Expectations in Experiments

Florian Wagener[*] [†] [‡]

20th November 2013

Abstract

The rational expectations hypothesis is one of the cornerstones of current economic theorising. This review discusses a number of experiments that focus on expectation formation by human subjects and analyses the implications for the rational expectations hypothesis. The experiments show that most agents are weakly rational and that their expectations coordinate quickly; but the strong rational expectations hypothesis poorly describes the expectational dynamics and is outperformed by other hypotheses.

Keywords

Rational expectations, expectation formation, laboratory experiments, human subjects

1 Two letters

The theoretical debate whether economic agents have rational expectations is not new: an early documented exchange of views occurred between Henri Poincaré and Léon Walras at the beginning of the twentieth century (Guesnerie, 2002).

Poincaré formulated the common sense objection to the hypothesis of rational expectations, which reduces in the absence of stochasticity to the hypothesis that economic agents have perfect foresight:

[*] CeNDEF, Department of Economics and Econometrics, University of Amsterdam, and Tinbergen Institute, Amsterdam, The Netherlands. E: F.O.O.Wagener@uva.nl

[†] I wish to thank Te Bao, Buz Brock, Cars Hommes and Jan Tuinstra for valuable critical comments on an earlier draft of this article. Some of the material in this review has been presented at the 2012 INEXC conference in Paris.

[‡] When citing this paper, please use the following: Wagener FOO. 2013. Expectations in Experiments. Annu. Rev. Econ. 6: Submitted. Doi: 10.1146/annurev-economics-080213-040935

> Par exemple, en mécanique, on néglige souvent le frottement (...). Vous, vous regardez les hommes comme infiniment égoïstes et infiniment clairvoyants. La première hypothèse peut être admise dans une première approximation, mais la deuxième nécessiterait peut-être quelques réserves.
>
> For instance, friction is often neglected in mechanics (...). In your case, you consider men as infinitely selfish and infinitely clairvoyant. The first assumption may be accepted as a first approximation, but the second may call for some reservations.
>
> Letter of Henri Poincaré to Léon Walras, 30 September 1901 (Jaffé, 1965, p. 164f)

Walras' response is indirect, but it seems that he is unwilling to let his assumption go:

> En réalité il y a des frottements dans le méchanisme économique; et d'autre part les hommes sont ni parfaitement égoïstes ni parfaitement clairvoyants. Il en résulte que la théorie (...) doit indiquer avec soin ces frottements (...)
>
> In reality, there are frictions in the economic mechanism; and agents are moreover neither infinitely selfish nor infinitely clairvoyant. It results that (...) theory (...) should indicate these frictions carefully.
>
> Letter of Léon Walras to Henri Poincaré, 3 October 1901 (Jaffé, 1965, p. 167).

In this exchange, the non-economist argued the assumption of rational expectations to be mildly absurd; the economic theoretician considered the advantages of this assumption for theoretical modelling to be sufficiently important to retain it. Unfortunately Walras did not formulate his defence more explicitly; it may be that he was, at least rudimentarily, aware of the game-theoretic ramifications of the hypothesis, as it expresses an expectational Nash equilibrium of all agents in the economy. Given his long-standing familiarity with Cournot's work, which in his time was not usual, this cannot be ruled out.

2 Degrees of rationality

It is a fundamental tenet of any theory of economics that agents behave rationally to some degree. The specific question of how they form expectations and whether they do this rationally is central to any theory of economic dynamics.

The origins of the rationality debate reach far back, at least to Hobbes and Descartes (cf. Hirschman, 1977). Muth (1961) gave an explicit formulation of what I shall call the 'ensemble' version of the rational expectations hypothesis: the average of all individual expectations is correct at any point in time. Lucas & Prescott (1971) formulated a 'homogeneous' version ten years later: all agents have the same expectations, which take the form of the correct future distribution of the expected variable. This is considerably stronger than the common sense notion of rationality, that there is no systematic bias in expectations, in this article referred to as the 'unbiased expectations' hypothesis.

Rational expectation models were theoretically highly attractive; empirical work provided a certain measure of support for them (see e.g. Turnovsky (1970); Turnovsky & Wachter (1972); Lucas (1973); Sargent et al. (1973); McCallum (1976)); as a consequence these models gained broad acceptance. Their radical implications simultaneously attracted severe criticism (Arrow, 1986). In the debate, the unbiased expectations hypothesis was, and is, often used to justify the homogeneous version, based on the (unjustified) assumption of equivalence of the two.

The Smith-Suchanek-Williams experiment (Smith et al., 1988) broke new ground. In an experimental setting, where the economic environment could be controlled, and where both the information available to the agents as well as their actions were perfectly known, groups of agents eventually arrived collectively at the rational expectations equilibrium through market interactions. This happened however only after a prolonged learning phase, lasting several runs of the experiment with the same subjects, which was characterised by strong speculative behaviour, and whose occurrence was considered at the time to be surprising (Sunder, 1995). In a refined set-up the behaviour of some of the agents was shown to be irrational, at least during the learning phase, in the sense that some of their purchasing and selling decisions resulted with certainty in a loss (Lei et al., 2001).

It has never been questioned that some kind of learning has to occur before agents can arrive at a rational expectations equilibrium. But experiments like Smith et al. (1988) emphasised that learning is central to economic behaviour; experimental evidence indicated moreover that the institutional structure of the market has major influence on the speed of the learning process (Smith, 1991).

A number of learning theories has been proposed. Perhaps closest to the spirit of the rational expectations hypothesis is Guesnerie's concept of 'eductive learning' (Guesnerie, 1992, 2002), where agents, to form their expectations, make full use of all available information about their economic environment, as well of the 'common-knowledge' hypothesis that all other agents do the same. This type of learning does not necessarily require full knowledge of the system, though the demands on the reasoning capabilities of the agents are substantial.

These demands are much lighter in 'evolutive learning' theories, where agents are modelled as econometricians or adaptive learners. The former have a perceived law of motion of the system on which they base their actions, estimating its parameters continuously as new data becomes available. Under broad conditions such an econometric learning procedure can converge to the rational expectations equilibrium (Bray & Savin, 1986; Marcet & Sargent, 1989; Evans & Honkapohja, 1995). Adaptive learners consider a number of different adaptive expectation rules, using that which fits the observed data best at that particular instant of time; the 'discrete choice' learning models of Brock & Hommes (1997, 1998) are of this type. Models where agents can choose between both adaptive and econometric learning rules have been considered as well (Tuinstra & Wagener, 2007; Branch & Evans, 2010).

Learning has been studied widely in the laboratory, and there is a tendency to refine experimental setups in order to elicit expectations of agents directly, rather than having to derive these from the agents' actions on the basis of a putative 'inner model' of their behaviour. This class of setups has been denoted as 'learning-to-forecast' experiments (Marimon & Sunder, 1994).

The aim of the present survey is to discuss a number of learning-to-forecast experiments that focussed on the evolution of agents' expectations of an economic variable in the situation that the aggregate expectations determined the evolution of this variable (Sonnemans et al., 2004; Hommes et al., 2005, 2007, 2008; Heemeijer et al., 2009; Bao et al., 2012); see also Hommes (2011). In these experiments subjects have to give their best predictions of a price variable; as their pay-off depends directly on the quality of the prediction, these predictions can be taken as the subjects' individual expectations. Through the law of motion of the economic system, these expectations feed back into the realisations of the economic variable observed by the subjects and are presumably used by them to form their predictions of its next realisation. This highly abstract setting forms a natural testing ground for different interpretations of the rational expectations hypothesis, some of which are explored in the following.

