

Expectations in Experiments

Florian Wagener^{* †‡}

16th August 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

^{*}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

agents have perfect foresight:

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 mécanisme 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 Muthian, weak and strong 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 ‘collective’ version of the rational expectations hypothesis; Lucas & Prescott (1971) formulated a ‘strong individual’ version ten years later. Outside of economics, yet another definition of rationality exist, which in this article will be referred to as the ‘weak individual’ hypothesis.

As rational expectation models outperformed models based on naïve or adaptive expectations in empirical work, the hypothesis as formulated by Lucas and Prescott gained broad acceptance. Its radical implications simultaneously attracted severe criticism (Arrow, 1986). The weak hypothesis is regularly used to justify the strong version, based on an implicit assumption of equivalence of the two. Theoretical investigations however invariably use the strong hypothesis.

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, it was shown that a group of agents eventually arrives collectively at the rational expectations equilibrium through market interactions. This occurred after a learning phase, taking 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, it was shown that at least during the learning phase, the behaviour of some of the agents is irrational 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 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. The two approaches can be combined (e.g. Tuinstra & Wagener, 2007; Branch & Evans, 2010).

An assumption that is usually treated as a direct corollary of the rational expectations hypothesis, but which is actually separate from it, is that of homogeneity of behaviour. Arrow (1986) points out that the Mantel-Sonnenschein-Debreu theorem implies that rationality of agents' behaviour by itself puts no restrictions on possible economic dynamics. Only adding the homogeneity hypothesis allows to make powerful predictions.

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).

Recently, a number of learning-to-forecast experiments have been conducted 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 (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 presumably used 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 COLLECTIVE 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). First Muth:

“(...) 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 a time series of expectations of an economic variable is known, as well as a time series of its realisations, Muth's collective rational expectations hypothesis can be tested.

2.2 THE WEAK (INDIVIDUAL) RATIONAL EXPECTATIONS HYPOTHESIS

The common sense notion of rationality is well expressed by Rawls (1971):

“One might reply that the rationality of a person's choice does not depend upon how much he knows, but only upon how well he reasons from whatever information he has, however incomplete. Our decision is perfectly rational provided that we face up to our circumstances and do the best we can.” (quoted after Friedman, 1979, p. 23)

This notion has variously been called ‘informational efficiency’ (Hommes et al., 2005) or ‘internal rationality’ (Adam & Marcet, 2011). In the context of expectation formation, it requires individual expectation errors to be unbiased and to have no significant autocorrelation structure. Also this hypothesis is testable on the basis of observed data.

2.3 THE STRONG (INDIVIDUAL) 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 rational expectations are 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, \dots, 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 collective version, the strong individual hypothesis assumes an internal model of economic agents, and it states explicitly that the expectations of all agents are the same.

A different but equivalent formulation has been given by Sargent (1987) in terms of perceived and actual laws of motion: agents base their actions on their observations, on their expectations of certain key economic variables, and on their dynamic model of the economy, their ‘perceived law of motion’. The actions in turn determine the aggregate dynamics of the economy, the ‘actual law of motion’ and the realisations of the key variables. In this context expectations are rational if the expected distribution and the realised distribution of the key variables coincide.

The strong 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. Moreover, the wording implies that the hypothesis only describes 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 strong 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 strength of the strong individual hypothesis, it is legitimate to inquire after the secrets of its phenomenal success. One of these must be that it was introduced at a 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 strong individual 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. In fact, the hypothesis is the strongest statement of what it means for economic agents to be rational, and it fulfils an important benchmark function in any kind of economic model.

Another reason for its success 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

Since the establishment of CeNDEF (Center for Nonlinear Dynamics in Economics and Finance) in 1998, a large number of learning-to-forecast experiments have been conducted there. They all share the characteristic that the forecasts determine the evolution of the realised variable: this creates a feedback loop from observations to data and then back to observations, which the subjects try to learn.

The foremost result of these experiments is that the structure of the economic system determines the stability of the rational expectations equilibrium, and if it is stable, the speed of convergence towards it.

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 and the split in learning-to-forecast and learning-to-optimize experiments. 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 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. 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 as experimental subjects.

