Waldman (1989), Fehr and Tyran study the adjustment of nominal prices after an anticipated money shock in a price setting game with positively (complements) or negatively sloped (substitutes) reaction curves, and find much faster convergence in the case of substitutes.[2] They argue that differences in the "stickiness of price expectations" are key for understanding these different aggregate outcomes. In our experiments, although participants have less information as we will see below, the results are similar: the market with negative expectations feedback quickly converges while the market with positive expectations feedback converges very slowly or oscillates and exhibits significant excess volatility. Another distinguishing feature of our approach is that we estimate individual forecasting rules and thus obtain detailed information about the *relation between individual expectations and aggregate outcomes* in strategic complementary market environments (positive expectations feedback) and strategic substitutability market environments (negative expectations feedback).

Since the seminal work of Muth (1961) and Lucas (1972), rational expectations became the dominating expectations hypothesis in the 1970s and 1980s. More recently models of bounded rationality and adaptive learning have become popular in macroeconomics and finance, as discussed e.g. in Sargent (1993), Honkapohja (1993), Evans and Honkapohja (2001), Barberis and Thaler (2003) and Hommes (2006). The state of the art on adaptive learning has recently been nicely formulated in an overview by Bullard (2006, p. 205):

> The state of affairs is thus that some rational expectations equilibria are learnable while others are not. Furthermore, convergence will in general depend on all the economic parameters of a given system, [...] (that is, it depends on the entire economic structure).

Laboratory experiments on expectations and forecasting have e.g. been performed by Schmalensee (1976), Hey (1994), Dwyer et al. (1993) and Kelley and Friedman (2002). In particular, several experiments in environments with *expectations feedback* between individual forecasts and aggregate market outcomes have been studied in Marimon and Sunder (1993), Marimon et al. (1993), Peterson (1993), Gerber et al. (2002), Hommes et al. (2005b, 2007, 2008), Adam (2007) and Sutan and Willinger (2009). The current paper is motivated by striking differences in earlier experiments on expectation formation in an asset pricing framework in Hommes et al. (2005b, 2008) and a cobweb framework in Hommes et al. (2007). In the asset pricing framework coordination on trend following strategies caused oscillatory price movements with persistent deviations from fundamentals, whereas in the cobweb framework prices were relatively stable, often moving close towards the equilibrium level. Can these differences be attributed to the sign of the expectations feedback structure? To answer this from the earlier experiments is difficult, because there are a number of other important differences in the

[1] Another example of negative expectations feedback is provided by the so-called *minority game* (see e.g. Challet and Zhang, 1997). In this game players choose between two actions and only the players making the minority choice receive a positive payoff. Interestingly, in the econophysics literature this highly stylized game has been used extensively as a model of a financial market.

[2] Another related study concerning strategic substitutes versus strategic complements is Potters and Suetens (2009). Their focus, however, is more on social behavior than on convergence and coordination. Cournot games and Betrand games are both social dilemma situations (the Nash equilibrium is Pareto-inefficient), but in Cournot games actions are strategic substitutes while in Bertrand games they are strategic complements. Potters and Suetens design two games with the same Nash equilibrium, the same social optimum and the same absolute (but opposite) slope of the reaction curve. They find more cooperation in the case of strategic complements than when actions are strategic substitutes. Also see Fehr and Tyran (2005), who review and interpret other results from the experimental economics literature in light of the difference between strategic substitutes and complements.

experimental designs. For example, the temporary equilibrium asset pricing framework required two-period ahead forecasts, whereas in the cobweb framework participants had to make one-period ahead forecasts. Another important difference is the presence of fundamental "robot traders" in the asset pricing framework of Hommes et al. (2005b). The main goal of the current paper is to investigate whether a different expectations feedback structure (positive versus negative) *alone*, leads to differences in individual forecasting rules and aggregate market behavior.

The paper is organized as follows. In Section 2 we discuss the positive and negative feedback systems and set out the experimental design. Section 3 describes aggregate market behavior, Section 4 investigates individual forecasting rules, and Section 5 concludes. Finally, Appendices A–C contain a more detailed description of the models underlying the positive and negative expectations feedback systems, a description of the experimental instructions and detailed estimation results of the individual prediction strategies.

2. Expectations feedback and experimental design

In both treatments of the experiments the price adjustment rule will be of the simple form

$$p_t = f(\overline{p}^e_t), \tag{1}$$

where $\overline{p}^e_t$ is the *average price forecast* of all participants in the market. Appendix A describes how this pricing rule can be derived from the "law of supply and demand" for a cobweb model (negative feedback) and an asset pricing model (positive feedback). In fact, since we assume demand and supply to be linear, in the experiments the map *f* in (1) will be linear, with negative, respectively, positive slope. As discussed above, positive versus negative feedback is related to the concepts of strategic complements versus strategic substitutes. We use positive/negative feedback to stress the fact that the difference is entirely due to the sign of the derivative of the expectations feedback map in (1).[3]

For the *negative feedback treatment* the price adjustment rule is given by

$$p_t = \frac{20}{21}(123 - \overline{p}^e_t) + \varepsilon_t, \tag{2}$$

while for the *positive feedback treatment* it is given by

$$p_t = \frac{20}{21}(\overline{p}^e_t + 3) + \varepsilon_t. \tag{3}$$

In both cases, $\overline{p}^e_t = \frac{1}{6}\sum_{h=1}^{6} p^e_{h,t}$ is the average prediction of the six participants in the experimental market, and $\varepsilon_t \sim N(0, 1/4)$ is a random term, representing e.g. small random fluctuations in the supply of the risky asset. The two treatments are perfect symmetric opposites. One can easily check that, for both (2) and (3), $p^* = 60$ corresponds to the steady state equilibrium price: if every participant predicts $p^e_{h,t} = 60$ the resulting realized market price is given by $p_t = 60 + \varepsilon_t$. Both price series (2) and (3) are generated as a linear function of the average predictions of six participants, the realization of the random shocks is exactly the same and the absolute value of the slope of the relation between p_t and $\overline{p}^e_t$ is equal to 20/21 for both treatments.[4] The only difference between the treatments is the sign of this slope. Note also that $p^* = 60$ is a stable steady state under naïve expectations, $p^e_{h,t} = p_{t-1}$, since the absolute value of the slope of both (2) and (3) is given by 20/21 and is smaller than one. However, observe also that since the absolute value of the slope is close to one, convergence will be relatively slow.[5] A smaller (absolute) value for the slope induces much faster convergence, and would therefore obfuscate the differences between the two treatments.[6]

Thirteen experimental markets of 50 periods were created, six with negative and seven with positive feedback. In each market six students participated, who earned more money if they predicted market prices more successfully. The computerized experiments were conducted in the CREED Laboratory at the University of Amsterdam on February 18 and 19, 2003. Fig. 1 shows the main experimental computer screen. Participants have information about past realized market prices and past own predictions, but have no information about other individual predictions. Participants were not explicitly informed about the price generating mechanisms ((2) or (3)), but did obtain qualitative information about the market, in particular that the market price was determined by the "law of demand and supply", that is, the market price will increase (decrease) when there is excess demand (supply). The reason for not providing participants with more information is that we wanted to mimic the situation on real markets as closely as possible. In those real markets typically market participants only have qualitative information and e.g. know the direction of price changes caused by market expectations, but not the exact magnitude of these changes, and also have no (or little) information about the beliefs of other agents.[7] Appendix B contains a detailed description of the experimental instructions.

[3] Alternatively, Haltiwanger and Waldman (1985) use the terminology "congestion effects" for strategic substitutes/negative feedback and "synergistic effects" for strategic complements/positive feedback.

[4] In a dynamic asset pricing model this slope is typically 1/(1+*r*), where *r* is the risk free interest rate. This slope is close to +1, and for an interest rate of 5%, i.e. $r = 0.05$, the slope becomes 20/21.

[5] In particular, under naïve expectations, starting at an initial expectation of 50 and in the absence of random shocks this means that in both treatments prices will be within a 3% distance of the equilibrium price only in period 37 and remain there from then on.

[6] See Sonnemans and Tuinstra (2008) who report on a positive feedback treatment where the slope of the price generating mechanism equals 2/3 and convergence to the steady state is indeed very fast.

[7] Sonnemans and Tuinstra (2008) also run a positive feedback treatment where participants have full information about the price generating mechanism. This does not qualitatively change the results.

Please cite this article as: Heemeijer, P., et al., Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. Journal of Economic Dynamics and Control (2009), doi:10.1016/j.jedc.2008.09.009

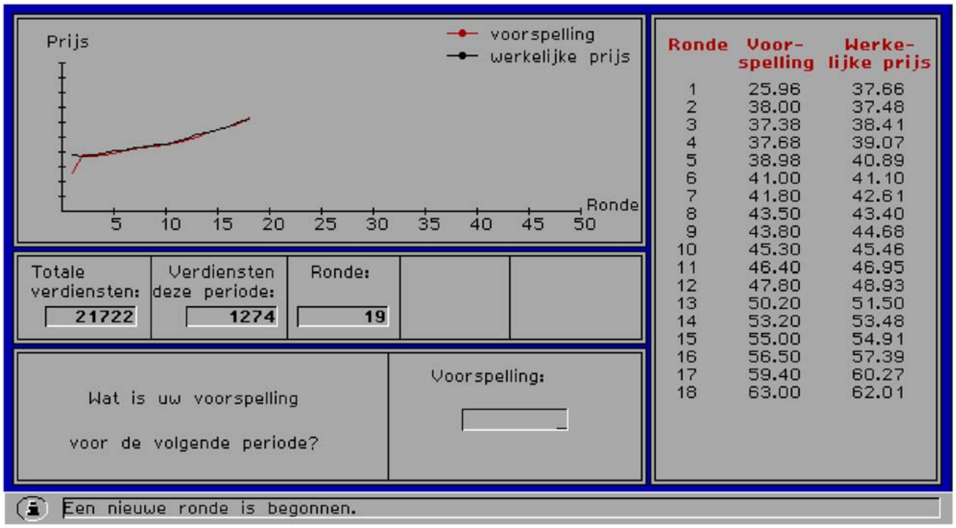

Fig. 1. The main experimental computer screen for one of the markets in the asset pricing treatment, as seen by one of the participants. Participants have information, both graphically and in a table, about previous realized market prices and previous own predictions. The Dutch labels of the computer screen translate as follows: "prijs" = price; "voorspelling" = prediction; "werkelijke prijs" = market price; "ronde" = round; "totale verdiensten" = total earnings; "verdiensten deze periode" = earnings this period; "Wat is uw voorspelling voor de volgende periode?" = What is your prediction for the next period?; and "Een nieuwe ronde is begonnen" = A new round has started.

The computer screen in Fig. 1 shows the actual development of one of the experimental markets, in this case for the positive feedback treatment, from the perspective of one of the participants. The participant observes both a graphical and a numerical representation of the realized market prices and his previous price predictions, in the upper left and right panel, respectively. In the middle left panel information is displayed regarding the total earnings so far of the participant, his earnings in the last period and the present time period of the experiment. The participant submits his price prediction for the next period in the lower middle panel. The individual earnings E per period are based on the quadratic error of the prediction:

$$E = Max\left\{1300 - \frac{1300}{49}(p_t - p^e_{h,t})^2, 0\right\} \tag{4}$$

and 1300 points corresponds to 0.5 euro. Earnings were on average about 22 euros in approximately 90 minutes.

3. Aggregate market behavior

Results of the experiments are shown in Fig. 2 (6 markets with negative feedback) and Fig. 3 (7 markets with positive feedback). Each individual panel shows, for one experimental market, the realized prices and the six time series of individual predictions. Two characteristics of the data catch the eye immediately.

First, in the negative feedback market prices tend to go through an initial phase of high volatility, neatly converging afterwards to the equilibrium price 60, only to be disturbed occasionally by the impact of a mistake by one of the group members. Under rational expectations, and complete information, all participants would predict $p^e_{h,t} = 60$, with resulting market price $p_t = 60 + \varepsilon_t$. We will investigate whether this rational expectations benchmark gives a good description of aggregate outcomes. Allowing for a short initial learning phase,[8] for all six negative feedback markets average prices are statistically not significantly different at a 5% level from what the rational expectations hypothesis predicts. Volatility is not significantly different at a 5% level from the volatility under rational expectations for the first (N1) and fourth (N4) market.