2.1 Muth's 'ensemble' rational expectations hypothesis

There is a subtle but highly important difference between Muth's original formulation of the rational expectations hypothesis and its subsequent interpretation by Lucas & Prescott (1971). Muth's hypothesis states that the ensemble of all expectations is distributed about the theoretical value:

> "(...) expectations of firms (or, more generally, the subjective probability distribution of outcomes) tend to be distributed, for the same information set, about the prediction of the theory (or the 'objective' probability distributions of outcomes). (...) It [i.e. the hypothesis] *does not* assert that the scratch work of entrepreneurs resembles the system of equations in any way; nor does it state that predictions of entrepreneurs are perfect or that their expectations are all the same". (Muth, 1961, p. 316)
>
> "Allowing for cross-sectional differences in expectations is a simple matter, because their aggregate effect is negligible as long as the deviation from the rational forecast for an individual firm is not strongly correlated with those of the others. Modifications are necessary only if the correlation of the errors is large and depends systematically on other explanatory variables". (Muth, 1961, p. 321)

The comment in square brackets is mine; the emphasis is Muth's.

Some care is needed in interpreting this definition, in particular the term 'subjective probability distribution'. As Muth notes emphatically that the hypothesis does not imply that the predictions of entrepreneurs are perfect, or that their expectations are all the same, 'subjective probability distribution' has to refer to the distribution of the observable predictions of the firms involved. In other words, Muth's hypothesis is phenomenological: it does not have an internal model of the entrepreneurs, nor does it impose that all entrepreneurs behave the same.

The point is crucial, as it implies that Muth's hypothesis is testable without having to impose a model of the internal workings of the economic agents, or even without having to have a model of the economy at all: if time series of expectations of an economic variable are known for a representative ensemble of agents, as well as a time series of its realisations, Muth's rational expectations hypothesis can be tested.

2.2 The unbiased expectations hypothesis

In statistical mechanics, a system is 'ergodic' if the average of an observable over an ensemble of systems, taken at a specific point in time, is equal to the time average of the same observable, computed for a single system. Analogously, to Muth's hypothesis of ensemble rationality there corresponds a 'time-averaged' version of the rational expectations hypothesis, stating that for every single agent the time-average of the mismatch between expectations and realisations vanishes; that is, over time the expectations of every agent are unbiased with respect to the theoretical value. A stronger formulation of this hypothesis also requires that the expectation errors do not display any autocorrelation structure. Also this hypothesis is testable on the basis of observed data.

This notion has variously been called 'informational efficiency' (Hommes et al., 2005) or 'internal rationality' (Adam & Marcet, 2011); for the purposes of this survey, it will be called the 'unbiased expectations hypothesis'.

2.3 The 'homogeneous' rational expectations hypothesis

The formulation of Muth's hypothesis given by Lucas & Prescott (1971) is usually taken for a mere rewording, which it is not, and it is this form in which the rational expectations hypothesis is ordinarily used in theoretical economic analysis:

> "(...) we shall (...) go to the opposite extreme, assuming that the actual and anticipated prices have the same probability distribution, or that price expectations are rational." (Lucas & Prescott, 1971, p. 660)
>
> "Specifically, we assume that expectations of firms are rational, or that the anticipated price at time t is the same function of $(u_1, ..., u_t)$ as is the actual price. That is, we assume that firms know the true distribution of prices for all future periods". (Lucas & Prescott, 1971, p. 664).

In contrast to Muth's ensemble version, this form of the hypothesis assumes an internal model of economic agents, and it states explicitly that the expectations of all agents are the same. That is, Lucas and Prescott strengthen Muth's definition by explicitly adding the assumption of full homogeneity of agents.

Arrow formulated the Mantel-Sonnenschein-Debreu theorem, that rationality on the micro level has few implications on the macro level, succinctly as follows: "In the aggregate, the hypothesis of rational behaviour has in general no implication" (Arrow, 1986, Section II). Adding the assumption of homogeneity is therefore by no means innocent; it is precisely this assumption that allows to make powerful theoretical predictions.

The homogeneous rational expectations hypothesis itself is not testable based on sequences of predicted and expected prices. It is a statement about the equality of two distributions, of which one, the 'anticipated distributions', only exists in the mind of the agents, and is not directly observable: it can only be constructed if models of the economy and of the agents are available. Only in the special case that there is no stochasticity in the system, the distributions degenerate, and the hypothesis becomes testable; but unless expectations and realisations are always exactly equal, every statistical test will reject homogeneous rationality in this situation.

The wording of the hypothesis implies that it is only meant to describe a time-asymptotic state of the economic system, making no statements about the rate of convergence towards this state (Modigliani, 1977; Friedman, 1979). This is acknowledged by Lucas and Prescott (cf. also the discussion in Lucas & Sargent, 1979):

> "Thus we surrender, in advance, any hope of shedding light on the process by which firms translate current information into price forecasts."
> (Lucas & Prescott, 1971, p. 660)

Finally, the homogeneous hypothesis implicitly assumes that obtaining rational expectations is costless (see Shaw, 1984, chapter 5, especially figure 5.1). Weakening this assumption is at the basis of the rational inattention literature (Sims, 1998, 2003, 2005; Woodford, 2009).

2.4 The success of rational expectations

Given the radicality of the homogeneous rational expectations hypothesis, it is legitimate to inquire after the secrets of its phenomenal success. One of these must be that it was introduced at an opportune moment in time. In the early days of econometric modelling, expectations were assumed to be formed adaptively: price forecasts were assumed to be conditioned on (a fixed segment of) the time series of past prices. Adaptive expectation formation implies, for instance, a theoretically exploitable trade-off between inflation and output in macroeconomic models, which is absent under rational expectations (Sargent & Wallace, 1975). In practice, inflationary policies triggered immediate demands for higher wages, nullifying most of the trade-offs and implying that expectations are based on future rather than past conditions.

The homogeneous rational expectations hypothesis solves the modelling problem involved. In its context, actions are rational if agents do not have an incentive to change them in view of the future evolution of the system given these actions. This brings out its game-theoretic background: a rational expectations equilibrium is a Nash equilibrium in a dynamic many-agent situation, where an individual agent cannot change the evolution of the aggregate variables unilaterally. Because of this, it fulfils an important benchmark function. Yet even if agents are perfectly informed about the system and about each others rationality, coordination on the rational expectations equilibrium is problematic (Guesnerie, 1992, 2002).

Another reason for the success of the homogeneous hypothesis is the very strength of the formulation, or what may be called its universality: it makes a definite prediction, which is moreover not dependent on any parameters that would have to be measured or fitted.

Finally, in practice the difference between adaptive and rational expectations amounts to the fact that under rational expectations agents forget irrelevant information quickly. This is a direct consequence of their forward looking nature; but it is also an often observed feature of human economic subjects, that they anticipate certain structural changes in their environment and adapt themselves to it.

3 Experiments

In the Smith et al. (1988) experiment, subjects have to perform two tasks: they have to form expectations about the future behaviour of the price, or, equivalently, about the future behaviour of the market participants, and then they have to make trading decisions. As both tasks are performed internally, the interpretation of the resulting trading data is complicated. Generally, subjects who have to perform both tasks simultaneously perform worse than those who only perform one of these (Bao et al., 2013).

This motivated separating the tasks into learning-to-forecast and learning-to-optimise. The remainder of this review will discuss a number of learning-to-forecast experiments that have been performed at CeNDEF, mostly in cooperation with the CREED laboratory, at the University of Amsterdam.

3.1 Convergence to rational expectations by competition?

One of the central arguments proposed to motivate rational expectations invokes evolutionary selection: rationally forecasting agents will outperform other agents in the long run and hence will take over the market eventually.

The Sonnemans et al. (2004) experiment investigated aspects of this argument. It was a classroom strategy experiment in the spirit of Axelrod (1984): a number of subjects, students in a dynamical systems course, devised functional forecasting rules of a Muthian cobweb economy with nonlinear supply functions. Strategies were allowed to condition on past realised prices and on their private past predictions; they had to be submitted in writing, were then translated into functional form and coded.

There were four rounds, each pair separated by two weeks. In each round at least twenty strategies were submitted. Per round, 620 simulations were run; for each simulation, six strategies were drawn at random from the pool of submitted strategies, and partook in a forecasting experiment.[1] After the round, time series of the predictions of the strategy and the price realisations were given to the subjects.