There were four rounds, to each of which at least twenty strategies were submitted. For the first three rounds, the strategy with the lowest average squared prediction error in a round won a prize of about 23 euros; in the last round, three prizes were awarded to the top three performing strategies, of about 114, 68 and 23 euros respectively.

The mean quadratic distance of the price dynamics to the rational expectations benchmark decreased monotonically over the four rounds, with the value in the last round about 25% of that of the first. But only 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 in dynamic and informational complexity as competition between the subjects intensifies. The fact that rational expectations may be the only surviving strategy in the limit of infinite time series and infinite computational resources of the subjects 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).

3.2 THE IMPACT OF MARKET STRUCTURE

The laboratory experiments Hommes et al. (2005, 2008) investigated the impact of institutional structure on expectation dynamics. They were direct forerunners of the Heemeijer et al. (2009) experiment discussed at some length below.

In both experiments, subjects were required to forecast price series; as in the Sonnemans et al. (2004) experiments, the dynamics of these series depended on the forecasts. In Hommes et al. (2005), the evolution equation, reminiscent of an asset pricing market, is given by

$$p_t = \frac{1}{R} \left((1 - n_t) \frac{1}{6} \sum_{i=1}^6 p_{i,t+1}^e + n_t p^* + \bar{y} + \varepsilon_t \right); \quad (1)$$

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 $p_{i,t+1}^e$ the individual expectations of the six experimental subjects of next period's price. Finally $n_t = 1 - \exp(-|p_{t-1} - p^*|/200)$ 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 (1) is that of a positive feedback loop: an increase in the average price expectations effects an increase of the realised price. The resulting expectation dynamics are of three types: 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 a brief coordination phase at the beginning of the run, 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

$$p_t = \alpha - \beta \sum_{i=1}^6 S(p_{i,t}^e) + \varepsilon_t; \quad (2)$$

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: $\sigma = -0.87$, -1.96 and -7.75 respectively. That is, the dynamics are eductively stable in the first treatment, but eductively unstable in the second and third treatment (cf. Guesnerie, 2002).

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

These experiments reinforce the observations of Smith (1991) that institutions determine whether individual behaviour can converge to rational behaviour as well the speed of the convergence. In fact, if the informational feedback is negative, even in a moderately unstable situation rapid convergence towards the rational expectations equilibrium can be obtained: this of course reflects the efficiency of most commodity markets.

Under positive informational feedback, the collective prediction behaviour of the experimental subjects is much more indeterminate. One of the factors that complicates the interpretation of the Hommes et al. (2005) experiment is that in the asset pricing experiment – without robot traders and without ceilings on predictions – subjects predict the price next period before the price in the current period is revealed. In particular, the system has a one-dimensional family of rational expectation solutions.

3.3 NEGATIVE VS POSITIVE INFORMATIONAL FEEDBACK

The desire to remove the indeterminateness of the Hommes et al. (2005) experiment and 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, strong 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 strong rational expectations hypothesis can be expected to have its best shot.

3.3.1 Setup

In the experiment, the subjects, mostly students, were required to make price predictions of a price time series. On the computer screen facing them, they were provided with a numerical table of their past predictions, the past price realisations, a graphical representation of these two time series, their earnings last period and the total earnings. When making their first prediction, they had no information whatsoever. The subjects were only qualitatively informed about the nature of the feedback structure in the price forming mechanism facing them. Each run of the experiment involved a group of six students making price predictions in each round, for fifty consecutive rounds. A round took on average slightly less than two minutes. Each subject participated in a single run only.

When all price predictions p_{it}^e were made, the realised price p_t was computed according to the relation

$$p_t = \alpha + \beta \left(\frac{1}{6} \sum_{i=1}^6 p_{it}^e \right) + \varepsilon_t. \quad (3)$$

The random terms ε_t were independently distributed normal variables with mean 0 and variance 1/4. After realisation of the price, individual period earnings E_{it} were determined according to

$$E_{it} = 1300 \max \left\{ 1 - \left(\frac{p_{it}^e - p_t}{7} \right)^2, 0 \right\}. \quad (4)$$

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; if all agents had rational expectations in all periods, expected earnings would have been 24.87 euros. The average of realised earnings over all runs was about 22 euros, that is, 88% of the 'rational expectations' value.

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 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.