In contrast, the results for the positive feedback markets are strikingly different. Although the heterogeneity of predictions decreases in a much shorter period, a quick convergence to the equilibrium price does *not* occur. Rather, most groups show a slow oscillatory movement around the equilibrium price of 60, and come close to it only in the very long run. Average prices and volatility in all positive feedback groups are significantly different at a 5% level from the price and volatility under rational expectations. We therefore have the following result.

Convergence of prices is demonstrated in more detail in the upper panel of Fig. 4 which shows the median over 6 or 7 markets of the absolute difference between the market price and the equilibrium price of 60 for both treatments. We find a much higher degree of convergence to the equilibrium price in the negative feedback treatment, already after period two (statistically significant at 5% in 44 of the 48 periods, Wilcoxon test).

[8] This initial learning phase is defined such that it ends when the majority of individual forecasts (i.e. at least 4 participants in each group) is within 5% of the realized market price.

Please cite this article as: Heemeijer, P., et al., Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. Journal of Economic Dynamics and Control (2009), doi:10.1016/j.jedc.2008.09.009

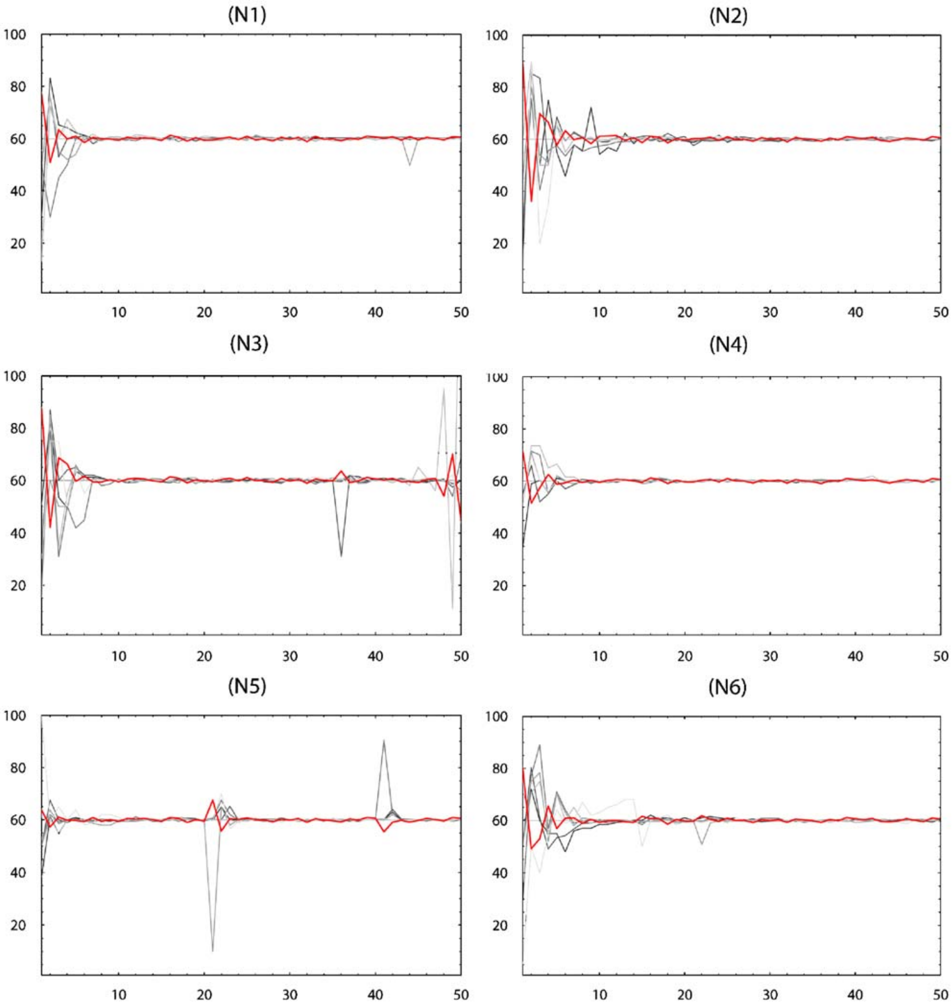

Fig. 2. Prices and predictions in the negative feedback treatment. Each panel contains, for one experimental market, time series for the realized price (in red) and the time series of individual prediction of the six participants.

Convergence can also be defined as a situation where the realized price enters a small neighborhood of the equilibrium price and remains within that neighborhood for the remaining periods. Consider as a neighborhood deviations of maximally 3% from the equilibrium price, that is the interval [58.2, 61.8].[9] Using this neighborhood we find that for the negative feedback treatment prices in group 3 have not converged by the end of the experiment and for the five other groups convergence takes place by periods 4, 7, 5, 42 and 23, respectively. For the positive feedback treatment we do not have convergence for any of the groups.[10]

Second, in both treatments there is little dispersion between individual predictions within experimental markets, which is particularly remarkable for the non-converging positive feedback treatment.[11] Participants in the positive feedback treatment *quickly coordinate* on a common non-equilibrium prediction rule.

Coordination of expectations is measured by the standard deviation of the individual expectations of the market participants. The lower panel of Fig. 4 shows the median over 6 or 7 markets of these standard deviations for each period. A low standard deviation implies a high level of consensus among the participants about the future price. We find that

[9] Under rational expectations the lower and upper bound of the realized price are given by 59.15 and 60.91, respectively, that is, deviations of up to 1.52% of the equilibrium price of 60. As a neighborhood we chose an interval that is twice as large, i.e. allowing deviations from the equilibrium price up to 3%. Qualitative similar results are obtained if we take a slightly smaller or larger neighborhood (e.g. allowing deviations of, say, 5% of the equilibrium price).

[10] One remaining issue is whether the groups from the positive feedback treatment would converge eventually. Given the limited number of periods this is difficult to answer from our data but one indication that this is not very likely is the following. For four of the seven positive feedback groups the average deviation from the equilibrium price in the last five periods of the experiment (periods 46–50) is larger than the average deviation in the 25 periods before that (i.e. in periods 21–45), suggesting that typically realized prices are not approaching the equilibrium price. A more formal approach to test asymptotic convergence based upon a limited number of periods is outlined in Noussair et al. (1995).

[11] Remarkably, for the negative feedback treatment the realized price seems to converge *before* the price predictions do. This is similar to the finding by Ehrblatt et al. (2006) that actions might converge before beliefs in a repeated game experiment.

Please cite this article as: Heemeijer, P., et al., Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. Journal of Economic Dynamics and Control (2009), doi:10.1016/j.jedc.2008.09.009

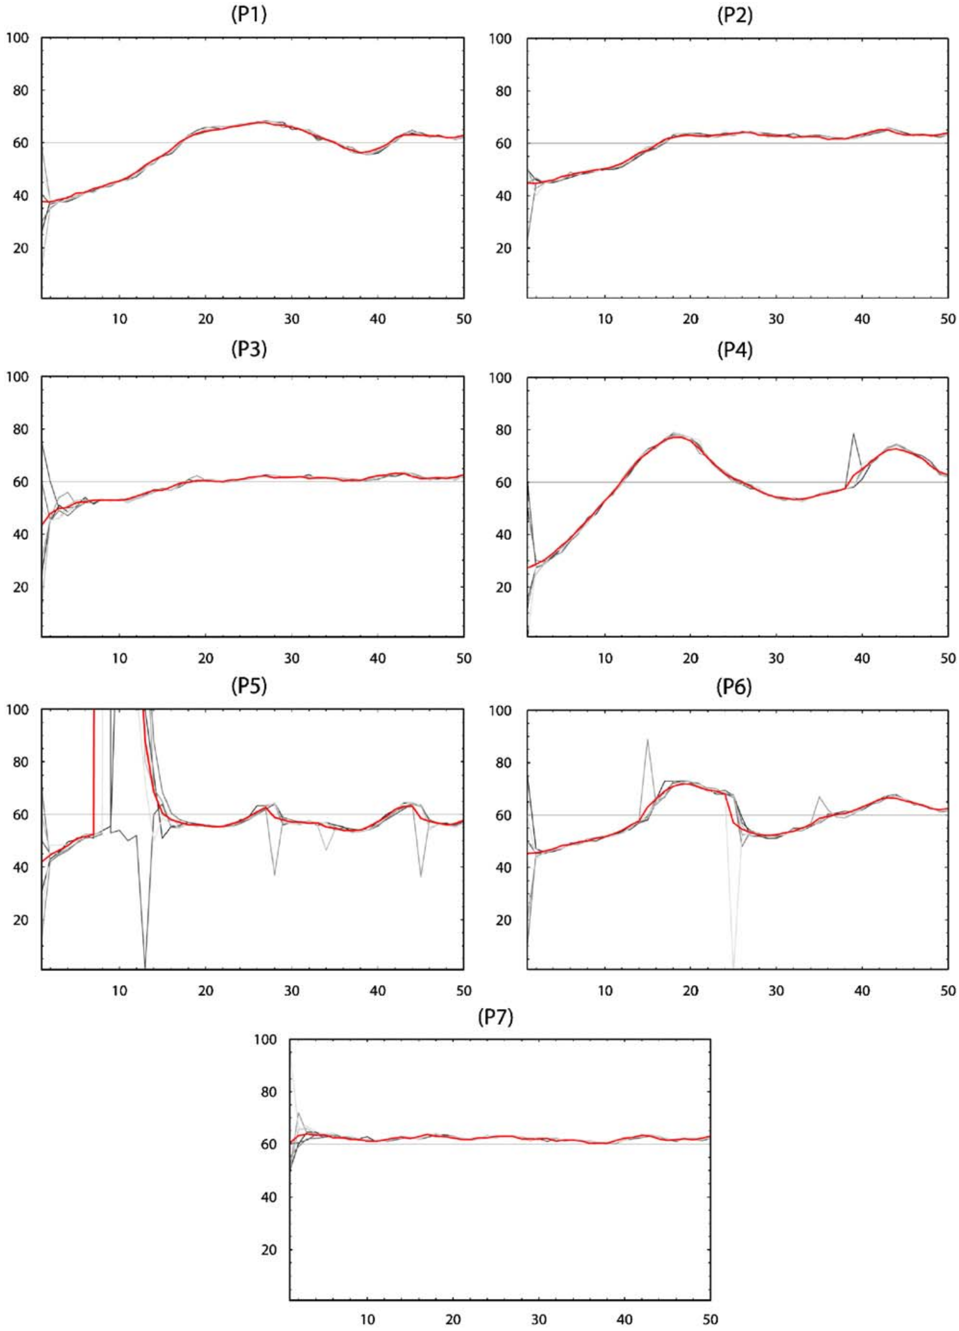

Fig. 3. Prices and predictions in the positive feedback treatment. Each panel contains, for one experimental market, time series for the realized price (in red) and the time series of individual prediction of the six participants.

the standard deviation is higher (and therefore coordination is less) for the negative feedback treatment in the early periods 2–7 (statistically significant at 5%, Wilcoxon test). After period 7, coordination is very high in both treatments. Note that, outside of equilibrium, it pays off for participants in the negative feedback treatment to "disagree" with the majority: if the average prediction is high, the realized price will be low. This drives the heterogeneity in predictions in the early periods and the fast convergence to the equilibrium price. In the positive feedback treatment, on the other hand, "agreeing" with the majority pays off since the market price will be close to the average price prediction. This quick coordination of price predictions in the positive feedback treatment is surprising, since participants were not able to observe each others predictions during the experiment, making the coordination itself "blind". Individual price predictions and aggregate market prices can concisely be summarized as exhibiting "slow coordination and fast convergence" in the negative feedback treatment, and "fast coordination and slow convergence" in the positive feedback treatment.