Parameters were chosen in such a way that the value of the ratio of marginal total supply to marginal total demand at the market clearing price equalled approximately -7.5; this means that the expectations feedback system was far from being stable under naïve expectations — that is, under the rule that expects the next price realisation to be equal to the most recent realisation — and hence far from being eductively stable in the sense of Guesnerie (1992).

The mean quadratic distance of the realised prices to the rational expectations equilibrium price decreased monotonically over the four rounds, with the value in the last round about 25% of that of the first. About 10% of the simulations converged to a steady state (near) the rational expectations state. In contrast to this, about 50% of the simulations exhibited chaotic dynamics. Although these numbers were roughly constant over the rounds, the highest fraction of chaotic dynamics occurred in the final round, where the average distance of the dynamics to the rational expectations equilibrium was smallest. This increase in dynamical complexity was also reflected in the average length of the code per strategy, which increased monotonically over the rounds.

These results do not support the evolutionary motivation of the emergence of rational expectations as the effect of increasing competition. They rather point to the increase

[1] Six is a popular group size in experimental economics: it is sufficiently large to practically exclude collusion when participants, or their strategies, cannot communicate; conversely, it is sufficiently small to allow to draw many different groups from a pool of moderate size.

in dynamic and informational complexity as competition between the subjects intensifies. Even if homogeneous rational expectations is the only surviving strategy in the limit of infinite time series and infinite computational resources of the subjects, this does not guarantee that it is the outcome of the process of increasing these resources indefinitely (cf. Brock et al., 2009, for a similar non-convergence result).

The results do however not contradict the 'common knowledge' approach to expectation formation. A possible interpretation the results is that the traders on the market could only establish a certain band of 'possible prices' as common knowledge. Even the for the lowest estimate of this price band, more than 90% of the variation of the supply function occurs within it; so only very little information about the price has been established as common knowledge.

3.2 Common features of the laboratory experiments

The experiments reported by Hommes et al. (2005, 2008) and Heemeijer et al. (2009) investigated the impact of institutional structure on expectation dynamics. Unlike the strategy experiment of Sonnemans et al. (2004), these were conducted in a laboratory and ran for at most three hours. Also, whereas Sonnemans et al. (2004) elicited prediction rules, the laboratory experiments only elicited point predictions. This subsection describes the common features of these three experiments.

In all of them, subjects participated in a forecasting task that lasted for 50 periods. At the beginning of each experiment, subjects were told that they had to act as advisers to a firm that was operating on a certain market. In order to make production decisions, the firm needed a price forecast of the market price next period. Some qualitative information was given on the market behaviour, whether an increase in price expectations would occasion the price to go down or to go up. The subjects were not given any quantitative information; in particular, they had not the information necessary to derive the rational equilibrium price.

During the experiment, a computer screen provided the subjects with a complete list of their own past forecasts and of past price realisations, both in numerical and in graphical form, but not the forecasts, past or present, of the other subjects. Also their total earnings, the earnings in the most recent period, as well as the total number of remaining rounds were displayed.

Groups of six subjects each were formed. At each point in time, the price expectations of all the subjects in a group generated a price realisation through a rule that was particular to the experiment. Consequently, the information on the screen was updated.

In this way, the forecasts of the subjects determined the dynamics of the price series. On average, one round of expectation formation took about two minutes.

After realisation of the price, individual period earnings E_{it} were determined according to

$$E_{it} = 1300 \max \left\{ 1 - \left(\frac{p^e_{it} - p_t}{7} \right)^2, 0 \right\} : \tag{1}$$

here p^e_{it} is subject i's expected value of the realised price p_t. At the end of the experiment, participants' earnings were converted to euros according to an exchange rate of 2600 earning points to 1 euro. The theoretic maximum of possible earnings was therefore 25 euros in each experiment. For a typical experiment (Heemeijer et al., 2009), expected earning were 24.87 if all agents had formed their expectations according to the homogeneous rational expectations hypothesis. The average of actual realised earnings over all runs was about 22 euros, that is, 88% of the 'rational' value.

3.3 The impact of market structure

In Hommes et al. (2005), the price realisation rule, reminiscent of an asset pricing market, was given by

$$p_t = \frac{1}{R} \left((1 - n_t) \frac{1}{6} \sum_{i=1}^{6} p^e_{it+1} + n_t p^* + \bar{y} + \varepsilon_t \right); \tag{2}$$

here $R = 1.05$ is the gross interest rate, $\bar{y} = 3$ the mean dividend pay-off, $\varepsilon_t \sim N(0, 1/4)$ independently distributed dividend fluctuations, $p^* = \bar{y}/(R-1) = 60$ the fundamental price and p^e_{it+1} the individual expectations of the six experimental subjects of next period's price. Finally $n_t = 1 - \exp\left(|p_{t-1} - p^*|/200\right)$ is a fraction of 'robot' fundamental traders that always predict the fundamental price p^*; these traders act as a 'stabilising' force if the deviation of the price from the fundamental becomes large. Without robot traders this kind of market is prone to having strong prolonged price bubbles, introducing 'ceiling events' as the predictions hit a previously unknown upper bound.

The informational structure of (2) is that of a positive feedback loop: an increase in the average price expectations effects an increase of the realised price. Three types of qualitatively different expectation dynamics were observed: monotone convergence to the rational expectations equilibrium, oscillatory decay towards the rational expectations equilibrium and persistent non-decaying oscillations. Which of these would obtain in a run was the outcome of an initial brief coordination phase, taking up to five periods.

After that, price predictions were strongly coordinated over the group.

The Hommes et al. (2007) experiment considered a negative feedback structure instead, by taking a nonlinear Muthian cobweb evolution

$$p_t = \alpha - \beta \sum_{i=1}^{6} S\left(p_{it}^{e}\right) + \varepsilon_t; \tag{3}$$

the supply functions are given as $S(p) = \tanh(\lambda(p-6)) + 1$. The experiment consisted of three treatments, best characterised by the derivative $\sigma = \beta S'(p^*)$ of the (deterministic part of the) evolution equation at the rational equilibrium price p^*: $\sigma = -0.87$, -1.96 and -7.75 respectively. That is, the dynamics are stable under naïve expectations in the first treatment, but unstable in the second and third.

Yet the observed expectational dynamics converged to rational expectations not only in the first treatment, but also in the second; in the third treatment convergence did not seem to obtain. Moreover, convergence in the first two treatments was rapid, occurring almost instantly.

Smith (1991) observed that in very many experiments, prices and allocations converged to their predicted rational expectations value, in spite of knowledge being either incomplete or private or both. He concluded that there is no support for the attainment of rational equilibria by conscious cognitive efforts, but that it is achieved rather by the institutional structure, though noting "What is imperfectly understood is the precise manner in which institutions serve as social tools that reinforce, even induce, individual rationality".

The Hommes et al. (2007) experiment identifies negative informational feedback as a mechanism that can induce individual rationality. Even in the moderately unstable situation rapid convergence towards the rational expectations equilibrium is obtained: this reflects the efficiency of most commodity markets.

Under positive informational feedback, the collective prediction behaviour of the experimental subjects is much more unstable, as price bubbles and persistent price oscillations occurred, which were sustained over the complete run.

3.4 Negative vs positive informational feedback

The comparison of the Hommes et al. (2005) and the Hommes et al. (2007) experiments is complicated by two facts: the price formation rule of the former, equation (2), is linear; that of the latter, equation (3), is nonlinear. Moreover, in the former experiment the realised price this period depends on the expected price next period, while in the latter the realised price this period depends on its expected value. The wish to be able to analyse positive and negative informational feedback in a symmetric setting informed the design of the Heemeijer et al. (2009) experiment, which will be discussed at some length.

In this experiment, homogeneous rationality implies that subjects expect a unique value for the expected price, rather than a forward looking dynamic law of which they would have to evaluate the dynamic consequences correctly. In this sense, it is the simplest of the learning-to-forecast experiments, and the one in which the homogeneous rational expectations hypothesis can be expected to have its best shot.