3.3.2 Results

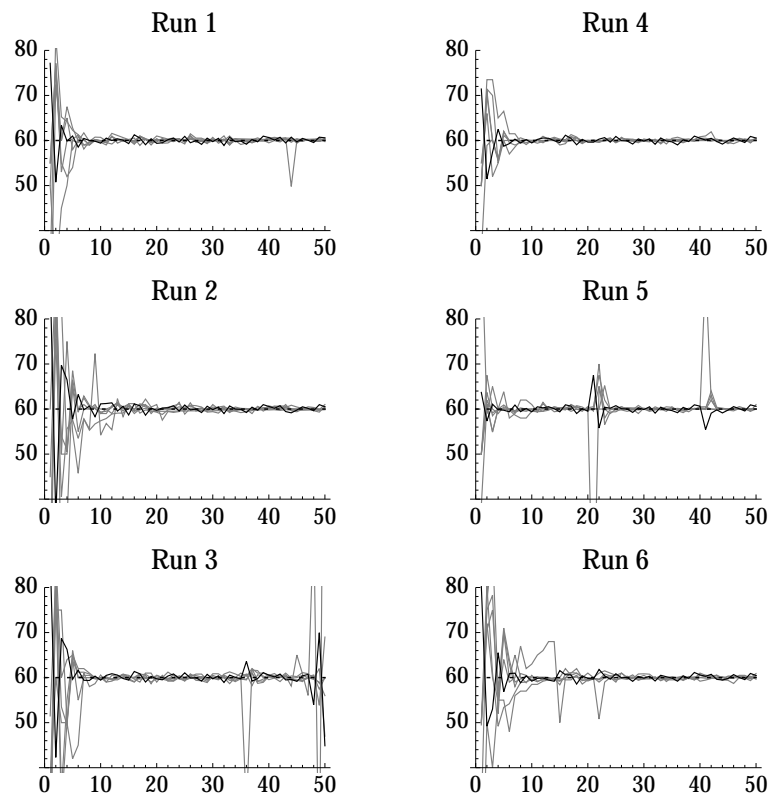


Figure 1: Negative feedback treatment of the (Heemeijer et al., 2009) experiment: price expectations (grey lines), realised prices (black line) and rational expectations benchmark (dotted line).

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

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

Clearly, some of the subjects tried to experiment for some periods, but were probably discouraged by their low earnings. In one run agents were expecting an end-of-period effect, which destabilised the price in the last few periods.

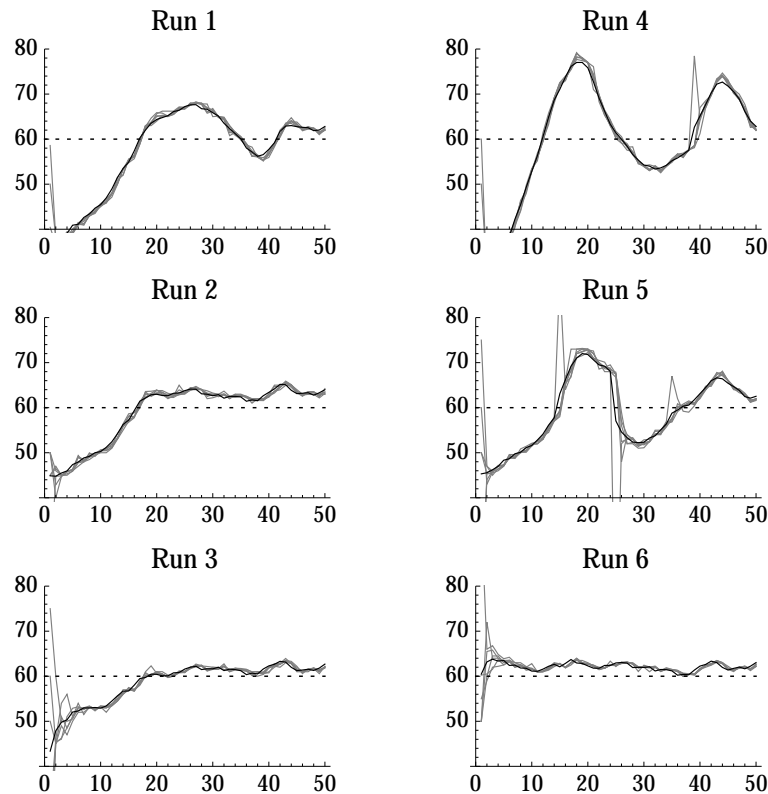


Figure 2: As figure 2: positive feedback treatment.