Please cite this article as: Heemeijer, P., et al., Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. Journal of Economic Dynamics and Control (2009), doi:10.1016/j.jedc.2008.09.009

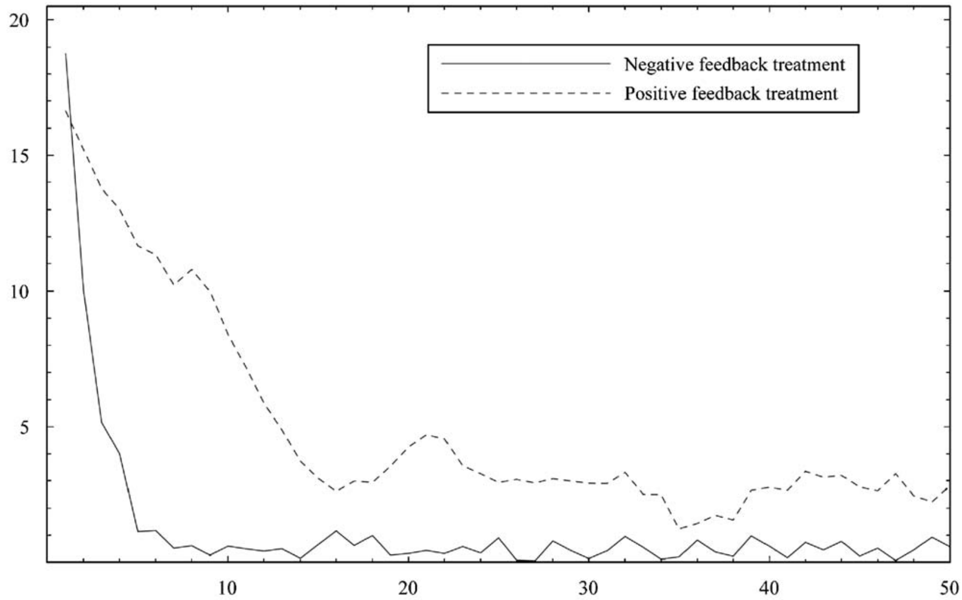

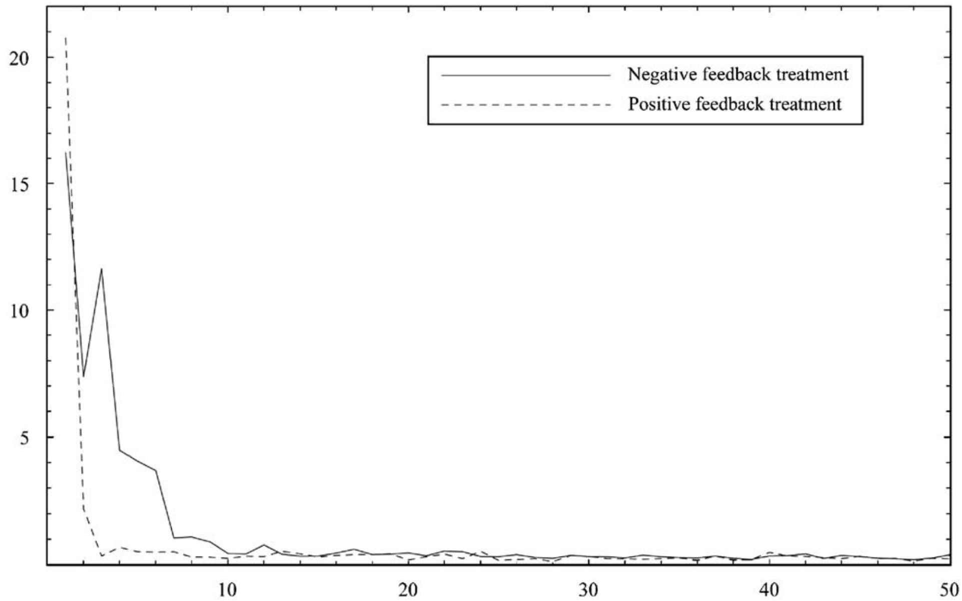

Fig. 4. Upper panel gives the median, over the different groups, of the absolute difference between the market price and the equilibrium price; the lower panel gives the median, over the different groups, of the standard deviations of individual predictions. Solid lines correspond to the negative feedback treatment, broken lines correspond to the positive feedback treatment.

4. Expectations rules

Before investigating individual forecasting strategies in the experiment, we consider two benchmarks of expectations, naïve and average price expectations.

4.1. Two simple benchmark rules

Naïve expectations means that all participants use the last observation as their forecast, i.e.

$$p^e_{ht} = p_{t-1} \tag{5}$$

and *average price expectations* refers to the case when all participants take the average price as forecast, i.e.[12]

[12] Note that these average price expectations are equivalent with participants that believe that the price is a constant plus noise (which is correct in equilibrium) and in every period estimate this constant by an (unweighted) least squares regression on past prices.

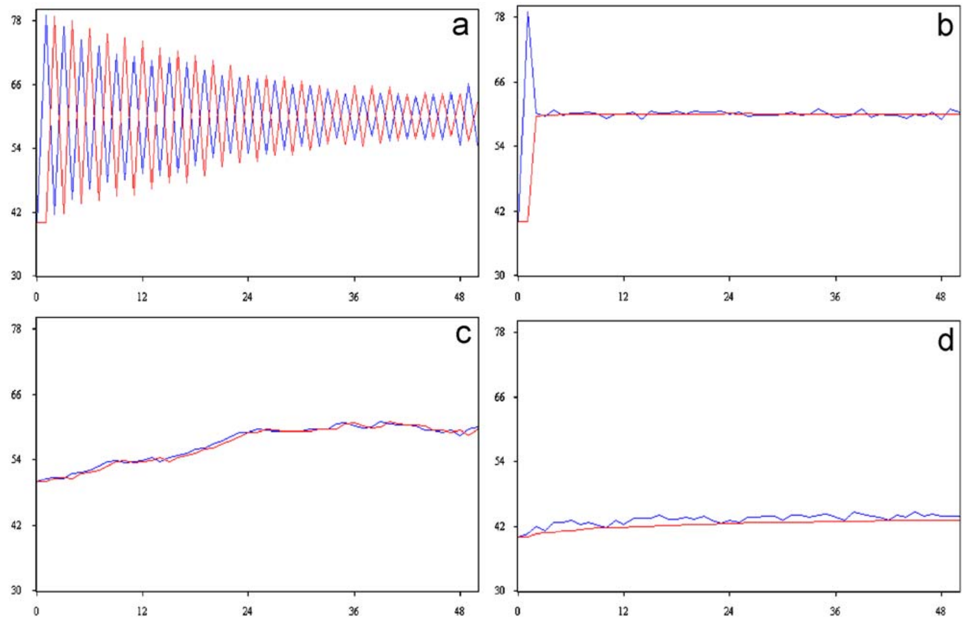

Fig. 5. Realized market prices (blue) and predictions (red) for both treatments in the benchmark cases when all agents have naïve expectations (left panels) and average price expectations (right panels): (a) negative feedback, naive expectations; (b) negative feedback, average expectations; (c) positive feedback, naive expectations; and (d) positive feedback, average expectations.

$$p^e_{ht} = \frac{1}{t}\sum_{j=0}^{t-1} p_j. \tag{6}$$

Fig. 5 shows the realized market prices under naïve and average price expectations for both treatments. This figure suggests that naïve and average price expectations are *not* consistent with the laboratory experiments in Figs. 2 and 3. For the negative feedback treatment with naïve expectations (top left panel) convergence to the equilibrium price is much slower than in the experiments, and the persistent up and down price oscillations are much more regular than what has been observed in the experiments. Moreover, in contrast to the experiments the naïve forecast is systematically wrong as it is high (low) when realized prices turn out to be low (high). Average price expectations in the negative feedback treatment (top right panel) lead to an extremely quick convergence to the equilibrium price, in fact already after two periods, which is faster than in the experiment. For the positive feedback treatment with naïve expectations (bottom left panel) slow monotonic convergence to the equilibrium price occurs. The key difference with the experiments is the absence of (slow) oscillatory movements under naïve expectations. The positive feedback treatment with average expectations (bottom right panel) shows an extremely slow movement in the direction of the fundamental equilibrium, and, in contrast to the experiments, the market price stays far below 60 for all 50 time periods. From these simple simulations we conclude that naïve and average expectations lead to quite different results for both treatments. None of these benchmark forecasting rules gives a good description of individual forecasting in the experiments.

4.2. Individual prediction strategies

Linear prediction rules, with three lags in prices and expected prices, of the form

$$p^e_{h,t} = c + \sum_{i=1}^{3} o_i p_{t-i} + \sum_{i=1}^{3} s_i p^e_{h,t-i} + v_t \tag{7}$$

have been estimated for each individual participant. Predictions of 71 out of 78 participants could be described successfully (i.e. without autocorrelation in the residuals) this way (see Appendix C, Tables C1 and C2).[13] The mean squared error (MSE) of the estimated linear forecasting rules has mean 1.15 for the negative feedback treatment and mean 0.66 for the positive feedback treatment, which would correspond to earnings of 1282 points for the negative feedback treatment and 1269

[13] In the estimations we have excluded a short initial learning phase (see footnote 8). We have also excluded a few time periods because of outliers in individual forecasts. These outliers are likely to be due either to typing errors or to experimentation of participants. Incorporating them into the estimations would give a biased image of the prediction strategies actually used by the participants. The excluded periods are the following (see Figs. 2 and 3). For the six negative feedback groups: {44}, {8,…,12}, {36,45,…,50}, ∅, {21,41} and {8,…,15,22}, respectively. For the seven positive feedback groups: ∅, ∅, ∅, {39}, {7,…,16, 28,34,45}, {15,25,35} and ∅, respectively.

points for the positive feedback treatment, both close to the maximum of 1300 points (= 0.5 euro) per time period. These results suggest that participants use linear forecasting rules based on recent information to form predictions. The time series of individual forecasts already suggested that individual forecasts are close to the equilibrium price level for the negative feedback treatment (see Fig. 2) and fluctuate slowly around the equilibrium level for the positive feedback treatment (see Fig. 3). For the estimated individual linear forecasting rules (7) this can be checked, by looking at the corresponding long run equilibrium price level

$$\hat{p} = \frac{c}{1 - \sum_{i=1}^{3} o_i - \sum_{i=1}^{3} s_i}. \tag{8}$$

Fig. 6 shows frequency distributions of the long run equilibrium price level $\hat{p}$ in (8), corresponding to the estimated individual forecasting rules for the negative feedback treatment (top panel) and the positive feedback treatment (bottom