3.4.1 *Price formation.*

In the Heemeijer et al. (2009) setup, the price formation rules (2) and (3) of the earlier experiments were replaced by

$$p_t = \alpha + \beta \left(\frac{1}{6} \sum_{i=1}^{6} p_{it}^e \right) + \varepsilon_t \tag{4}$$

The random terms ε_t were independently distributed normal variables with mean 0 and variance $1/4$.

There were two treatments, representing two different kinds of 'market structure'. In both treatments, the parameters α and β were chosen such that the rational expectations equilibrium price equalled $p^* = 60$. In the first treatment, the 'feedback' parameter β was set to $\beta = -0.95$; in the second, $\beta = 0.95$. Accordingly, the treatments are denoted 'negative' and 'positive' feedback, respectively.

The negative feedback treatment is an exact implementation of a Muthian cobweb that is stable under naïve expectations and hence under eductive learning. The positive feedback treatment shares structural features with asset pricing models; but in most such models, agents make predictions about the price of the next period, rather than the current period.

3.4.2 *Results.*

The data of the experiments are published on the CeNDEF web site, and are freely accessible. The raw results of the negative feedback treatment are shown in figure 1. Individual predictions coordinated and converged rapidly towards values that are close to the rational expectations value. This is remarkable given the 'weak' nature of the

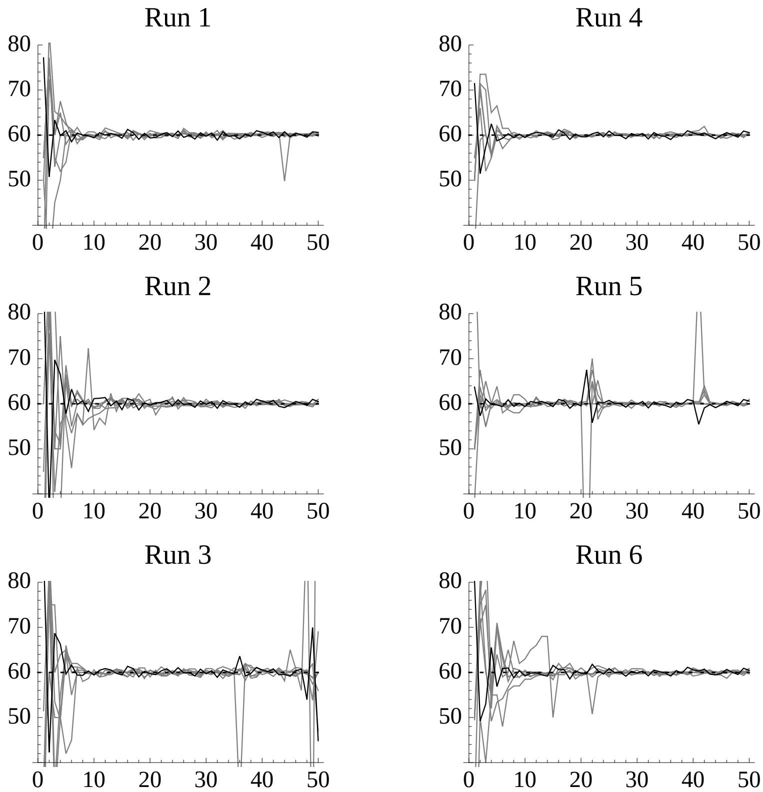

Figuur 1: *Negative feedback treatment of the (Heemeijer et al., 2009) experiment: price expectations (grey lines), realised prices (black line) and rational expectations equilibrium price (dotted line).*

feedback: the value $\beta = -0.95$ is close to the value -1 bounding the eductive stability region. A number of agents actually ended up by constantly predicting the rational expectations value $p^* = 60$ exactly for most of the run.

In runs 5 and 6 some of the subjects tried to experiment for some periods, e.g. around $t = 20$ and $t = 40$ in run 5. They were presumably discouraged from pursuing this by low earnings. In run 3 agents were probably expecting an end-of-period effect, which destabilised the price in the last few periods.

Figure 2 illustrates the results of the positive feedback treatment. One of the runs has been left out: it is the run marked P5 in figure 3 of Heemeijer et al. (2009), where expectations and realisations went far over the value of 100 in some time periods.

In this treatment again there is rapid coordination of individual predictions around a common prediction; but that common prediction deviates from the rational expectations prediction more than in the negative feedback case, in some cases substantially

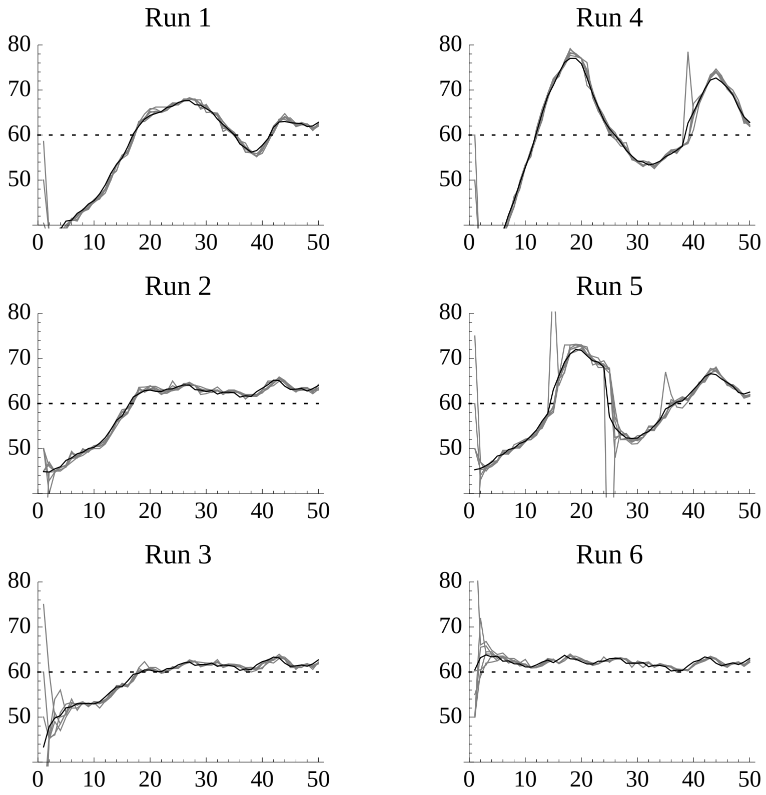

Figuur 2: *As figure 1: positive feedback treatment.*

more. There seems to be substantial endogenous dynamics of the common prediction, which moreover varies qualitatively: in some runs it converges monotonously towards the rational expectations value; in others it converges but oscillates. Again, as in the negative feedback treatment, subjects tend to experiment occasionally (runs 4 and 5).

Testing the unbiased expectations hypothesis.. As all individual predictions are available, the 'unbiased expectations hypothesis' can be tested. The individual expectation errors $e_{it} = p_t - p_{it}^e$ is the difference between realised and expected price. In order to ensure that expectational coordination has taken place, only the realisations of e_{it} for $11 \leq t \leq 50$ have been considered. Under the unbiased expectations hypothesis, the mean of the e_{it} is expected to be zero.

For the negative feedback treatment, the hypothesis is rejected 5 times, out of 36, at the 10% level; for positive feedback, there are 3 rejections, out of 36, at the same level. This supports rather strongly the hypothesis that the subjects have unbiased expectations.

Convergence to homogeneous rational expectations.. As the homogeneous rational expectations hypothesis is a statement about an asymptotic steady state, it needs to be complemented with a theory of the rate of convergence towards that state. Only if the rate of advent of new information to the system is much slower than the convergence rate, there is a case for homogeneous rational expectations to be the correct description of actual expectations; if information arrives at a faster rate, transient dynamics will play a significant role.

Information about these transients is obtained as follows: a model of the subjects behaviour is estimated; then non-stochastic price dynamics are simulated using the estimated behaviour. These dynamics turn out to converge for all runs: the asymptotic steady state prices as well as the deterministic rate of convergence towards that price are estimated for the simulated dynamics.