The results of the positive feedback treatment are given in figure 2. 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, but the speed of convergence towards the rational expectations value differs greatly over the runs. There is substantial endogenous dynamics, which moreover varies qualitatively: in some runs it converged monotonously towards the rational expectations value; in others it converged but oscillated. Again, as in the negative feedback treatment, sub-

jects tend to experiment occasionally.

Weak rationality of experimental subjects. As all individual predictions are available, the hypothesis that all agents are weakly rational can be tested. The expectation errors $e_t = p_t - p_t^e$ have been computed; in order to ensure that coordination has taken place, only the realisations of e_t for $11 \leq t \leq 50$ have been considered. Under the weak rational expectations hypothesis, the mean of the e_t is zero. The p-values of the t-test of this hypothesis have been computed.

For the negative feedback treatment, the hypothesis is rejected 5 times, out of 36, on the 10% level; for positive feedback, there are 3 rejections, out of 36, at the same level. This suggests rather strongly that almost all agents are weakly individually rational; moreover, the cases where weak rationality is rejected are probably due to experimentation. It may be argued that this is rational behaviour as well, though costly, to find out about the behaviour of the system in as yet unexplored regions of the phase space.

Testing for autocorrelations in the prediction errors using the Durbin-Watson test rejects weak rationality for 8 out of 72 subjects (both treatments) on a 5% significance level. This is slightly higher than the expected rejection rate of 4, but not enough to overturn the conclusion that all subjects are weakly rational, especially taken into account the fact that some subjects were occasionally experimenting.

Convergence to strongly rational expectations. As the strong 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 strong rational expectations to be the correct description of actual expectations; if information arrives at a faster rate, transient dynamics will play a significant role.

To obtain information about these transients, a model of the subjects behaviour has to be estimated. Heemeijer et al. (2009) have estimated 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_{it}^e in (3) and running the price evolution with the same noise realisations as used in the experiment gives fits of varying quality: a particularly bad one is shown in figure 3a. This is due to some estimation rules not being rational in the sense that there is significant correlation in the autocorrelations. If these are removed, the fits improve greatly (figure 3b). For

instance, the rule of the eliminated nonrational agent in figure 3a, reads as $p_t^e = 1.956 + 0.9112p_{t-1}$, and it implies a fundamental value of 1.02. What is more, in the two runs that diverge most from the estimated trajectories, that is, run 4 after $t = 40$ and run 5 after $t = 15$, the expectation of a single subject is radically deviating from all others, either due to a typing mistake (as probably occurred in run 5 at $t = 25$) or due to the wish to explore other regions. If this is taken into account, the fit of the rational rules to the data is remarkably good.

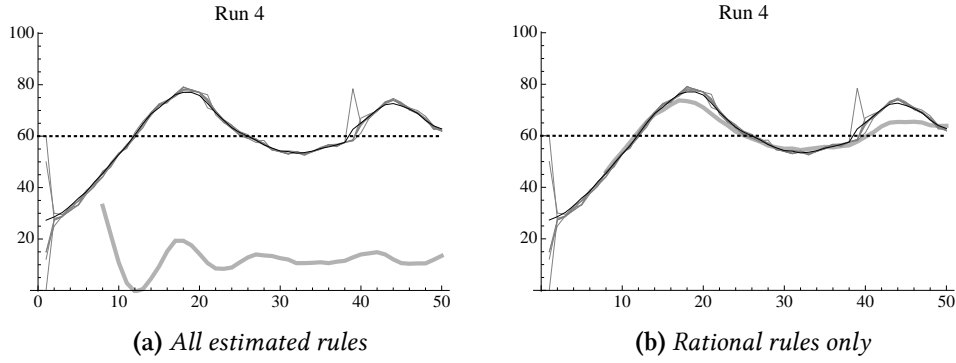


Figure 3: Positive feedback experiment, run 4: price expectations (thin), realised prices (thick) and evolution based on estimated prediction rules (thick grey) of all agents (left) and rational agents (right).