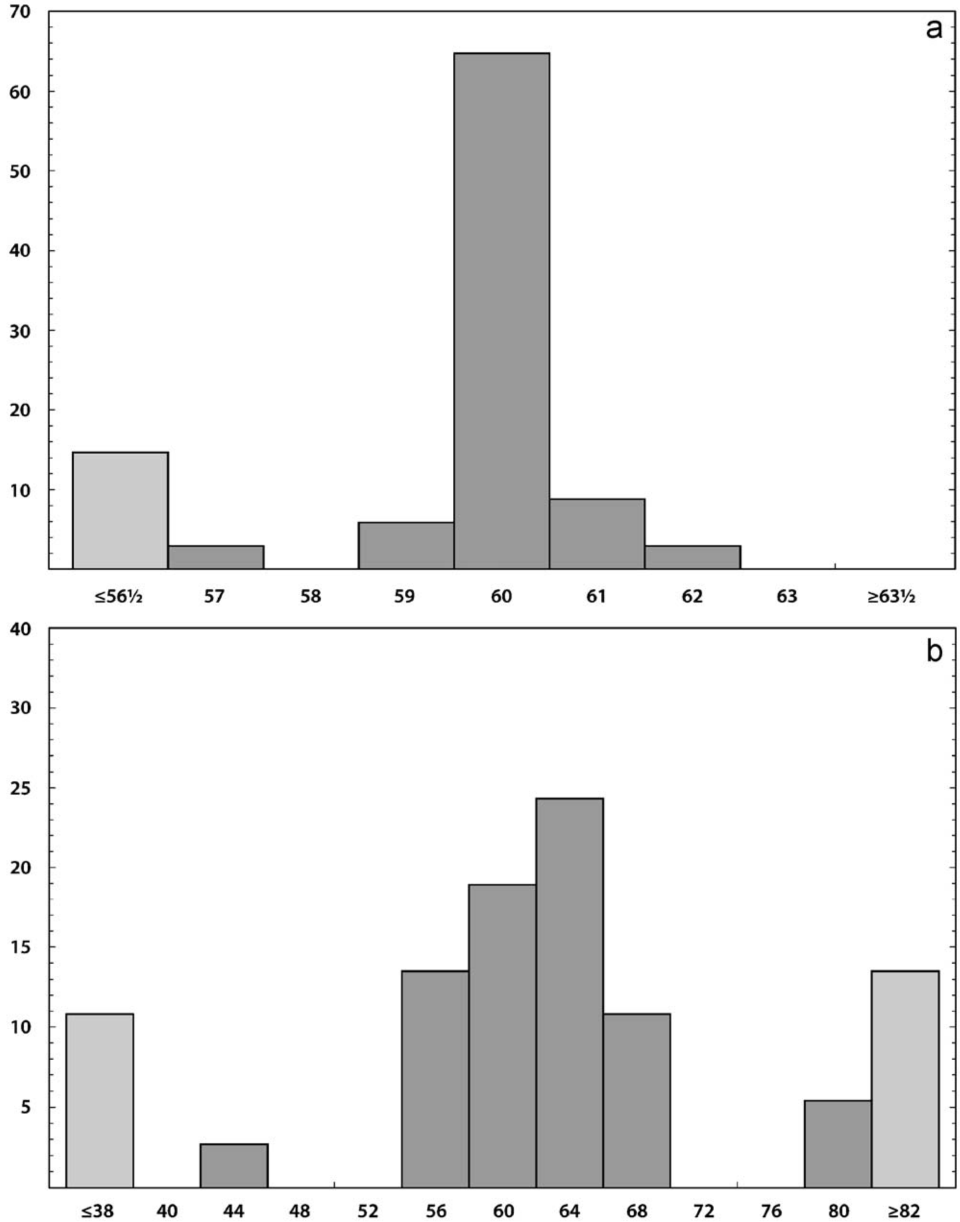

Fig. 6. Frequency distribution of long run equilibrium price level $\hat{p}$ in (8), corresponding to estimated individual forecasting rules for negative feedback treatment (top panel) and positive feedback treatment (bottom level). For the negative feedback treatment 34 forecasting rules of Table C1 (excluding 2 rules with residual autocorrelations) have been used, and more than 60% has an equilibrium level very close to the true equilibrium price 60. For the positive feedback treatment 37 forecasting rules of Table C2 (excluding 5 rules with residual autocorrelations) have been used, and most rules have an equilibrium level fairly close to the true equilibrium price 60.

Please cite this article as: Heemeijer, P., et al., Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. Journal of Economic Dynamics and Control (2009), doi:10.1016/j.jedc.2008.09.009

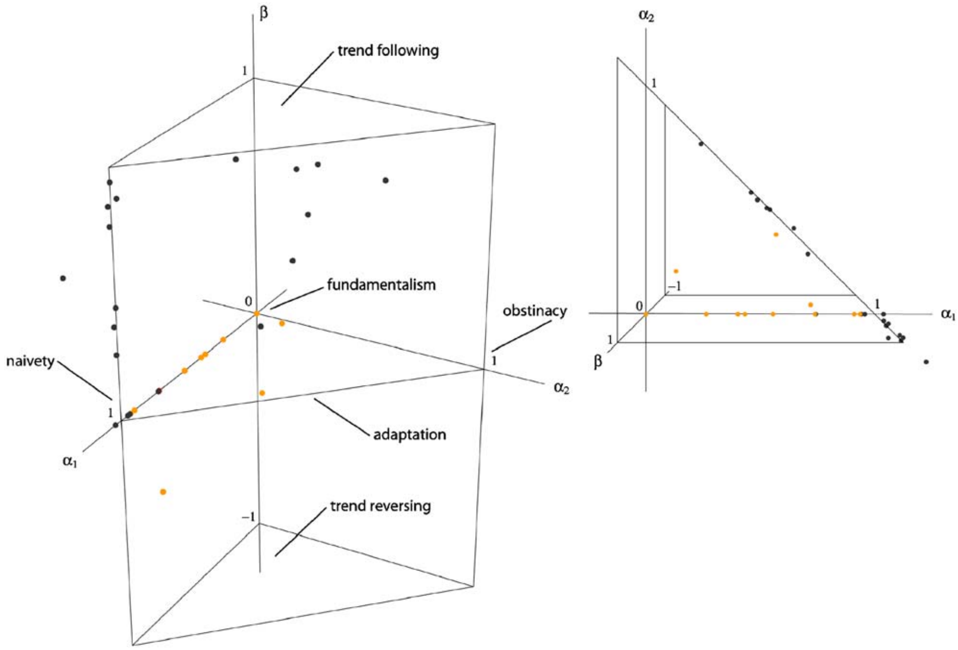

Fig. 7. Prism of first-order heuristics containing the parameter vectors of the prediction rules of the form $p^e_{h,t} = \alpha_1 p_{t-1} + \alpha_2 p^e_{h,t-1} + (1-\alpha_1-\alpha_2)60 + \beta(p_{t-1}-p_{t-2}) + v_t$. The smaller graph on the right is a top-down view of the prism. Yellow dots depict prediction rules from participants in the negative feedback treatment and black dots depict rules from participants in the positive feedback treatment. Positive (negative) values of β correspond to a trend following (trend reversing) prediction rule. The special cases "naivety", "fundamentalism" and "obstinacy" correspond to $p^e_{h,t} = p_{t-1}$, $p^e_{h,t} = 60$ and $p^e_{h,t} = p^e_{h,t-1}$, respectively. Finally, "adaptation" refers to a prediction rule of the form $p^e_{h,t} = \alpha p_{t-1} + (1-\alpha)p^e_{h,t-1}$, with $0<\alpha<1$.

panel). For the negative feedback treatment 34 forecasting rules from Table C1 (excluding the 2 rules with residual autocorrelations) have been used, and more than 60% of these have an equilibrium level $\hat{p}$ very close (within distance 0.5) to the fundamental equilibrium price $p^* = 60$. The mean equilibrium level is 56.90, with standard deviation 9.79; the null hypothesis that the equilibrium level equals $\hat{p} = 60$ cannot be rejected (p-value 0.078). For the positive feedback treatment 37 forecasting rules from Table C2 (excluding the 5 rules with residual autocorrelations) have been used, and most rules have an equilibrium level fairly close to the fundamental equilibrium price 60. The mean equilibrium level is 64.97, with a much higher standard deviation of 30.34. Also for the positive feedback treatment the null hypothesis that the equilibrium level is $\hat{p} = 60$ cannot be rejected (p-value 0.332).

Apparently, in both treatments individuals are able to learn the fundamental, long run equilibrium level $p^* = 60$, for example by observing that the average price is close to 60. This motivates a simpler linear individual forecasting rule

$$p^e_{h,t} = \alpha_1 p_{t-1} + \alpha_2 p^e_{h,t-1} + (1-\alpha_1-\alpha_2)60 + \beta(p_{t-1}-p_{t-2}) + v_t. \tag{9}$$

This rule has a simple *behavioral* explanation, and is in fact an *anchoring* and *adjustment heuristic* as in Tversky and Kahneman (1974). The forecast rule uses a (time varying) *anchor* or reference point, given by a weighted average of the last observed price p_{t-1}, the last individual forecast $p^e_{h,t-1}$ and the long run equilibrium level 60, and extrapolates the last price change $\beta(p_{t-1}-p_{t-2})$ from there. Since the rule (9) is based upon one lag in price, forecast and price change, we refer to it as the "*first order heuristic*". For 40 of the 78 participants the first order heuristic can be successfully estimated, and the estimated parameters are given in Table C3 in Appendix C and are represented graphically in Fig. 7.

Concerning individual prediction strategies, there are two important differences between the positive and negative feedback treatments. In the case of positive feedback (21 prediction rules, black dots in Fig. 7), the trend parameter β is typically highly significant (15 rules) and strongly positive (>0.5, for 11 rules). In the negative feedback case, only 3 rules have a significant β, one slightly positive (0.06) and two negative values (-0.44 and -0.38, thus acting as contrarians). Trend following strategies are thus more important in the case of positive feedback, while some individuals behave as contrarians in the negative feedback case. The second difference between the treatments is that, in the case of positive feedback, the sum $\alpha_1+\alpha_2$ is typically close to 1 (>0.9 for 19 rules), implying that the equilibrium level $p^* = 60$ gets little weight in individual forecasting. In contrast, for the negative feedback treatment, only for 3 rules $\alpha_1+\alpha_2>0.9$, and typically it is much smaller, implying that the equilibrium price level 60 gets high weight in individual forecasting.[14]

[14] From Table C3 it can be seen that for 8 individual rules in the negative feedback treatment, all estimated coefficients α_1, α_2 and β are 0, so that the individual forecasts in (9) reduce to a constant forecast 60. Comparing these results to Table C1, it can be seen that for 7 of these individuals the corresponding estimated linear rules in (10) reduce to a constant close to 60.

Summarizing, one can say that in the positive feedback environment participants tend to base their prediction on a weighted average of the last price and the last prediction, and extrapolate trends in past prices from there ($\alpha_1, \alpha_2 > 0$, $\alpha_1 + \alpha_2$ close to 1 and $\beta > 0$), without taking the equilibrium price into account. Individual forecasting in the positive feedback treatment may thus be described as *naïve and adaptive trend following*. On the other hand, most of the estimated prediction rules from the negative feedback treatment (19 prediction rules, yellow dots in Fig. 7) lie along the positive α_1-axis (α_2 and β close to 0), implying that typically predictions in that treatment are a weighted average between the last observed price and the equilibrium price level. Individual forecasting in the negative feedback treatment may thus be described as *adaptive-average price expectations*.

5. Discussion

Neo-classical economic theory assumes that individuals form expectations rationally (Muth, 1961; Lucas, 1972). This implies that on average market participants make correct price forecasts, and that prices quickly converge to their market clearing equilibrium values, thereby leading to an efficient allocation of resources (Fama, 1970, 1991). However, large fluctuations of prices on financial markets have fueled the debate whether this is indeed a good description of economic behavior (Shiller, 1981; De Bondt and Thaler, 1985; Garber, 1990). Our findings show that whether rational expectations (i.e. when *all* agents have rational expectations) gives a good description of aggregate market behavior depends upon the underlying expectations feedback structure. In fact, commonly observed differences between experimental commodity and financial markets (Smith, 1962; Smith et al., 1988; Gode and Sunder, 1993; Lei et al., 2001; Hommes et al., 2005b, 2007, 2008) may be attributed to a large extent to this feedback structure. Prices in commodity markets where speculative demand only plays a limited role will be much more stable and closer to the equilibrium value when the product (and production technology) has been around for a while, and commodity prices can fluctuate wildly only for relatively new products (e.g. computer chips, see The Economist, 1996a, b, 2001). The fact that some established commodity markets regularly exhibit fluctuations is consistent with our conclusions, since these fluctuations may be attributed to the presence of demand-driven speculators (Cashin et al., 2002; Canoles et al., 1998). On the other hand, some features of actual financial markets, such as the presence of contrarians or certain programmed trading strategies might introduce elements of negative feedback in asset markets.[15] Typically, however, due to a dominating positive feedback structure driven by speculative demand, financial markets can easily diverge from the equilibrium price and be relatively unstable and excessively volatile. Of course real markets are much more complicated than our simple, stylized experimental setup. But what our experiments show is that the difference, positive versus negative, in expectations feedback structure *alone* may significantly affect price stability and market volatility.