The details of this programme were executed as follows. Heemeijer et al. (2009) provided best estimates of individual forecasting rules of at most three lags in past forecasts and past price realisations (tables C1 and C2 in Heemeijer et al. (2009)). Substituting these for the price expectations p^e_{it} in (4) and running the price evolution with the same noise realisations as used in the experiment gives fits of varying quality. It was usually sufficient to remove one individual estimated rule per run, and in one instance two, to improve the fits to a satisfactory degree. What is more, in the two runs that diverge most from the estimated trajectories, the expectation of one of the subjects was radically deviating from all others, either due to a typing mistake or due to the wish to explore. If this is taken into account, the fit obtained by the remaining rules to the data is remarkably good.

Treating the remaining forecasting rules as describing the actual behaviour of the subjects, a new realised price time series is obtained for each run by running the evolution law (4) without noisy perturbations. Then a linear model of the form

$$p_t = c_0 + \sum_{i=1}^{3} c_i p_{t-i} \tag{5}$$

is estimated on the resulting time series, giving for run k estimates $\hat{c}_{ki}$ of the parameters c_i. The implied steady state price p^*_k is computed according to

$$p^*_k = \frac{\hat{c}_{k0}}{1 - \sum_{i=1}^{3} \hat{c}_{ki}}. \tag{6}$$

The implied rate of convergence r_k is defined as

$$r_k = -\log|\lambda_k|, \tag{7}$$

where $\lambda = \lambda_k$ is the root having the largest complex absolute value of the equation

$$\lambda^3 - \sum_{i=1}^{3} \hat{c}_{ki}\lambda^{3-i} = 0. \tag{8}$$

The results for p_k^* and r_k are very robust with respect to the number of lags taken in equation (5).

Run k	p_k^*	r_k
Run 1	60.07	0.1022
Run 2	62.20	0.1892
Run 3	60.71	0.0832
Run 4	59.49	0.0617
Run 5	58.83	0.1916
Run 6	61.10	0.0882
Mean	60.4 ± 1.2	0.119 ± 0.057

Tabel 1: *Estimated steady states and convergence rates (positive feedback treatment), their sample means and sample standard deviations.*

The results are shown in table 1. All estimated steady state values p_k^* of the realised price are close to the rational expectations value. The hypothesis that they equal the rational value of 60 cannot be rejected even at the 10% level. This supports the homogeneous rational expectations hypothesis asymptotically, in that prices converge to the rational expectations equilibrium values.

The convergence rates are low. It might be conjectured that they equal the feedback strength $r_\beta = -\log 0.95 \approx 0.0513$. But this hypothesis is rejected at the 5% level. Nevertheless, data from Sonnemans & Tuinstra (2010), where the same experiment was conducted with a feedback strength $r_\beta = -\log 0.667 \approx 0.4$, show that the convergence rate increases with the feedback strength. It may be conjectured that in a positive feedback situation this dependence is linear.

Of course, a single data point cannot decide this conjecture. If it were however true, it would have interesting implications: asset markets have typically a positive feedback structure, with the feedback strength proportional to the risk free interest rate. If the convergence rate is of the same order of magnitude as the risk free rate, the time scale of convergence towards the rational expectations equilibrium, which is of the order of the inverse of the convergence rate, is of the order of magnitude between 6 and 50 years, for risk free rates between 0.15 and 0.02 respectively.

3.4.3 Homogeneous rationality.

One of the major virtues of the homogeneous rational expectations hypothesis is that it is universal and that it provides very strong restrictions on predicted agent behaviour; for instance, in the simple context of the Heemeijer et al. (2009) experiment, the only price prediction that is consistent with the homogeneous rational expectations (HRE) hypothesis is

$$p^e_{\mathrm{HRE,t}} = p^* = 60 \tag{9}$$

for all t.

Any other model of expectation formation of agents should share the property of being universal. Such a model can contain parameters that are unknown beforehand; but these parameters should be determined in such a way that they fit all situations in which the homogeneous rational hypothesis and its alternative are to be compared; in particular, they should not be estimated anew for each different context. This section tries to construct a competing universal model that aims to explain the experimental data of the Heemeijer et al. (2009) experiment better than the homogeneous rational expectations hypothesis.

In constructing this model, two goals are aimed at. Firstly, the model should be 'externally consistent' with the observed time series: this is the usual natural requirement that it should provide good in-sample descriptions and good out-of-sample predictions of the observed data. Secondly, the model should be 'internally consistent' with the data: price predictions at a certain point in time should be conditioned only on data that are available to the experimental subjects at that time, or, put differently, the prediction rule should be such that it could be used by the subjects during the experiment. In particular, in the context of the Heemeijer et al. (2009) experiment, expectations derived from the homogeneous rational expectations hypothesis do not fit the second requirement, as the subjects had been given only general qualitative information about the market they were facing, and in particular lacked the data to compute the rational expectations equilibrium price.

Precisely because of the uncertainty about the environment which they are facing, the experimental subjects cannot use eductive learning, and they have to use some form of evolutionary learning. They are assumed to construct a number of competing prediction rules and select that rule which has been performing best, according to some criterion, in the past. As this is a restricted exercise, for the predictors three ordinary least square learning rules are chosen, in the spirit of Branch & Evans (2010). This choice highlights the central problem of any alternative theory to rational expectations: which forecast rules, or more generally, which perceived laws of motion to choose, as it is inconsistent with our internal evidence of human behaviour to assume that all possible prediction rules should qualify. It is more reasonable to let agents use a criterion to decide at which point in time they are dissatisfied with the currently available rules, and then to have a theory how an additional rule is chosen. As remarked before in the context of the Sonnemans et al. (2004) strategy experiment, it is not at all clear that rational expectations will be the eventual outcome of an ongoing competition between different forecasting rules.

The alternative will be denoted as the 'switching OLS' (sOLS) model. Motivated by the rapid coordination of expectations which is observed in the experiments, agents are assumed to be homogeneous. Moreover, it is assumed that they condition their price predictions only on the publicly available realised prices. At the beginning of time period t, agents have access to the series of past realised prices: that is, their information set is

$$I_{t-1} = \{p_1, \dots, p_{t-1}\}. \tag{10}$$

They keep track of three prediction rules, indexed by the number of time lags $\ell = 0, 1, 2$; these are ordinary least squares predictions of the form

$$p^e_{\ell,t} = p^e_{\ell,t}(I_{t-1}) = \hat{c}_{\ell 0,t} + \sum_{s=1}^{\ell} \hat{c}_{\ell s,t}\, p_{t-s}, \tag{11}$$

where the $\hat{c}_{\ell s,t} = \hat{c}_{\ell s,t}(I_{t-1})$, for $s = 0, \dots, \ell$, are obtained by fitting the lag-ℓ prediction rule (11) on the price time series I_{t-1}. Let $L_t = L_t(I_{t-1})$ be the lag of the rule with the lowest total squared prediction error $\sum_{s=1}^{t-1}(p^e_{\ell,s} - p_s)^2$. Then

$$p^e_{\text{sOLS},t}(I_{t-1}) = p^e_{L_t,t}(I_{t-1}) \tag{12}$$

is the new prediction of the sOLS model.

One-step-ahead. First, the one-step-ahead prediction quality of the homogeneous rational expectations model and the switching OLS model are compared; that is, for the switching OLS model the prediction rules (11) are progressively estimated on the time series $I^{\text{lab}}_{t-1} = \{p^{\text{lab}}_1, \dots, p^{\text{lab}}_{t-1}\}$ of realised prices of the laboratory experiments. Let $\overline{p}^e_{\text{lab},t}$ denote the corresponding average price prediction of the experimental subjects at time t. To compare the quality of the predictions, for each model and for each run the mean of the absolute prediction errors

$$\bar{e}_{\text{model}} = \frac{1}{46} \sum_{t=5}^{50} \left| \overline{p}^{e}_{\text{lab},t} - p^{e}_{\text{model},t}(I^{\text{lab}}_{t-1}) \right| \tag{13}$$

have been computed for both models, as well as the probability of these mean absolute prediction errors being drawn from the same distribution.