Convergence rates are obtained as follows. Runs are indexed by $k = 1, \dots, 6$. Using estimated forecasting rules for ‘rational’ agents only, for each run price time series are obtained by running the evolution law (3) without noisy perturbations. Then a linear model of the form

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

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

$$p_k^* = \frac{\hat{c}_{k0}}{1 - \sum_{i=1}^3 \hat{c}_{ki}}. \quad (6)$$

The implied rate of convergence r_k is defined as

$$r_k = -\log |\lambda_k|, \quad (7)$$

where λ_k is the root λ of the homogeneous equation

$$\lambda^3 - \sum_{i=1}^3 \hat{c}_{ki} \lambda^{3-i} = 0 \quad (8)$$

having the largest complex absolute value. The results are shown in table 1.

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

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

All implied steady states are close to the rational expectations value. The p-value of a t-test of the hypothesis that the population mean of the runs equals the rational value of 60 is equal to 0.45. This supports the strong rational expectations hypothesis, in that prices converge to the rational expectations equilibrium values.

The convergence rates are low, however, and it might be conjectured that their population mean equals the feedback strength $r_\beta = -\log 0.95 \approx 0.0513$. But this is rejected by a t-test at a p-value of 0.032. Nevertheless, it may be conjectured that in a positive feedback situation the actual convergence rate is of the same order of magnitude as the feedback strength.

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; the lower the risk free rate, the longer it takes to converge to the rational price.

3.3.3 *Strong rationality*

One of the major virtues of the rational expectations hypothesis is that it provides a universal model for agent behaviour; in the simple situations considered in the context of the Heemeijer et al. (2009) experiment, rational behaviour is even uniquely determined.

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 strong rational hypothesis and its alternative are to be compared. 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 rational expectations hypothesis.

In constructing this model, two goals are aimed at: firstly, the model should be externally consistent with the observed time series, that is, it should give a good prediction of the observed data. Secondly, the model should be internally consistent with the data, that is, it should use only data that are known to the experimental subjects. In particular, in the context of the Heemeijer et al. (2009) experiment, rational expectations do not fit the second requirement, as the subjects have been given only general qualitative information about the market they are facing.

Furthermore, the alternative model is based on the time-honoured induction principle: as the subjects do not have access to the actual law of motion of the system, they 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 model will be denoted in the following as the “switching OLS” (sOLS)

model. Agents keep track of three prediction rules, indexed by $k = 0, 1, 2$; these are ordinary least squares predictions of the form

$$p_{k,t}^e = \hat{c}_{k0,t} + \sum_{i=1}^k \hat{c}_{ki,t} p_{t-i}, \quad (9)$$

where the $\hat{c}_{ki,t}$ are obtained by fitting the available priced time series up to p_{t-1} on the prediction rule (9). Let k_t be the index of the rule with the lowest total squared prediction error based on the data up to and including p_{t-1} ; that is, the value of k minimising $\sum_s (p_{k,s}^e - p_s)^2$. Then

$$p_{\text{sOLS},t}^e = p_{k_t,t}^e \quad (10)$$

is the new prediction of the sOLS model.

One-step-ahead. First, the one-step-ahead prediction quality of the two models is compared; that is, the prediction rules (9) are progressively estimated on the price time series p_t realised in the experiments. To compare the quality of the predictions, the mean of the absolute prediction error has been computed for both models, as well as the p-value of the hypothesis that these means are equal. The results are shown in tables 2 and 3.

Negative feedback				
Run	Rational	sOLS	difference	p-value
1	0.2555	0.5235	-0.2680	2.0×10^{-4}
2	0.4628	0.9316	-0.4688	7.8×10^{-6}
3	1.1057	1.6313	-0.5257	2.4×10^{-5}
4	0.1429	0.5313	-0.3885	2.6×10^{-8}
5	0.5733	0.8808	-0.3075	2.7×10^{-7}
6	0.4003	0.8403	-0.4399	6.7×10^{-5}

Table 2: Mean absolute one-step-ahead prediction errors of strongly rational expectations and sOLS predictions respectively, their differences, and the p-value on the test whether these means are equal. Negative feedback case.