What causes these differences in aggregate outcomes? Two important differences between the treatments have been observed: (i) in the positive feedback treatment market prices are characterized by *slow oscillatory movements*, while in the negative feedback treatment market prices converge to the steady state $p^* = 60$, within about ten periods and (ii) in the positive feedback treatment *coordination of individual forecasting rules* is extremely *quick*, in fact within 3 periods, while in the negative treatment heterogeneity in forecasts persist for about 10 periods, and coordination only occurs after the market price has converged to its steady state. Our estimation results of individual forecasting rules clearly provides an explanation for the oscillatory price movements in the positive feedback treatment. In a market with positive feedback, trend following behavior is much more prominent than in a market with negative feedback. In markets dominated by positive feedback, when a majority of individuals uses a trend following strategy, other individuals have an incentive to use such a strategy too, thus reinforcing trends in prices. In contrast, in markets dominated by negative feedback when many individuals adopt trend following strategies, it becomes more profitable for other agents to go against the trend. As a consequence, trend following strategies do *not* survive in a market dominated by negative feedback, but play an important role in markets dominated by positive feedback, such as speculative asset markets. In particular, the survival of trend following strategies contributes to oscillatory price movements, persistent deviations from the fundamental benchmark and excess volatility.

Explaining the differences observed in the first 10 periods of the experiments is more difficult, but the realized market prices in Figs. 2 and 3 suggest the following tentative explanation. In the first period of an experiment, participants make heterogeneous forecasts, leading to some average forecast, say around 50. In the positive feedback treatment, the realized market price will then be fairly close to, but slightly above the average forecast, i.e. slightly above 50. In the second period, if most participants put large weight to the observed market price (i.e. their forecasts are close to naïve expectations), this simple but plausible forecasting behavior will already lead to coordination of forecasts, say close to 50, and a second realized market price close to, but slightly above the first observation. After this quick coordination of forecasts, realized market prices move slowly and, as our estimation of individual forecasting rules have shown, trend following behavior may reinforce slow oscillatory price movements. The initial stage of an experiment in the negative feedback treatment is quite different. Assuming heterogeneous forecasts in the first period with an average forecast of say 50, due to the negative feedback, market price will be pushed to the other side of the steady state price 60, leading to a realized market price close to 70. In the second period, when most participants put large weight to the observed market price (i.e. their forecast is close to naïve expectations), once more due to the negative feedback the next realized market price will be on the other side, that

[15] Note however that our analysis in Section 4 revealed that only very few participants can actually be classified as contrarian, as opposed to a substantial number of trend followers. The latter therefore seems to be the more relevant behavioral strategy.

Please cite this article as: Heemeijer, P., et al., Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. Journal of Economic Dynamics and Control (2009), doi:10.1016/j.jedc.2008.09.009

is, below the steady state price 60, with a value of say somewhat above 50. Quick coordination of individual forecasts therefore is unlikely in the case of negative feedback. Due to the negative feedback, prices fluctuate considerably in the initial phase of the experiments and, as shown in Fig. 2 individual forecasts remain *heterogeneous* for a while. Naïve expectations and trend following rules perform bad in these circumstances, and it is more likely that participants adopt some average or adaptive expectations rule. Fig. 2 indeed shows that individual forecasts become less extreme and move in the direction of the average price, thus stabilizing price fluctuations and leading to convergence to the steady state 60 after about 10 periods.

At this point it is useful to stress some similarities and differences between our experiments and those of Fehr and Tyran (2008). Fehr and Tyran emphasize the impact of the strategic environment in games on the aggregate outcome. They compare games with strategic complements (positive feedback) and strategic substitutes (negative feedback), (see also Camerer and Fehr, 2006 for an extensive recent discussion). Although in their experiment participants have to form expectations about the average prices of the other players their main task (and the only one for which they are rewarded) is to set prices, whereas in our experiment participants only have to submit forecasts of the actual price. Nevertheless, the prices chosen by participants in the Fehr and Tyran experiments are consistent with the price expectations they submitted. Another important difference concerns the information available to participants. In their setting participants have information about other participants' past average forecasts. Moreover, a payoff table gives the best reply to the expected average forecasts of others. In our experiments participants do *not* have any information about other participants' forecast, nor do they know what exactly the best reply is. Participants only observe realized market prices and, for example, do not even know exactly how many other participants affect the realized price. Despite the limited information, similar results are obtained: quick convergence in the case of negative expectations feedback and slow convergence or oscillating prices and significant excess volatility in the case of positive feedback. Fehr and Tyran (2008, p. 347) attribute the differences in aggregate outcome between their different strategic environments to "stickiness of price expectations". A novel feature of our analysis is the estimation of individual forecasting strategies, which shows that in our experiments individual expectations are indeed important in explaining the differences between our two different market environments. In particular, the fact that trend following strategies perform well in the positive feedback treatment is key in explaining the different aggregate outcomes.[16]

A final major difference[17] with the experiments of Fehr and Tyran (2008) is that the former have one large anticipated (negative and permanent) monetary shock,[18] whereas we consider a series of small random shocks, which do not change the location of the steady state. In the pre-shock phase in Fehr and Tyran (2008) there is little evidence for deviations from the equilibrium for both treatments whereas we find quick convergence under negative feedback and slow convergence and oscillations in the positive feedback treatment. We conjecture that this difference in findings between the experiments may be attributed to the local stability properties of the steady state. In the neighborhood of the steady state the slope of the best-reply function is zero in the experiments of Fehr and Tyran, whereas the slope of the price generating mechanism is close to one (in absolute value) in our experiments. In fact, Sonnemans and Tuinstra (2008) discuss a positive feedback treatment with a smaller slope of the price generating mechanism (2/3 instead of 20/21) and find that in that case prices also converge quickly to the steady state.

In the post-shock phase in Fehr and Tyran (2008) there is a difference between treatments with respect to the speed of adjustment towards the new steady state, with much more inertia when actions are strategic complements. An interesting additional treatment for our experiment would be to have a similar large permanent shock, displacing the location of the steady state, in order to investigate whether convergence to the new steady state equilibrium is slower in the positive feedback treatment than in the negative feedback treatment. For these experiments we could use a lower value of the derivative (e.g. 2/3) in order to be sure that there will indeed be convergence of prices to the steady state in both treatments eventually. We leave this interesting additional treatment for future work.

The results from our experiments are consistent with noise trader models (DeLong et al., 1990a, b) and herding models in finance (Kirman, 1993; Lux, 1995; Brock and Hommes, 1998; De Grauwe and Grimaldi, 2006), showing that simple trend following strategies may explain investors' behavior in financial markets. The estimated individual forecasting rules suggest an explanation for the differences between the treatments. Models with *heterogeneous expectations* and *evolutionary* and *reinforcement* learning, e.g. as in Brock and Hommes (1997) and Branch (2004), may explain these different outcomes in markets with positive and negative feedback. Only in an environment where positive feedback is sufficiently strong, trend following strategies have a chance of surviving evolutionary competition and agents may coordinate on trend following strategies. We leave the development of a learning model to explain these and related experiments for future research.

[16] Note that where we identify an "anchor and adjustment heuristic", also Fehr and Tyran argue that "anchoring" can explain part of their results. In their case the pre-shock steady state is used as an anchor for the post-shock price, in our case the (time-varying) anchor corresponds to an estimate of the unknown but fixed steady state.

[17] Other differences between the experiments of Fehr and Tyran (2008) are the number of participants per game/market (4 versus 6), the number of available strategies (30 versus 10,000 prices, between 0 and 100 in two decimals) and the length of the experiment (30 versus 50 periods).

[18] Fehr and Tyran (2001) show, in an environment where actions are strategic complements, that price adjustment is much quicker for a *positive* monetary shock than for a negative monetary shock.

Please cite this article as: Heemeijer, P., et al., Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. Journal of Economic Dynamics and Control (2009), doi:10.1016/j.jedc.2008.09.009

Acknowledgment

Earlier versions of this paper have been presented at the ESA 2004 North American Regional Meeting, Tucson, USA, November 4–7, 2004, the third EU *Complex Markets* workshop, Cagliari, Italy, February 21, 2006, the ESA 2006 European Regional Meeting, Nottingham, UK, September 7–10, 2006, the ECB-conference *Monetary Policy, Asset Markets and Learning*, Frankfurt, Germany, November 6–7, 2006 and seminars at the University of Antwerp, University of Amsterdam, University of Copenhagen, Scuola Superiore Sant'Anna, Pisa, and Université Strasbourg. We thank all participants, especially John Duffy and Oliver Kirchkamp, as well as two referees and the editor of this special issue for stimulating discussions and helpful comments. We thank the Dutch Science Foundation (NWO) and the EU, through a STREP grant, for financial support.

Appendix A. Expectations feedback models

This appendix describes the cobweb and asset pricing model underlying the negative and positive feedback systems. In both treatments, the price adjustment rule is of the general form

$$p_t = p_{t-1} + \lambda \sum_{h=1}^{H} ED(p_{t-1}, p^e_{h,t}), \tag{A.1}$$

where the excess demand function ED depends upon the market price p_{t-1} and individual forecasts $p^e_{h,t}$ and λ is the speed of adjustment. In real markets, the price adjustment rule (A.1) depends upon lagged market price p_{t-1} and individual forecasts $p^e_{h,t}$. In the experiments however, parameters have been fixed in such a way that the dependence of p_t on p_{t-1} cancels out, in order to keep the setting as simple as possible and to isolate the effect of the expectations feedback. In the experiments therefore, (A.1) is of an even simpler form

$$p_t = f(\overline{p^e_t}), \tag{A.2}$$

where $\overline{p^e_t}$ is the *average price forecast* of all participants in the market. Since we assume demand and supply to be linear, in the experiments the map f in (A.2) is linear, with positive, respectively, negative slope.

A.1. Negative feedback: the cobweb framework

The market with negative expectations feedback used in the experiment is based on the classical cobweb or hog cycle model. The cobweb model has served as a benchmark framework to study various expectations schemes, e.g. naïve expectations (Ezekiel, 1938), adaptive expectations (Nerlove, 1958), rational expectations (Muth, 1961), OLS learning (Bray and Savin, 1986), genetic algorithm learning (Arifovic, 1994) and heterogeneous expectations (Brock and Hommes, 1997). Originally the model was mainly applied to agricultural commodities such as corn or hogs, but later it has e.g. also been used as a model for the demand for lawyers, engineers (Freeman, 1975, 1976) and public school teachers (Zarkin, 1985) and the demand for oil (Krugman, 2001). The key feature of the model is a fixed production lag, so production decisions by price-taking firms are based on a forecast of the market price in the next time period. Let $D(p_t)$ be a nonnegative and monotonically decreasing demand function and let $S(p^e_{h,t})$ be the nonnegative supply function of firm h, derived from expected profit maximization. Supply of firm h in the next period depends upon the expected price, $p^e_{h,t}$, for that period by that firm. The second order condition for profit maximization implies that S is a non-decreasing function. Moreover, we assume that all firms have the same supply function. In the model equality between demand and supply in each period is not required, but in each period the market price is adjusted in the direction of the excess demand, the trade gap itself being absorbed by a hypothetical market maker.[19] An increase of the price forecast leads to increasing production, a decreasing excess demand and therefore a lower market price. The following price-adjustment formula was used:

$$p_t = p_{t-1} + \lambda\left(D(p_{t-1}) - \sum_{h=1}^{H} S(p^e_{h,t})\right) + \varepsilon_t. \tag{A.3}$$

The expression between brackets represents excess demand, where the market maker uses demand in period $t-1$ as a proxy for the (unknown) demand in period t. We assume there are H suppliers, only differing in the way they form

[19] Our presentation of the cobweb model differs from the classical textbook treatment with market clearing. Here we present the model with a market maker using a price adjustment rule based upon excess demand. The market maker price adjustment rule has often been used in an asset pricing framework with positive feedback (see the next subsection). In the experimental instructions for both the positive and negative feedback treatments the qualitative information about the market and price formation is formulated in terms of price adjustment according to excess demand, to ensure that the information in the two treatments is the same, with only symmetrically opposite differences in the type of feedback. Our presentation of the cobweb model thus gives a mathematical representation closely following the participants' qualitative information. By choosing a suitable adjustment coefficient λ, the lagged price p_{t-1} drops from the price adjustment rule (A.3), and the model (A.4) in the experiments is equivalent to the classical cobweb model with the price p_t determined by market clearing.

expectations. In the laboratory experiment we take $H = 6$. The market maker adjusts the market price proportionally to the excess demand, with a positive price adjustment coefficient λ. The term ε_t is a random term, representing e.g. small uncertainties in demand; in the laboratory experiment $\varepsilon_t \sim N(0, 1/4)$. In the cobweb treatment in the experiment, the participants act as advisors to the producers in the market, so their individual price expectations $p^e_{h,t}$ determine aggregate supply. The demand function is taken to be linearly decreasing, $D(p_t) = a - bp_t$, and the supply function is linearly increasing, $S(p^e_{h,t}) = sp^e_{h,t}$. For the laboratory experiments we chose parameters $s = 1/6$ and $b = 21/20$, $\lambda = 1/b = 20/21$ and $a = 123$. The (unknown) price adjustment rule then becomes:

$$p_t = \frac{20}{21}(123 - \overline{p}^e_t) + \varepsilon_t, \tag{A.4}$$

with $\overline{p}^e_t = \frac{1}{6}\sum_{h=1}^{6} p^e_{h,t}$ the average prediction of the six participants in the experimental market.