The results depend strongly on the feedback structure. Homogeneous rational expectations outperform switching OLS expectations significantly in all situations of the negative feedback treatment. Yet both models perform well: if the absolute modelling error is defined as the mean difference between values that are observed and those that are predicted by the model, the largest absolute modelling error of the homogeneous rational expectations rule is 1.1057, or about $1.1057/60 \approx 2\%$ of the rational expectations equilibrium price; the modelling error of the switching OLS rule is significantly larger, but still rather good at 3% of p^*.

In the positive feedback treatment however, the switching OLS rule outperforms homogeneous rational expectations in all situations, as the modelling error turns out to be significantly smaller. For instance, the largest modelling error of the switching OLS model is again about 3% of the fundamental price, while it equals 14% for the homogenous rational expectations model.

Figures 3 and 4 show, for the negative and the positive feedback treatment respectively, the price realisations and the one-step ahead predictions of both models. The predictions of the switching OLS model, based on the time series realised in the experiment, track the coordinated actual predictions much closer than the, constant, homogeneous rational predictions.

Many-steps-ahead. To assess the performance of the switching OLS model for many-steps-ahead predictions, the rules (11) are estimated on an initial segment $\{p^{\text{lab}}_{t_0}, \ldots, p^{\text{lab}}_{t_1}\}$ of the time series obtained in the experiment. For $t > t_1$, simulated price realisations p^{sim}_t are obtained by computation, that is, by substituting price expectations p^e_t obtained from (11) into (4), with the noise realisations ε_t the same as in the experiment. The information set of the simulated agents is thus given as

$$I^{\text{sim}}_{t-1} = \left\{ p^{\text{lab}}_{t_0}, \ldots, p^{\text{lab}}_{t_1}, p^{\text{sim}}_{t_1+1}, \ldots, p^{\text{sim}}_{t-1} \right\}. \tag{14}$$

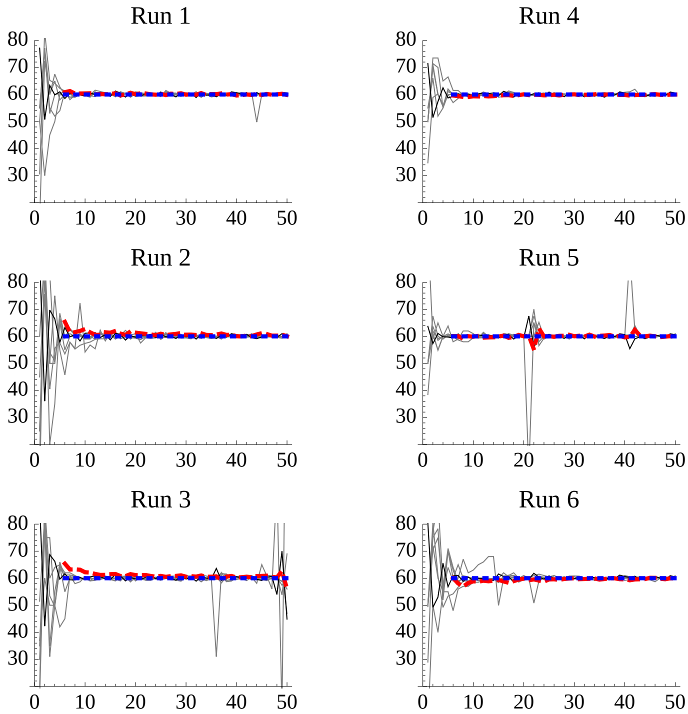

Figuur 3: *One-step ahead predictions for the negative feedback treatment. Homogeneous rational expectations (blue, dotted), switching ordinary least squares predictions (red, dashed), and actual predictions of subjects (black, solid).*

The beginning of the training interval has been chosen at $t_0 = 4$ to ensure rough initial coordination of beliefs. Its end has been set equal to $t_1 = 24$, in order to avoid having in the training set the presumable typing error that occurred at $t = 25$ of one of the participants in run 5 of the positive feedback treatment (cf. figure 2).

Figure 5 shows the many-steps-ahead predictions of the homogeneous rational expectations rule and the switching OLS rule under the same noise realisations as in the experiments. The switching OLS predictions track the predictions of the experimental subjects significantly better in most cases. The main exception, run 5, is interesting: one of the experimental subjects, by error or by design, made at $t = 25$ a very low prediction, which caused a large downward jump of the actual price realisation and the subsequent price predictions. The price predictions returned to the many-steps-ahead prediction later on. Also in run 4 there are significant differences between the switching OLS model and the true realisation; but even there, the alternative model tracks the actual price evolution much better than the rational expectations model.

The simulated predictions $p^e_{\text{model},t}$ are again compared, for both models, to the average

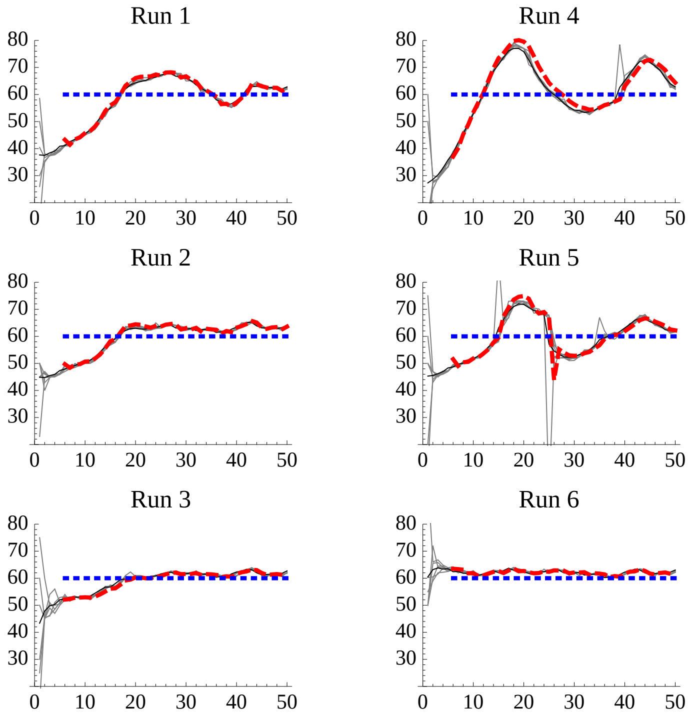

Figuur 4: *As figure 3, but for the positive feedback treatment.*

predictions $\overline{p}^{e}_{\text{lab},t}$ obtained from the experimental subjects, by computing the mean absolute prediction error (13). In the negative feedback treatment, equality of the mean absolute prediction errors of the two models is rejected only for one of the runs. This is slightly surprising, as the homogeneous rational expectations model outperformed the switching OLS model significantly in one-step-ahead forecasting. It is explained by the fact that the compared error time series are much shorter in the many-steps-ahead comparison. Moreover, most of the 'irregularity' of the time series is in the initial segment, which helps the switching OLS rule to perform better.

For the positive feedback treatment, equality is rejected for all but one of the runs, the switching OLS predictions outperforming rational predictions significantly in all other runs. This is even the case in runs 3 and 6, where the dynamics has all but converged on the rational expectations predictions.