The results depend strongly on the feedback structure. Rational expectations outperform switching OLS expectations significantly in all situations of the negative feedback treatment; yet both models perform well, with maximal mean absolute errors of about 2% of the fundamental for rational expectations, and about 3% for the alternative.

Positive feedback				
Run	Rational	sOLS	difference	p-value
1	5.9409	0.7303	5.2106	2.0×10^{-13}
2	4.2762	0.6413	3.6349	1.7×10^{-15}
3	2.5054	0.5743	1.9310	3.8×10^{-9}
4	8.2404	1.8618	6.3787	1.7×10^{-9}
5	6.1893	1.5175	4.6718	6.7×10^{-10}
6	2.0796	0.4837	1.5959	7.7×10^{-18}

Table 3: Same as table 2, positive feedback case.

In the positive feedback treatment however, the alternative model outperforms rational expectations in all situation. Moreover, its maximal mean absolute error is again about 3% of the fundamental price, while this is 14% for rational expectations.

Figure 4 shows, for the positive feedback treatment, the price realisations and the one-step ahead predictions of both models.

Many-steps-ahead. To assess the performance of the switching OLS model for many-steps-ahead predictions, the rules (9) are estimated on an initial segment $t_0 \leq t \leq t_1$ of the time series obtained in the experiment. For $t > t_1$, price realisations p_t are obtained by substituting price expectations p_t^e obtained from (9) into

$$p_t = \alpha + \beta p_t^e + \varepsilon_t, \quad (11)$$

where α and β are as in (3), and where ε_t are the same noise realisations as in the experiment. That is, after ‘training’ the model on an initial segment, its long-range predictive power is assessed.

The beginning of the training interval has been chosen at $t_0 = 4$ to ensure 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.

The results are given in tables 4 and 5. 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 slightly surprising result is explained by the fact that the compared error time series are 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

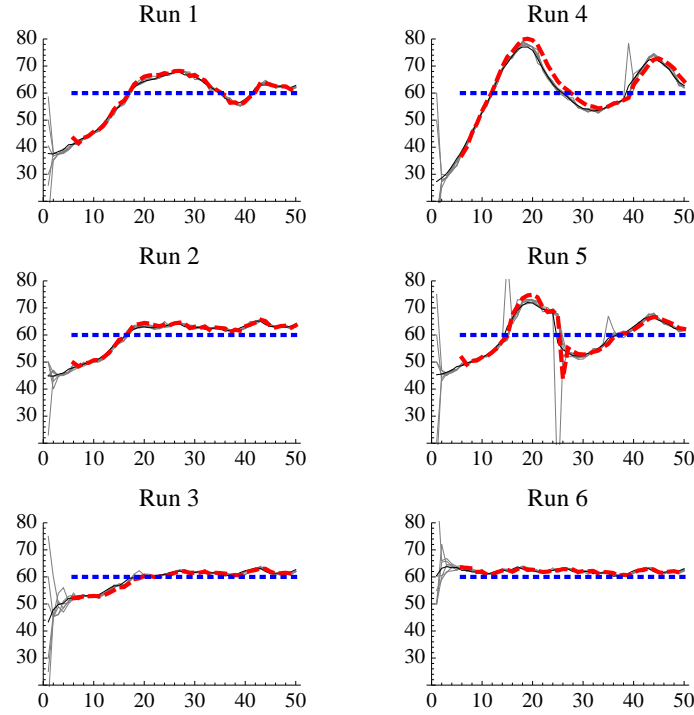


Figure 4: *One-step ahead predictions for the positive feedback treatment. Rational expectations (blue, dotted) and switching ordinary least squares predictions (red, dashed).*

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.

3.3.4 Interpretation

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

To understand the underlying mechanism, assume for the moment that all agents make the same prediction p_t^e at time t , and introduce the deviation x_t^e of the expected price

Negative feedback				
Run	Rational	sOLS	difference	p-value
1	0.2361	0.2704	-0.0343	0.7627
2	0.1609	0.2317	-0.0706	0.1070
3	1.6520	1.5838	0.0682	0.9927
4	0.1287	0.1487	-0.0200	0.6671
5	0.3496	0.5855	-0.2359	1.7×10^{-3}
6	0.1509	0.1837	-0.0328	0.2710

Table 4: Mean absolute many-steps-ahead prediction errors of rational expectations and sOLS predictions respectively, their differences, and the p-value on the test whether these means are equal. Negative feedback case.