A.2. Positive feedback: an asset pricing framework

The second market in our laboratory experiment consists of a standard asset-pricing model (see e.g. Cuthbertson, 1996; Campbell et al., 1997; Brock and Hommes, 1998 for detailed treatments). Demand of speculators for a risky asset depends positively upon the asset's expected price increase. Investors can either invest in a risk-free asset (e.g. a government bond) with a fixed gross return 1+r, or invest in a risky asset (e.g. a stock) paying an uncertain dividend y_t in each period. In this model, a market maker adjusts prices according to excess demand, that is, excess demand for the risky asset leads to an increase in the asset price, and excess supply to a decrease. This market maker price adjustment rule has become one of the standard approaches for modeling price formation in asset markets, see e.g. Beja and Goldman (1980), Day and Huang (1990) and Lux (1995). Since in the experiment demand was taken to be an increasing function of the price predictions, as is natural in the case of a stock or some other financial asset, our asset pricing treatment is driven by positive expectations feedback. As in the cobweb model above, a market maker adapts prices in proportion to excess demand and the actual development of prices is given by (see e.g. Chiarella et al., 2002; Hommes et al., 2005a for more details)

$$p_t = p_{t-1} + \lambda\left(\sum_{h=1}^{H} \frac{E_{h,t}(p_t + y_t - (1+r)p_{t-1})}{a\sigma^2} - z^s\right) + \varepsilon_t. \tag{A.5}$$

Excess demand is given by aggregate demand, consisting of the sum of individual demand functions of all traders, minus the (constant) supply z^s of shares of the risky asset. The demand function is standard and obtained from mean-variance maximization of next period's expected wealth, where a is the risk aversion coefficient, $E_{h,t}(p_t + y_t - (1+r)p_{t-1})$ is the belief of trader h about excess return, and σ^2 is the belief about the conditional variance of excess returns, here assumed to be constant for all periods and all traders. Each participant in the experiment acts as an advisor to a trader (say a large pension fund) informing them of their price prediction $p^e_{h,t} = E_{h,t}(p_t)$, which the trader then uses to calculate her demand function. Six participants form a market in the experiment, i.e. $H = 6$.

To achieve symmetry between the positive and negative feedback treatment, we choose parameters such that the treatments only differ in the sign of the expectations feed-back. Fixing parameters $r = 1/20$, $a\sigma^2 = 6$, $z^s = 1$, $\lambda = 20/21$ and assuming $E_{h,t}(y_t) = 3$ for all h and t we get, for the positive feedback treatment, the price adjustment rule:[20]

$$p_t = \frac{20}{21}(\overline{p}^e_t + 3) + \varepsilon_t, \tag{A.6}$$

where as before, $\overline{p}^e_t$ defines average price expectations over all participants, and $\varepsilon_t \sim N(0, 1/4)$ is a random term, representing e.g. small random fluctuations in the supply of the risky asset.

Appendix B. Experimental instructions

During each of the four experimental sessions, a short welcoming message was read aloud from paper, after which the participants were randomly assigned to a cubicle in the computer lab. In each cubicle there was a computer, some experimental instructions on paper and some blank paper with a pen. The two treatments had similar instructions that differed only in their description of the market environment. When all the participants were seated, they were asked to read the instructions on their desks. After a few minutes, they were given the opportunity to ask questions regarding the instructions, after which the experiment started. When the 50 time periods were completed, the participants were asked to remain seated and fill in the questionnaire, which was subsequently handed out to them. After a reasonable amount of

[20] As in the cobweb treatment, for this choice of the parameters, the lagged price p_{t-1} drops from the price adjustment rule (A.5), so that for both treatments the price p_t is determined by average market expectations.

time, the participants were called to the ante-room one by one to hand in the questionnaire and receive their earnings, in cash. The participants left the computer lab after receiving their earnings.

The experimental instructions the participants read in their cubicles consisted of three parts, totalling five pages. The first part contained general information about the market the experiment was about to simulate, which was of course treatment-specific. The second part contained an explanation of the computer program used during the experiment. The third part displayed a table relating the absolute prediction error made in any single period to the amount of credits earned in that period. The conversion rate between credits and euros, being 2600 credits to 1 euro, was made public by announcement, since it was not listed with the table. The questionnaire after the experiment contained 19 questions, the first 10 of which could be answered only by the integers 1 through 5. The experimental instructions will be translated below.[21]

B.1. Translation of experimental instructions for negative feedback treatment

B.1.1. Experimental instructions

The shape of the artificial market used by the experiment, and the role you will have in it, will be explained in the text below. Read these instructions carefully. They continue on the backside of this sheet of paper.

B.1.2. General information

You are an advisor of an importer who is active on a market for a certain product. In each time period the importer needs a good prediction of the price of the product. Furthermore, the price should be predicted one period ahead, since importing the good takes some time. As the advisor of the importer you will predict the price $P(t)$ of the product during 50 successive time periods. Your earnings during the experiment will depend on the accuracy of your predictions. The smaller your prediction errors, the greater your earnings.

B.1.3. About the market

The price of the product will be determined by the law of supply and demand. The size of demand is dependent on the price. If the price goes up, demand will go down. The supply on the market is determined by the importers of the product. Higher price predictions make an importer import a higher quantity, increasing supply. There are several large importers active on this market and each of them is advised by a participant of this experiment. Total supply is largely determined by the sum of the individual supplies of these importers. Besides the large importers, a number of small importers is active on the market, creating small fluctuations in total supply.

B.1.4. About the price

The price is determined as follows. If total demand is larger than total supply, the price will rise. Conversely, if total supply is larger than total demand, the price will fall.

B.1.5. About predicting the price

The only task of the advisors in this experiment is to predict the market price $P(t)$ in each time period as accurately as possible. The price (and your prediction) can never become negative and always lies between 0 and 100 euros in the first period. The price and the prediction in period 2 through 50 is only required to be positive. The price will be predicted one period ahead. At the beginning of the experiment you are asked to give a prediction for period 1, $V(1)$. When all participants have submitted their predictions for the first period, the market price $P(1)$ for this period will be made public. Based on the prediction error in period 1, $P(1)-V(1)$, your earnings in the first period will be calculated. Subsequently, you are asked to enter your prediction for period 2, $V(2)$. When all participants have submitted their prediction for the second period, the market price for that period, $P(2)$, will be made public and your earnings will be calculated, and so on, for 50 consecutive periods. The information you have to form a prediction at period t consists of: All market prices up to time period $t-1$: $\{P(t-1), P(t-2), \ldots, P(1)\}$; All your predictions up until time period $t-1$: $\{V(t-1), V(t-2), \ldots, V(1)\}$; Your total earnings at time period $t-1$.

B.1.6. About the earnings

Your earnings depend only on the accuracy of your predictions. The better you predict the price in each period, the higher will be your total earnings. The attached table lists all possible earnings.

When you are done reading the experimental instructions, you may continue reading the computer instructions, which have been placed on your desk as well.

[21] We aimed at acquiring the highest degree of symmetry in the instructions for the two treatments. Nevertheless, as one of the referees pointed out, it might be cognitively easier for participants to come up with a price prediction in the negative feedback treatment than in the positive feedback treatment.

Please cite this article as: Heemeijer, P., et al., Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. Journal of Economic Dynamics and Control (2009), doi:10.1016/j.jedc.2008.09.009

B.2. Translation of experimental instructions for positive feedback treatment

B.2.1. Experimental instructions

The shape of the artificial market used by the experiment, and the role you will have in it, will be explained in the text below. Read these instructions carefully. They continue on the backside of this sheet of paper.

B.2.2. General information

You are an advisor of a trader who is active on a market for a certain product. In each time period the trader needs to decide how many units of the product he will buy, intending to sell them again the next period. To take an optimal decision, the trader requires a good prediction of the market price in the next time period. As the advisor of the trader you will predict the price $P(t)$ of the product during 50 successive time periods. Your earnings during the experiment will depend on the accuracy of your predictions. The smaller your prediction errors, the greater your earnings.

B.2.3. About the market

The price of the product will be determined by the law of supply and demand. Supply and demand on the market are determined by the traders of the product. Higher price predictions make a trader demand a higher quantity. A high price prediction makes the trader willing to buy the product, a low price prediction makes him willing to sell it. There are several large traders active on this market and each of them is advised by a participant of this experiment. Total supply is largely determined by the sum of the individual supplies and demands of these traders. Besides the large traders, a number of small traders is active on the market, creating small fluctuations in total supply and demand.

B.2.4. About the price

The price is determined as follows. If total demand is larger than total supply, the price will rise. Conversely, if total supply is larger than total demand, the price will fall.

B.2.5. About predicting the price

The only task of the advisors in this experiment is to predict the market price P(t) in each time period as accurately as possible. The price (and your prediction) can never become negative and lies always between 0 and 100 euros in the first period. The price and the prediction in period 2 through 50 is only required to be positive. The price will be predicted one period ahead. At the beginning of the experiment you are asked to give a prediction for period 1, $V(1)$. When all participants have submitted their predictions for the first period, the market price $P(1)$ for this period will be made public. Based on the prediction error in period 1, $P(1)-V(1)$, your earnings in the first period will be calculated. Subsequently, you are asked to enter your prediction for period 2, $V(2)$. When all participants have submitted their prediction for the second period, the market price for that period, $P(2)$, will be made public and your earnings will be calculated, and so on, for 50 consecutive periods. The information you have to form a prediction at period t consists of: All market prices up to time period $t-1$: $\{P(t-1), P(t-2), \ldots, P(1)\}$; All your predictions up until time period $t-1$: $\{V(t-1), V(t-2), \ldots, V(1)\}$; Your total earnings at time period $t-1$.

B.2.6. About the earnings

Your earnings depend only on the accuracy of your predictions. The better you predict the price in each period, the higher will be your total earnings. The attached table lists all possible earnings.

When you are done reading the experimental instructions, you may continue reading the computer instructions, which have been placed on your desk as well.

B.3. Translation of computer instructions

B.3.1. Computer instructions

The way the computer program works that will be used in the experiment, is explained in the text below. Read these instructions carefully. They continue on the backside of this sheet of paper.

The mouse does not work in this program.

Your earnings in the experiment depend on the accuracy of your predictions. A smaller prediction error in each period will result in higher earnings.

To enter your prediction you can use the numbers, the decimal point and, if necessary, the backspace key on the keyboard.

Your prediction can have two decimal numbers, for example 30.75. Pay attention not to enter a comma instead of a point. Never use the comma. Press enter if you have made your choice.

The better your prediction, the more credits you will earn. On your desk is a table listing your earnings for all possible prediction errors.

Please cite this article as: Heemeijer, P., et al., Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. Journal of Economic Dynamics and Control (2009), doi:10.1016/j.jedc.2008.09.009

For example, your prediction was 13.42. The true market price turned out to be 12.13. This means that the prediction error is: 13.42–12.13 ≈ 1.30. The table then says your earnings are 1255 credits (as listed in the third column [this is a typing error, it should be second column]).