3.4.4 *Interpretation.*

The comparison of the homogeneous rational expectations model to the switching OLS model stacks the cards in favour of the former: to compute the rational benchmark, the complete structure of the model has to be known, and this knowledge is denied to the

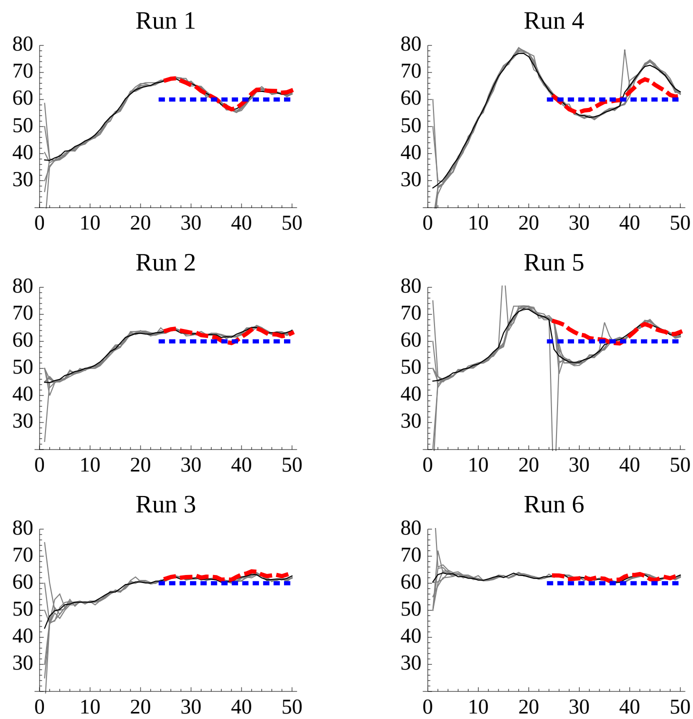

Figuur 5: *Many-steps-ahead predictions for the positive feedback treatment. Homogeneous rational expectations (blue, dotted), switching ordinary least squares predictions (red, dashed), and actual predictions by experimental subjects (black, solid).*

subjects by design of the experiments. Put this way, the surprising aspect of the results is not that the rational benchmark performs poorly in the positive feedback treatment, but that it performs so well in the negative feedback treatment.

A heuristic order-of-magnitude argument may help to understand the underlying mechanism. Consider the evolution equation (4), assume for the moment that all agents make the same prediction p_t^e at time t, and introduce the quantities x_t^e and e_t: the former is the deviation of the expected price from the rational benchmark $p^* = \alpha/(1-\beta)$; the latter is the prediction error:

$$x_t^e = p_t^e - p^* \quad \text{and} \quad e_t = p_t - p_t^e. \tag{15}$$

Subtracting p_t^e from both sides of equation and rewriting transforms equation (4) into

$$e_t = (\beta - 1)x_t^e + \varepsilon_t; \tag{16}$$

note that this representation is not available to the experimental subjects. Taking squares and expectations yields

$$\mathbb{E}e_t^2 = (1-\beta)^2 \left(x_t^e\right)^2 + \mathbb{E}\varepsilon_t^2. \tag{17}$$

As the pay-off of the subjects is a decreasing function of e_t^2, their actions will tend to push $\mathbb{E}e_t^2$ as far as possible towards zero. Now e_t consists of a 'signal' term $(\beta-1)x_t^e$ and a 'noise' term ε_t. Heuristically speaking, their actions will amount to reducing the signal term, over which they have some control, in such a way that it is at least not significantly larger than the noise term. Convergence towards the fundamental steady state is by this argument expected to be fast until the two terms are of equal order of magnitude; afterwards, the signal is harder to distinguish from the noise. In particular, the terms are equal if

$$|x_t^e| = \frac{\sqrt{\mathbb{E}\varepsilon_t^2}}{1-\beta}. \tag{18}$$

As $\mathbb{E}\varepsilon_t^2 = 1/4$, in the negative feedback treatment $\beta = -0.95$ equation (18) gives rise to a typical price deviation $|x_t^e| = 0.26$, whereas in the positive feedback treatment $\beta = 0.95$, this value is equal to 10. A visual comparison of the order of magnitudes of these deviations to those of the experimental deviations from the fundamental price recorded in figures 1 and 2 suggests that the orders are correct.

This is borne out by the results of Sonnemans & Tuinstra (2010), where the same dynamics is studied, but with $\beta = 0.667$. This corresponds to a typical price deviation of $|x_t^e| = 1.50$; the heuristic argument predicts fast convergence until this level is reached, which is corroborated by the results (see Sonnemans & Tuinstra, 2010, figure 6).

Concluding, the differences in the deviations from the rational expectations equilibrium are a consequence of the different institutional arrangements in the two treatments, and are not caused by rationality or lack of rationality of the experimental subjects.

3.5 Summary

In the Heemeijer et al. (2009) experiment, almost all agents made unbiased predictions: for a large majority of them, the mean expectation error was not significantly different from zero.

To test the homogeneous rational expectations hypothesis, the assumption had to be made that the preferences of the agents were perfectly described by the individual earnings function (1), and that agents did not put any value on leisure time while performing in the experiment.

First, in both treatments, the rational expectations equilibrium was obtained as an asymptotic limit of the expectations dynamics. In the negative treatment, convergence to this limit was fast and occurred within the observation period; in the positive treatment, it was weak and had to be established by extrapolation methods. It is conjectured that the convergence rate is of the same order of magnitude as the feedback strength; if true, this has far-reaching consequences for the dynamics of asset markets.

Second, the fit of an alternative parameter-free model of aggregate expectation formation to the data was compared to that of the homogeneous rational expectations model. For one-step-ahead predictions, the rational model outperformed the alternative significantly in the negative feedback treatment, while these roles were reversed in the positive feedback treatment. But while in the former case, the maximal relative error of the alternative model was about 3% compared to 2% of the rational model, in the latter case the maximal relative error of the rational model was about 14% against 3% of the alternative model.

Finally, for many-step-ahead predictions under the same noise realisations as in the experiments, both models performed more or less equally well in the negative feedback treatment, while the alternative model was again significantly better in the positive feedback treatment.

If the requirement is made that a model should be able to explain structurally different situations, then on basis of these experimental data it can be concluded that though rational expectations constitute a reasonable first approximation, the switching OLS model should be preferred, mainly on the strength of its superior performance in the positive feedback situation.

3.6 Structural breaks

The switching OLS model has been introduced only as a proof of the assertion that there are parameter-free learning models that can outperform the homogeneous rational expectation model. A more sophisticated approach should feature some monitoring device for structural changes in the environment. This point becomes clear immediately when the results of the structural break experiment of Bao et al. (2012) are considered.

The design of the experiment was almost identical to the Heemeijer et al. (2009) experiment, excepting the following: the number of rounds in a run was increased from 50 to 65, and the variance of the random disturbance ε_t was decreased to $(3/10)^2$. The major design change was however that the parameter α, which determines the level of the fundamental price, changed twice during the experiment. The fundamental price equalled 56 during the first 20 rounds, 41 during the next 23, and 62 during the final 22 rounds. The feedback strength β was constant over the experiment, equal to -0.95 in the negative feedback treatment, and equal to 0.95 in the positive feedback treatment.

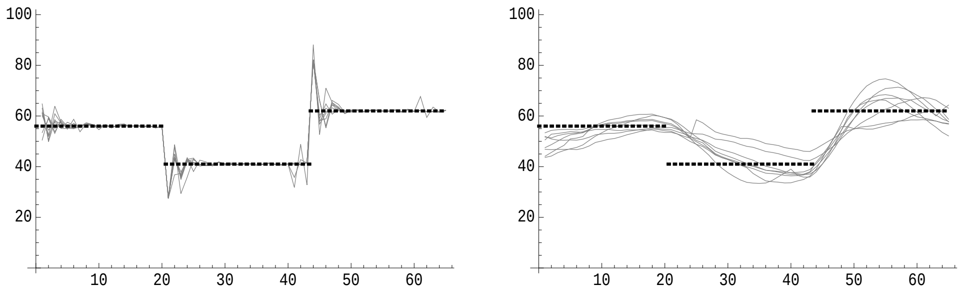

(a) *Negative feedback* **(b)** *Positive feedback*

Figuur 6: *Price realisations in the structural break experiment of Bao et al. (2012). Solid grey lines are realised prices per run; dotted black lines are fundamental prices (cf. Bao et al. (2012), figure 1).*

The resulting price realisations are shown in figure 6. The familiar pattern of fast convergence in the negative feedback treatment and transient oscillations in the positive feedback treatment can be observed here as well. What is remarkable is that the experimental subjects immediately notice the structural break — this is most readily apparent in the negative feedback treatment — and are prepared to forget what they have learned about the process almost instantly. In the positive feedback treatment, there is much more inertia of the predictive dynamics, but also there the agents notice the changed circumstances. The interpretation of these results runs along the same lines as for the Heemeijer et al. (2009) experiment.