Positive feedback				
Run	Rational	sOLS	difference	p-value
1	3.6864	0.4035	3.2830	1.5×10^{-7}
2	3.2007	0.8563	2.3444	4.6×10^{-14}
3	1.6935	0.9153	0.7783	1.5×10^{-5}
4	5.8694	2.9301	2.9393	1.2×10^{-3}
5	4.3558	4.1553	0.2006	0.3011
6	1.9024	0.3845	1.5179	1.8×10^{-8}

Table 5: As table 4, positive feedback case.

from the rational benchmark p^* and the prediction error e_t as

$$x_t^e = p_t^e - p^* \quad \text{and} \quad e_t = p_t - p_t^e. \quad (12)$$

In x_t^e and e_t , the system evolution equation (3) reads as

$$x_t = x_t^e + e_t = \beta x_t^e + \varepsilon_t; \quad (13)$$

subtracting x_t^e on both sides, taking squares and expectations yields

$$\mathbb{E}e_t^2 = (1 - \beta)^2 (x_t^e)^2 + \mathbb{E}\varepsilon_t^2. \quad (14)$$

As the pay-off of the subjects is a decreasing function of e_t^2 , their principal interest is to push $\mathbb{E}e_t^2$ as far as possible towards zero. They will strive to choose the first term on the right hand side of equation (14) in such a way that it is not significantly larger than the second term. Convergence towards the fundamental steady state is by this

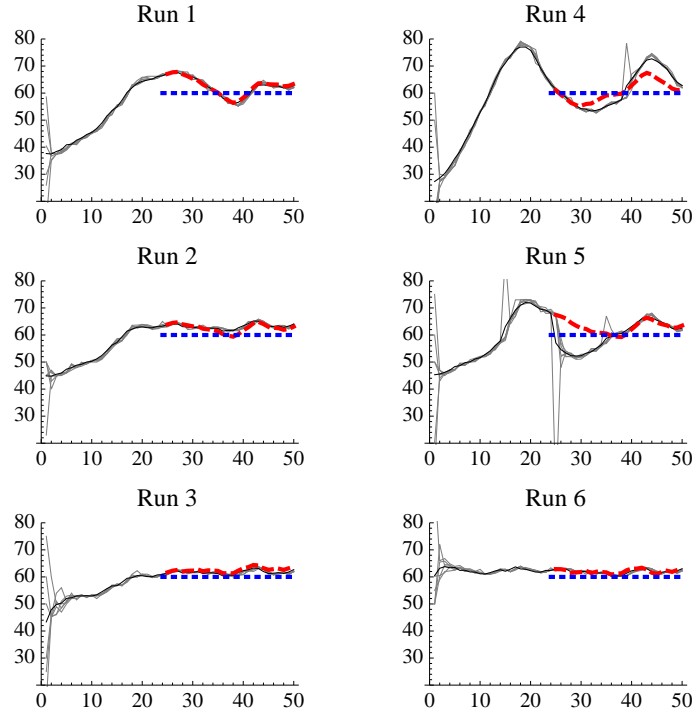


Figure 5: *Many-steps-ahead predictions for the positive feedback treatment. Rational expectations (blue, dotted) and switching ordinary least squares predictions (red, dashed).*

argument expected to be fast until the two terms are of equal order of magnitude. In particular, these terms are equal if

$$|x_t^e| = \frac{\sqrt{\mathbb{E}\varepsilon_t^2}}{1 - \beta}. \quad (15)$$

As $\mathbb{E}\varepsilon_t^2 = 1/4$, in the negative feedback treatment $\beta = -0.95$ equation (15) 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 with figures 1 and 2 confirms that these are the correct orders of magnitude.

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 $|x_t^e| = 1.50$ and fast convergence until this level is reached, both of 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 treat-

ments, and are not caused by rationality or lack of rationality of the experimental subjects.

3.4 SUMMARY

In the experiment, almost all agents made weakly rational predictions: for a large majority of the agents, the mean expectation error was not significantly different from zero, and there were no significant autocorrelations in expectation errors.