The available information for predicting the price of the product in period t consists of: All product prices from the past up to period $t-1$; Your predictions up to period $t-1$; Your earnings until then.

B.3.2. *The computer screen*

The instructions below refer to this figure.

In the upper left corner a graph will be displayed consisting of your predictions and of the true prices in each period. This graph will be updated at the end of each period.

In the rectangle in the middle left you will see information about the number of credits you have earned in the last period and the number you have earned in total. The time period is also displayed here, possibly along with other relevant information.

On the right hand side of the screen the experimental results will be displayed, that is, your predictions and the true prices for at most the last 20 periods.

At the moment of submitting your price prediction, the rectangle in the lower left side of the figure will appear. When all participants have subsequently submitted their predictions, the results for the next period will be calculated.

When everyone is ready reading the instructions, we will begin the experiment. If you have questions now or during the experiment, raise your hand. Someone will come to you for assistance.

Appendix C. Estimation results of individual forecasting rules

Estimation individual forecasting rules (negative and positive feedback) and estimation of individual first order heuristics are explained in Tables C1–C3, respectively.

Table C1
Estimation individual forecasting rules (negative feedback).

Participant	c	p_{-1}	p_{-2}	p_{-3}	p^e_{-1}	p^e_{-2}	p^e_{-3}	R^2	*AC*	*Eq.*	*MSE*
1	53.67	0.1080	0	0	0	0	0	0.2013	No	60.17	0.425
2	29.75	0.7002	0	0	0	−0.1957	0	0.8795	No	60.04	0.668
3	25.47	0	0.2431	0	0	0	0	0.0983	No	33.65	0.613
4	23.30*	0.4213	0	0	0	0	0	0.1385	No	43.72	0.667
5	32.90*	0.3919	−0.3136	0	0.3750	0	0	0.3077	No	60.18	0.925
6	39.48	0.3255	0.2009	0	−0.5089	0	0.3240	0.6504	No	59.95	0.593
7	87.60	0	0	0	0	−0.1772	−0.2876	0.3478	No	59.80	1.266
8	10.26*	0.0111	0	0	0.0306	0	0	0.1912	No	10.71	1.363
9	32.15	0.0953	0	0	0	0	0.3662	0.7756	Yes	–	0.709
10	29.38	0.2818	0.2317	0	0	0	0	0.2821	No	60.39	0.792
11	16.13	0.2697	0.1532	0	0	0.3088	0	0.4381	No	60.12	0.053
12	20.81	0.6534	0	0	0	0	0	0.5102	No	60.04	1.157
13	−0.489*	0.3003	0.4690	0	0	0.2218	0	0.7600	No	54.94	0.814
14	59.15	0	0	0	0	0	0	0.0000	No	59.15	1.126
15	7.433	0.8692	0	0	0	0	0	0.9412	No	56.83	1.100
16	31.26	0	0.4799	0	0	0	0	0.4220	No	60.10	1.232
17	−170.6	0	0	0	−1.356	1.538	3.671	0.9670	No	59.80	1.105
18	82.00	−0.7656	0.3995	0	0	0	0	0.7943	No	60.02	0.660
19	34.40	0.4264	0	0	0	0	0	0.5653	No	59.97	0.511
20	45.60	0.2423	0	0	0	0	0	0.3077	No	60.18	0.429
21	60.00	0	0	0	0	0	0	0.0000	No	60.00	0.292
22	20.97	0.6489	0	0	0	0	0	0.7385	Yes	–	0.519
23	16.56	0.3326	0.3946	0	0	0	0	0.2316	No	60.70	0.543
24	60.00	0	0	0	0	0	0	0.0000	No	60.00	0.292
25	44.31	0.2653	0	0	0	0	0	0.2074	No	60.31	2.867
26	23.03	−0.2041	0.4658	0	0.3586	0	0	0.6671	No	60.65	1.950
27	60.98	0	0	0	0	0	0	0.0000	No	60.98	3.960

Please cite this article as: Heemeijer, P., et al., Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. Journal of Economic Dynamics and Control (2009), doi:10.1016/j.jedc.2008.09.009

Table C1 (*continued*)

Participant	c	p_{-1}	p_{-2}	p_{-3}	p^e_{-1}	p^e_{-2}	p^e_{-3}	R^2	*AC*	*Eq.*	*MSE*
28	58.89	0	0	0	0	0	0	0.0000	No	58.89	2.926
29	45.38	0	0	−0.0898	0	0	0	0.5408	No	59.85	0.805
30	5.533*	0.9115	0	0	0	0	0	0.7367	No	52.52	6.158
31	5.767*	0.5157	0	0	0	0.3906	0	0.7284	No	61.55	0.687
32	27.21	0.4251	0.1179	0	0	0	0	0.6324	No	59.54	0.588
33	90.46	0	0	0	0	0	−0.5047	0.2533	No	60.12	0.482
34	59.66	0	0	0	0	0	0	0.0000	No	59.66	0.469
35	45.71	0.2338	0	0	0	0	0	0.2004	No	59.66	0.656
36	60.48	0	0	0	0	0	0	0.0000	No	60.48	0.910

Prediction rules for the 36 participants of the negative feedback treatment (least squares estimation of Eq. (7)). The first column shows participants' numbers, clustered according to group; the second through eighth column show estimates of the coefficients; the ninth and tenth columns show the R-squared statistic and report on autocorrelation in the residuals up to the 20th order (Ljung–Box Q-statistics, 5% level); the last two columns contain the long run equilibrium price and the mean squared error of the estimated rule. Insignificant variables were eliminated one by one, largest p value first, until all p values were below 5%. An asterisk in the second column indicates that the constant is insignificant.

Table C2
Estimation individual forecasting rules (positive feedback).

Participant	c	p_{-1}	p_{-2}	p_{-3}	p^e_{-1}	p^e_{-2}	p^e_{-3}	R^2	*AC*	*Eq.*	*MSE*
1	−0.790*	1.675	0	−0.4329	−0.2324	0	0	0.9965	No	81.44	0.622
2	−0.682*	1.340	−0.5007	0	0.4642	−0.2914	0	0.9980	No	56.36	0.708
3	−1.176*	1.724	0	−0.3995	−0.3069	0	0	0.9932	No	66.82	0.738
4	−1.121	1.893	−0.8748	0	0	0	0	0.9971	No	61.59	0.518
5	0.417*	1.443	−0.8745	0	0.4264	0	0	0.9975	No	81.76	0.529
6	−0.817*	1.787	−0.7724	0	0	0	0	0.9982	No	55.96	0.538
7	0.742*	1.184	0	−0.1698	0	0	0	0.9964	Yes	–	0.550
8	−0.179*	1.463	−0.4552	0	0	0	0	0.9938	No	22.95	0.536
9	0.657*	1.220	−0.7315	0	0.5006	0	0	0.9969	No	60.28	0.477
10	0.339*	1.285	0	0	0	−0.2887	0	0.9969	No	91.62	0.381
11	0.693*	1.368	−0.8523	0	0.4743	0	0	0.9948	No	69.30	0.472
12	0.223*	1.851	0	0	−0.3270	−0.3533	−0.1723	0.9926	No	139.4	0.521
13	0.040*	1.450	−0.4504	0	0	0	0	0.9870	No	100.0	0.545
14	0.164*	1.069	−0.4708	0	0.4000	0	0	0.9943	No	91.11	0.460
15	−0.251*	1.275	−0.2989	−0.2706	0	0.2984	0	0.9981	No	64.36	0.649
16	2.170*	1.232	0	0	0	−0.2662	0	0.9780	No	63.45	1.136
17	−0.985*	1.251	0	−0.2345	0	0	0	0.9900	No	59.70	0.452
18	−0.1026	1.219	−0.5430	0	0.4372	0	0	0.9942	No	−0.07	0.417
19	2.411	1.084	0	0	0.2635	0	−0.3910	0.9940	No	55.43	0.880
20	1.956*	−0.9115	0	0	0	0	0	0.8975	No	1.02	0.875
21	1.382	1.641	−0.9729	0	0.3084	0	0	0.9978	No	58.81	0.515
22	2.687	1.6274	−0.4900	0	0	0	−0.1816	0.9934	No	60.79	0.683
23	1.475	1.441	0	−0.4659	0	0	0	0.9948	No	59.24	0.858
24	0.062*	1.943	−0.9439	0	0	0	0	0.9953	No	68.89	0.977
25	34.27	0	0.1203	0	0.3421	0.2670	−0.3179	0.9892	Yes	–	1.578
26	173.7*	0	0	0	0	0	0	0.0000	No	173.7	0.780
27	2.601	1.000	0	−0.1972	−0.0384	0.1215	0.0682	1.0000	Yes	–	1.000
28	4.160	1.005	0	0	0	−0.1025	0	0.9981	No	42.67	0.705
29	15.71	1.004	0	0.5544	−0.2446	−0.4973	−0.1217	0.9981	Yes	–	1.385
30	13.52	1.062	−0.5319	0.3410	0.2280	−0.0978	−0.2084	0.9995	No	65.28	0.860
31	2.295*	0.8857	0	−0.4284	0.5064	0	0	0.9866	No	63.22	1.336
32	0.7813*	1.117	−0.7796	0	0.6513	0	0	0.9927	No	69.14	0.823

Table C2 (*continued*)

Participant	c	p_{-1}	p_{-2}	p_{-3}	p^e_{-1}	p^e_{-2}	p^e_{-3}	R^2	*AC*	*Eq.*	*MSE*
33	−0.946*	1.767	−0.8572	0.1052	0	0	0	0.9937	No	63.07	1.665
34	8.501*	1.130	0	−0.4372	0	0	0	0.6584	No	27.67	0.905
35	1.851	1.182	0	−0.5068	0	0.2952	0	0.9931	Yes	–	1.385
36	14.01*	0.7478	0	0	0	0	0	0.2058	No	55.55	0.751
37	−3.020*	1.0498	0	0	0	0	0	0.9363	No	60.64	0.391
38	1.560	0.9728	0	0	0	0	0	0.9316	No	57.35	0.391
39	6.501	1.1315	−0.2359	0	0	0	0	0.9656	No	62.27	0.342
40	2.584*	1.043	0	0	0	−0.1619	0.0780	0.9719	No	63.18	0.298
41	1.739*	1.383	−0.4099	0	0	0	0	0.9443	No	64.65	0.360
42	1.113*	0.9327	−0.2968	0	0.3471	0	0	0.9569	No	65.47	0.445

Prediction rules for the 42 participants of the positive feedback treatment (least squares estimation of Eq. (7)). See the caption for Table C1 for more information.

Table C3
Estimation of individual first order heuristics.