4 Conclusion

The homogeneous rational expectations hypothesis is not a testable statement about the expectation formation of real agents in a noisy economic framework. It is, or rather it should be, a device to be used by economic theoreticians to arrive at an upper limit of the economic performance of agents.

The experiments reviewed in this article elucidate the relation between institutional arrangements and the possibility that collective predictions of initially uninformed human subjects may converge towards the rational benchmark. In the simple context of the experiments, the institutional arrangements provide an information feedback structure that ties the prediction problem, which is facing the subjects directly, to the problem of convergence towards the rational benchmark, which is of no interest to the subjects, given their incentive structure. If the information feedback is weak, the experiments reviewed show that the homogeneous rational expectations benchmark performs poorly as a predictor of subjects behaviour and is easily outperformed by a model that takes into account learning behaviour.

In the context of the Heemeijer et al. (2009) experiment, which is favourable to the homogeneous rational expectations hypothesis, it seems that Walras was right: the hypothesis may be admitted as a first approximation. For other contexts, as the asset pricing context in Hommes et al. (2005), this is more doubtful. But even in this favourable context, we are moving to a stage where the frictions mentioned by Walras can be measured and modelled as well.

Referenties

Adam K, Marcet A. 2011. Internal rationality, imperfect market knowledge and asset prices. *Journal of Economic Theory* 146(3):1224–52

Arrow KJ. 1986. Rationality of Self and Others in an Economic System. *The Journal of Business* 59(4):S385–99

Axelrod R. 1984. *The evolution of cooperation.* New York: Basic Books

Bao T, Duffy J, Hommes CH. 2013. Learning, forecasting and optimizing: An experimental study. *European Economic Review* 61:186–4

Bao T, Hommes CH, Sonnemans J, Tuinstra J. 2012. Individual expectations, limited rationality and aggregate outcomes. *Journal of Economic Dynamics and Control* 36:1101–20

Branch WA, Evans GW. 2010. Asset return dynamics and learning. *Review of Financial Studies* 23(4):1651–80

Bray MM, Savin NE. 1986. Rational expectations equilibria, learning, and model specification. *Econometrica* 54(5):1129–60

Brock WA, Hommes CH. 1997. A rational route to randomness. *Econometrica* 65(5):1059–95

Brock WA, Hommes CH. 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control* 22:1235–74

Brock WA, Hommes CH, Wagener FOO. 2009. More hedging instruments may destabilize markets. *Journal of Economic Dynamics and Control* 33:1912–28

Evans GW, Honkapohja S. 1995. Adaptive Learning and Expectational Stability: An Introduction. In A Kirman, M Salmon, eds., *Learning Rationality in Economics*, chap. 4, pp. 102–26. Oxford: Basil Blackwell

Friedman BM. 1979. Optimal expectations and the extreme information assumptions of "rational expectations" macromodels. *Journal of Monetary Economics* 5(1):23–41

Guesnerie R. 1992. An exploration of the eductive justifications of the rational expectations hypothesis. *American Economic Review* 82(5):1254–78

Guesnerie R. 2002. Anchoring economic predictions in common knowledge. *Econometrica* 70(2):439–80

Heemeijer P, Hommes CH, Sonnemans J, Tuinstra J. 2009. Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. *Journal of Economic Dynamics and Control* 33(5):1052–72

Hirschman A. 1977. *The passions and the interests*. New Jersey: Princeton University Press

Hommes CH. 2011. The heterogeneous expectations hypothesis: Some evidence from the lab. *Journal of Economic Dynamics and Control* 35(1):1–24

Hommes CH, Sonnemans J, Tuinstra J, van de Velden H. 2005. Coordination of expectations in asset pricing experiments. *Review of Financial Studies* 18(3):955–80

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

Hommes CH, Sonnemans J, Tuinstra J, van de Velden H. 2008. Expectations and bubbles in asset pricing experiments. *Journal of Economic Behavior & Organization* 67(1):116–33

Jaffé W, ed. 1965. *Correspondence of Léon Walras and Related Papers*, vol. III. Amsterdam: North-Holland

Lei V, Noussair C, Plott C. 2001. Nonspeculative bubbles in experimental asset markets: Lack of common knowledge of rationality vs. actual irrationality. *Econometrica* 69(4):831–59

Lucas RE. 1973. Some international evidence on output-inflation tradeoffs. *The American Economic Review* 63(3):326–334

Lucas RE, Prescott E. 1971. Investment under uncertainty. *Econometrica* 39(5):659–81

Lucas RE, Sargent TJ. 1979. After Keynesian Macroeconomics. *Federal Reserve Bank of Minneapolis Quarterly Review* 3(2):1–16

Marcet A, Sargent TJ. 1989. Convergence of least-squares learning in environments with hidden state variables and private information. *The Journal of Political Economy* 97(6):1306–22

Marimon R, Sunder S. 1994. Expectations and learning under alternative monetary regimes: an experimental approach. *Economic Theory* 4(1):131–62

McCallum BT. 1976. Rational Expectations and the Natural Rate Hypothesis: Some Consistent Estimates. *Econometrica* 44(1):43–52

Modigliani F. 1977. The monetarist controversy; or, should we forsake stabilization policies? *The American Economic Review* 67(2):1–19

Muth J. 1961. Rational expectations and the theory of price movements. *Econometrica* 29(3):315–35

Sargent TJ, Fand D, Goldfeld S. 1973. Rational expectations, the real rate of interest, and the natural rate of unemployment. *Brookings Papers on Economic Activity* 1973(2):429–480

Sargent TJ, Wallace N. 1975. "Rational" Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule. *The Journal of Political Economy* 83(2):241–54

Shaw GK. 1984. *Rational expectations*. Brighton: Wheatsheaf Books

Sims CA. 1998. Stickiness. In *Carnegie-Rochester Conference Series on Public Policy*, vol. 49, pp. 317–56

Sims CA. 2003. Implications of rational inattention. *Journal of Monetary Economics* 50(3):665–90

Sims CA. 2005. Rational inattention: a research agenda. *Tech. rep.*, Volkswirtschaftliches Forschungszentrum der Deutschen Bundesbank

Smith VL. 1991. Rational choice: The contrast between economics and psychology. *Journal of Political Economy* 99(4):877–97

Smith VL, Suchanek GL, Williams AW. 1988. Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica* 56(5):1119–51

Sonnemans J, Hommes CH, Tuinstra J, van de Velden H. 2004. The instability of a heterogeneous cobweb economy: a strategy experiment on expectation formation. *Journal of Economic Behavior & Organization* 54(4):453–81

Sonnemans J, Tuinstra J. 2010. Positive expectations feedback experiments and number guessing games as models of financial markets. *Journal of Economic Psychology* 31(6):964–84

Sunder S. 1995. Experimental asset markets: A survey. In J Kagel, A Roth, eds., *The handbook of experimental economics*, pp. 445–500. Princeton: Princeton University Press

Tuinstra J, Wagener FOO. 2007. On learning equilibria. *Economic Theory* 30(3):493–513

Turnovsky SJ. 1970. Empirical Evidence on the Formation of Price Expectations. *Journal of the American Statistical Association* 65(332):1441–1454

Turnovsky SJ, Wachter ML. 1972. A Test of the "Expectations Hypothesis"Using Directly Observed Wage and Price Expectations. *The Review of Economics and Statistics* 54(1):47–54

Woodford M. 2009. Information-constrained state-dependent pricing. *Journal of Monetary Economics* 56:S100–24