To test the strong hypothesis, the assumption had to be made that the preferences of the agents were perfectly described by the individual earnings function (4), and that agents did not put any value on leisure time while performing in the experiment.

First, in both treatments, rational expectations were 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 equal, or at least 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 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 its far superior performance in the positive feedback situation.

3.5 STRUCTURAL BREAKS

The alternative switching model used to analyse the Heemeijer et al. (2009) experiment has been introduced only as a proof of the assertion that there are parameter-free learning models that can outperform the rational expectation benchmark. A more sophisticated model 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 is almost identical to the Heemeijer et al. (2009) experiment, excepting the following: Minor changes were that the number of rounds in a run was increased from 50 to 65 and that the variance of the random disturbance ε_t was decreased to $(3/10)^2$.

The major design change was 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 is constant over the experiment, equal to -0.95 in the negative feedback treatment, and equal to 0.95 in the positive feedback treatment.

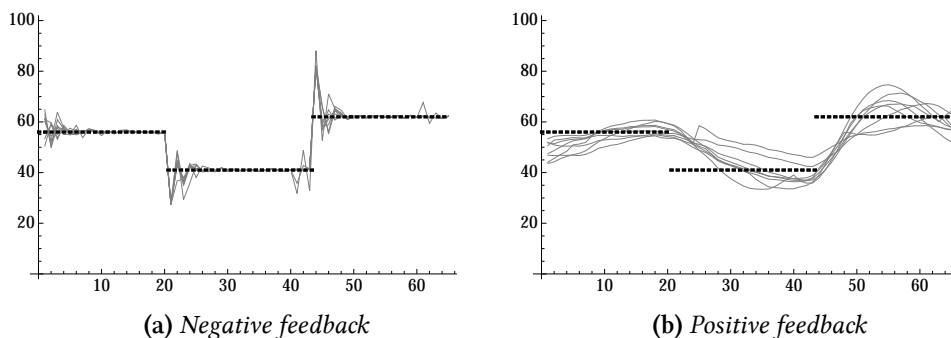


Figure 6: Price realisations in the structural break experiment of Bao et al. (2012). Solid grey lines are realised prices per run; dotted black lines are fundamental prices (cf. Bao et al. (2012), figure 1).

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 that given of the Heemeijer et al. (2009) experiment given above.

4 Conclusion

The rational expectations hypothesis, or rather what has been called here its strong individual incarnation, is not a testable statement about the expectation formation of real economic agents. It is rather a device to be used by economic theoreticians, whether they are outside observers or participating agents, to arrive at a rational benchmark when starting from full knowledge of the system, or at any rate of a model of the system.

Seen in this light, the comparison exercise between rational expectations and the switching model performed above seems to lose much of its point, as the experimental subjects were not equipped with any model of the system they were facing, and they had no possibility of deriving the rational solution to their expectational problem.

Once it is admitted that the strong individual hypothesis is not a statement that can be tested on the basis of observational data, and if moreover it is accepted, as has been made plausible by the evidence presented, that experimental subjects are weakly individually rational, then the experiments reviewed in this article gain a new aspect: they elucidate the relation between institutional arrangements and the possibility that collective predictions of initially uninformed but rational agents 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 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 the most favourable to the strong rational expectations hypothesis, it seems that Walras was right: the rational expectations 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 the most favourable situation, economic theory is moving to a stage where the frictions can be measured and modelled as well.

References

- 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 C. 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 C, 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 C. 2011. The heterogeneous expectations hypothesis: Some evidence from the lab. *Journal of Economic Dynamics and Control* 35(1):1–24
- Hommes C, 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 C, Sonnemans J, Tuinstra J, van de Velden H. 2007. Learning in cobweb experiments. *Macroeconomic Dynamics* 11(S1):8–33
- Hommes C, 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 R, 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
- 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
- Rawls J. 1971. *A Theory of Justice*. Cambridge: Harvard University Press
- Sargent TJ. 1987. Rational expectations. In J Eatwell, M Milgate, P Newman, eds., *The New Palgrave: A Dictionary of Economics*. Palgrave Macmillan
- 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 C, 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 F. 2007. On learning equilibria. *Economic Theory* 30(3):493–513

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