PFOH participant no.	α_1	α_2	β	Original participant no.	Original group no.	Label
1	0.7389	0	0	2	N1	Naive fundamentalist
2	0	0	0	3	N1	None
3	0.9362	0	0	4	N1	Naive and fundamentalist
4	0.1350	0.1923	0.0605	5	N1	Fundamentalist
5	0.5689	0.3480	0	8	N2	None
6	0.5553	0	0	12	N2	Naive fundamentalist
7	0.7391	0	−0.4444	13	N3	None
8	0	0	0	14	N3	Naive and fundamentalist
9	−0.3770	0	−0.3762	18	N3	Naive fundamentalist
10	0.4016	0	0	19	N4	Naive fundamentalist
11	0	0	0	21	N4	Fundamentalist
12	0	0	0	24	N4	Fundamentalist
13	0.2633	0	0	25	N5	Fundamentalist
14	0	0	0	27	N5	Naive and fundamentalist
15	0	0	0	28	N5	Naive and fundamentalist
16	0.9101	0	0	30	N5	Naïve
17	0	0	0	34	N6	Fundamentalist
18	0.4321	0	0	35	N6	None
19	0	0	0	36	N6	None
20	1.5096	−0.5238	0	3	P1	Naive trend Follower
21	1.0177	0	0.8591	4	P1	Naive trend follower
22	0.5227	0.4711	0.9118	5	P1	Adaptive trend follower
23	1.0142	0	0.7818	6	P1	None
24	0.4888	0.5000	0.7290	9	P2	Adaptive trend follower
25	0.4670	0.5269	0.9210	11	P2	Adaptive trend follower
26	0.9994	0	0.4609	13	P3	Naive trend follower
27	0.5369	0.4627	0.5587	14	P3	Adaptive trend follower
28	1.0090	0	0.2765	15	P3	Naive trend follower
29	0.9557	0	0	20	P4	Naive trend follower
30	0.6669	0.3089	0.9696	21	P4	None
31	0.9616	0	0.8678	22	P4	None
32	0.9989	0	0.9437	24	P4	Naive and adaptive trend follower
33	0	0	0	26	P5	Naive and adaptive trend follower
34	0.2831	0.7045	0.8266	32	P6	Adaptive trend follower
35	1.1366	−0.1226	0.6077	33	P6	Naive trend follower
36	0.7428	0	0	36	P6	Naive trend follower
37	1.0376	0	0	37	P7	None
38	0.9419	0	0	38	P7	None
39	1.0155	0	0.2907	41	P7	Naive trend follower
40	0.6370	0.3842	0.3182	42	P7	Adaptive trend follower

Prediction rules for both treatments (least squares estimation of Eq. (9)) for the prism of first-order heuristics. The first column is the number of relevant participants, clustered according to treatment; the second to fourth columns show estimates of the coefficients, estimated by eliminating the least significant variable until all p values were below 5%. This procedure was applied only to the linear prediction rules statistically equivalent to a rule in the prism (Wald restriction test, 5% level). The fifth and sixth columns show the participant's original number and group (cf. Tables C1 and C2); the seventh checks for statistical equivalence with canonical rules (Wald restriction test, 5%).

Please cite this article as: Heemeijer, P., et al., Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. Journal of Economic Dynamics and Control (2009), doi:10.1016/j.jedc.2008.09.009

References

Adam, K., 2007. Experimental evidence on the persistence of output and inflation. Economic Journal 117, 603–636.

Arifovic, J., 1994. Genetic algorithm learning and the cobweb model. Journal of Economic Dynamics and Control 18, 3–28.

Barberis, N., Thaler, R., 2003. A survey of behavioral finance. In: Constantinidis, G.M., Harris, M., Stulz, R. (Eds.), Handbook of the Economics of Finance. Elsevier, pp. 1051–1121.

Beja, A., Goldman, M.B., 1980. On the dynamic behavior of prices in disequilibrium. Journal of Finance 35, 235–248.

Branch, W.A., 2004. The theory of rationally heterogeneous expectations: evidence from survey data on inflation expectations. Economic Journal 114, 592–621.

Bray, M.M., Savin, N.E., 1986. Rational expectations equilibria, learning and model specification. Econometrica 54, 1129–1160.

Brock, W.A., Hommes, C.H., 1997. A rational route to randomness. Econometrica 65, 1059–1095.

Brock, W.A., Hommes, C.H., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. Journal of Economics Dynamics and Control 22, 1235–1274.

Bullard, J.B., 2006. The learnability criterion and monetary policy. Federal Reserve Bank of St. Louis Review 88, 203–217.

Camerer, C.F., Fehr, E., 2006. When does "economic man" dominate social behavior? Science 311, 47–52.

Campbell, J.Y., Lo, A.W., MacKinlay, A.C., 1997. The Econometrics of Financial Markets. Princeton University Press, Princeton, NJ.

Canoles, W.B., Thompson, S.R., Irwin, S.H., France, V.G., 1998. An analysis of the profiles and motivations of habitual commodity speculators. Journal of Futures Markets 18, 765–801.

Cashin, P., McDermott, J.C., Scott, A., 2002. Booms and slumps in commodity prices. Journal of Development Economics 69, 277–296.

Challet, D., Zhang, Y.C., 1997. Emergence of cooperation and organization in an evolutionary game. Physica A 246, 407–418.

Chiarella, C., Dieci, R., Gardini, L., 2002. Speculative behavior and complex asset price dynamics: a global analysis. Journal of Economic Behavior & Organization 49, 173–197.

Cuthbertson, K., 1996. Quantitative Financial Economics: Stocks, Bonds and Foreign Exchange. Wiley, Chichester, UK.

Day, R.H., Huang, H., 1990. Bulls, bears and market sheep. Journal of Economic Behavior & Organization 14, 299–329.

De Bondt, W.F.M., Thaler, R.H., 1985. Does the stock market overreact? Journal of Finance 40, 793–805.

De Grauwe, P., Grimaldi, M., 2006. Exchange rate puzzles: a tale of switching attractors. European Economic Review 50, 1–33.

DeLong, J.B., Shleifer, A., Summers, L.H., Waldmann, R.J., 1990a. Noise trader risk in financial markets. Journal of Political Economy 98, 703–738.

DeLong, J.B., Shleifer, A., Summers, L.H., Waldmann, R.J., 1990b. Positive feedback investment strategies and destabilizing rational speculation. Journal of Finance 45, 379–395.

Dwyer, G.P., Williams, A.W., Battalio, R.C., Mason, T.I., 1993. Tests of rational expectations in a stark setting. Economic Journal 103, 586–601.

Economist, 1996a. That astonishing microchip. Section Leaders 338 (7958), March 23.

Economist, 1996b. When the chips are down. Section Semiconductors 338 (7958), March 23.

Economist, 2001. The great chip glut. Section Business 360 (8234), August 11.

Ehrblatt, W.Z., Hyndman, K., Özbay, E.Y., Schotter, A., 2006. Convergence: an experimental study of teaching and learning in repeated games. Working Paper, New York University.

Evans, G.W., Honkapohja, S., 2001. Learning and Expectations in Macroeconomics. Princeton University Press, Princeton, NJ.

Ezekiel, M., 1938. The cobweb theorem. Quarterly Journal of Economics 52, 255–280.

Fama, E.F., 1970. Efficient capital markets: a review of theory and empirical work. Journal of Finance 25, 383–417.

Fama, E.F., 1991. Efficient capital markets II. Journal of Finance 46, 1575–1617.

Fehr, E., Tyran, J.-R., 2001. Does money illusion matter? American Economic Review 91, 1239–1262.

Fehr, E., Tyran, J.-R., 2005. Individual irrationality and aggregate outcomes. Journal of Economic Perspectives 19, 43–66.

Fehr, E., Tyran, J.-R., 2008. Limited rationality and strategic interaction: the impact of the strategic environment on nominal inertia. Econometrica 76, 353–394.

Freeman, R.B., 1975. Legal "cobwebs": a recursive model of the market for new lawyers. Review of Economics and Statistics 57, 171–179.

Freeman, R.B., 1976. A cobweb model of the supply and starting salary of new engineers. Industrial and Labor Relations Review 29, 236–248.

Garber, P.M., 1990. Famous first bubbles. Journal of Economic Perspectives 4, 35–54.

Gerber, A., Hens, T., Vogt, B., 2002. Rational investor sentiment. Working Paper 126, Institute for Empirical Research in Economics, University of Zurich.

Gode, D.K., Sunder, S., 1993. Allocative efficiency of markets with zero-intelligence traders: market as a partial substitute for individual rationality. Journal of Political Economy 101, 119–137.

Guesnerie, R., 2005. Strategic substitutabilities versus strategic complementarities: towards a general theory of expectational coordination? Revue d'Economie Politique 115, 393–412.

Haltiwanger, J., Waldman, M., 1985. Rational expectations and the limits of rationality: an analysis of heterogeneity. American Economic Review 75, 326–340.

Haltiwanger, J., Waldman, M., 1989. Limited rationality and strategic complements: the implications for macroeconomics. Quarterly Journal of Economics 104, 463–483.

Hey, J.D., 1994. Expectations formation: rational or adaptive or ...? Journal of Economic Behavior & Organization 25, 329–349.

Hommes, C.H., 2006. Heterogeneous agent models in economics and finance. In: Tesfatsion, L., Judd, K.L. (Eds.), Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics. North-Holland, Amsterdam, pp. 1109–1186.

Hommes, C.H., Huang, H., Wang, D., 2005a. A robust rational route to randomness in a simple asset pricing model. Journal of Economic Dynamics and Control 29, 1043–1072.

Hommes, C.H., Sonnemans, J.H., Tuinstra, J., van de Velden, H., 2005b. Coordination of expectations in asset pricing experiments. Review of Financial Studies 18, 955–980.

Hommes, C.H., Sonnemans, J.H., Tuinstra, J., van de Velden, H., 2007. Learning in cobweb experiments. Macroeconomic Dynamics 11 (S1), 8–33.

Hommes, C.H., Sonnemans, J.H., Tuinstra, J., van de Velden, H., 2008. Expectations and bubbles in asset pricing experiments. Journal of Economic Behavior & Organization 67, 116–133.

Honkapohja, S., 1993. Adaptive learning and bounded rationality. European Economic Review 37, 587–594.

Kelley, H., Friedman, D., 2002. Learning to forecast price. Economic Inquiry 40, 556–573.

Kirman, A., 1993. Ants, rationality and recruitment. Quarterly Journal of Economics 108, 137–156.

Krugman, P.R., 2001. The oil-hog cycle. The New York Times, November 18, p. 14.

Lei, V., Noussair, C.N., Plott, C.R., 2001. Non-speculative bubbles in experimental asset markets: lack of common knowledge of rationality vs. actual irrationality. Econometrica 69, 831–859.

Lucas, R.E., 1972. Econometric testing of the natural rate hypothesis. In: Eckstein, O. (Ed.), The Econometrics of Price Determination Conference. Board of Governors of the Federal Reserve System and Social Science Research Council, Washington, DC, pp. 50–59.

Lux, T., 1995. Herd behaviour, bubbles and crashes. Economic Journal 105, 881–896.

Marimon, R., Spear, S.E., Sunder, S., 1993. Expectationally driven market volatility: an experimental study. Journal of Economic Theory 61, 74–103.

Marimon, R., Sunder, S., 1993. Indeterminacy of equilibria in a hyperinflationary world: experimental evidence. Econometrica 61, 1073–1108.

Muth, J.F., 1961. Rational expectations and the theory of price movements. Econometrica 29, 315–335.

Nerlove, M., 1958. Adaptive expectations and cobweb phenomena. Quarterly Journal of Economics 73, 227–240.

Noussair, C.N., Plott, C.R., Riezman, R.G., 1995. An experimental investigation of the patterns of international-trade. American Economic Review 85, 462–491.

Peterson, S.P., 1993. Forecasting dynamics and convergence to market fundamentals. Evidence from experimental asset markets. Journal of Economic Behavior & Organization 22, 269–284.

Potters, J., Suetens, S., 2009. Cooperation in experimental games of strategic complements and substitutes. Review of Economic Studies, Forthcoming.

Sargent, T.J., 1993. Bounded Rationality in Macroeconomics. Clarendon Press, Oxford.

Schmalensee, R., 1976. An experimental study of expectation formation. Econometrica 44, 17–41.

Shiller, R.J., 1981. Do stock prices move too much to be justified by subsequent changes in dividends. American Economic Review 71, 421–436.

Smith, V.L., 1962. An experimental study of competitive market behavior. Journal of Political Economy 70, 111–137.

Smith, V.L., Suchanek, G.L., Williams, A.W., 1988. Bubbles, crashes and endogeneous expectations in experimental spot asset markets. Econometrica 56, 1119–1151.

Sonnemans, J., Tuinstra, J., 2008. Positive expectations feedback systems and beauty contests as models of financial markets. CeNDEF Working Paper 08–07.

Sutan, A., Willinger, M., 2009. Guessing with negative feedback: an experiment. Journal of Economic Dynamics and Control, this issue, doi:10.1016/j.jedc.2009.01.005.

Tversky, A., Kahneman, D., 1974. Judgment under uncertainty: heuristics and biases. Science 185, 1124–1131.

Zarkin, G.A., 1985. Occupational choice: an application to the market for public school teachers. Quarterly Journal of Economics 100, 409–446.