An experimental study on expectations and learning in overlapping generations models

Heemeijer, P.; Hommes, C.H.; Sonnemans, J.H.; Tuinstra, J.

Published in:
Studies in Nonlinear Dynamics and Econometrics

DOI:
10.1515/1558-3708.1944

Studies in Nonlinear Dynamics & Econometrics

Volume 16, *Issue* 4 2012 *Article* 1

An Experimental Study on Expectations and Learning in Overlapping Generations Models

Peter Heemeijer* Cars Hommes†
Joep Sonnemans‡ Jan Tuinstra**

*ABN AMRO Bank, p.heemeijer@uva.nl
†University of Amsterdam, c.h.hommes@uva.nl
‡University of Amsterdam, j.h.sonnemans@uva.nl
**University of Amsterdam, j.tuinstra@uva.nl
DOI: 10.1515/1558-3708.1944

An Experimental Study on Expectations and Learning in Overlapping Generations Models*

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

Abstract

A plethora of models of learning has been developed and studied in macro-economic models in recent years. In this paper we will try to discriminate between these learning models by running laboratory experiments with incentivized human subjects. Participants predict inflation rates for 50 successive periods in a standard overlapping generations model and are rewarded on the basis of their forecasting accuracy. The information set for each participant contains the past inflation rates and the participant's own past predictions which, in turn, determine the actual inflation rate. We consider two treatments, with a low and a high level of monetary growth, respectively. We find that the level of convergence to the monetary steady state is significantly lower and volatility of inflation rates higher in the second treatment. Constant gain learning algorithms, such as adaptive expectations with a low adjustment parameter, seem to provide a better description of the experimental data than decreasing gain algorithms, such as (ordinary) least squares learning. Moreover, many participants switch between prediction strategies during the experiment on the basis of poor performance of their initial prediction strategy.

*We thank the participants to the Workshop on Interacting Agents on Nonlinear Dynamics in Macroeconomics, in Udine, Italy, June 9-11, 2010 and an anonymous referee for helpful comments and suggestions.

1 Introduction

Dynamic macroeconomic models typically correspond to so-called *expectations feedback systems*, where agents' expectations about certain future variables, such as e.g. the inflation rate, are one of the key determinants of the actual realization of these variables. A convenient approach to model expectations in such an expectations feedback system is by assuming that agents have *rational expectations*, that is, that their predictions coincide with those of the relevant economic theory. In recent years, however, it has increasingly been perceived as unsatisfactory that this assumption endows economic agents with exact information about the structure of the economy and the beliefs of other agents as well as unbounded cognitive abilities to process this information. The rational expectations hypothesis might still be valid as a description of long run behavior, however, since economic agents learn over time and thus may eventually arrive at a rational expectations steady state. In the last two decades, therefore, many models of learning in macroeconomics have been introduced (see e.g. Lucas, 1986, and Marcet and Sargent, 1987, for early contributions and Evans and Honkapohja, 2001, for a comprehensive overview). In such a learning model *boundedly rational* agents are generally assumed to have no structural information about their economic environment other than time series observations on certain economic variables. They use these observations to form beliefs about their economic environment and base their decisions on these beliefs.

A wide range of different learning models has been suggested and studied, and different models may lead to outcomes that differ substantially. Bullard (1994), for example, studies a standard overlapping generations model with money growth, where economic agents believe that the inflation rate equals an unknown constant which they estimate by a least squares regression on past prices. This may lead to so-called *learning equilibria*: equilibrium paths that do not converge to the steady state, but that are characterized by endogenous fluctuations in inflation rates and agents beliefs. In the same model Tuinstra and Wagener (2007) show that, by changing the estimation procedure only slightly, the monetary steady state becomes globally stable for almost any specification of the fundamentals of the overlapping generations economy. Other learning models may give other outcomes altogether. Obviously, this "wilderness of bounded rationality" is a serious drawback of departing from rational expectations. Unfortunately, no convincing theoretical arguments or empirical validation that favour one learning model over the others have, thus far, been provided.

In this article we study, in an attempt to cut through this "wilderness of bounded rationality", expectation formation and learning in an overlapping generations model by means of laboratory experiments with incentivized human subjects. The overlapping generations framework that we use is identical to the one studied

Published by De Gruyter, 2012

by Bullard (1994) and Tuinstra and Wagener (2007). The only task of the participants in the experiment is to predict the next inflation rate, for 50 consecutive periods. Predicted inflation rates, in turn, are the main determinant of the realized inflation rate. Participants are paid according to their forecasting accuracy. We consider two treatments, which differ in the (exogenous) monetary growth rate. These treatments are chosen in such a way that different learning models predict different behavior of participants in the two treatments, thus allowing us to use the experimental data to discriminate between these learning models.

This type of 'Learning to Forecast' experiment has previously been done for some other economic environments, such as asset pricing and cobweb models (see e.g. Hommes et al., 2005, 2007, and Heemeijer et al., 2009, for an overview see Hommes, 2011). There are three important differences with respect to these earlier experiments. First, in keeping in line with the vast majority of the literature on learning in macroeconomics, in the experiment discussed in this article the inflation rate is determined by the expectation of a single 'representative' agent. Second, the realized inflation rate for the current period depends upon the expected inflation rate for the current period *as well as* the expected inflation rate for the next period. Finally, the law of motion governing the relationship between the expected and realized inflation rates is highly nonlinear. Examples of earlier OLG experiments are Marimon and Sunder (1993, 1994) and Marimon et al. (1993). More recent experiments on learning in macroeconomic environments can be found in Adam (2007), Pfajfar and Žakelj (2010) and Assenza et al. (2011).

We find that convergence to the monetary steady state occurs much less frequently in the treatment with the high monetary growth rate. This suggests that constant gain learning algorithms, such as adaptive expectations with a low adjustment parameter, provide a better description of the experimental data than decreasing gain algorithms, such as (ordinary) least squares learning. A second important result is that participants tend to switch between prediction strategies during the experiment, in particular when their initial prediction strategy provides poor forecasts.

The remainder of this article is structured as follows. Section 2 introduces the overlapping generations model that is used in the experiment and analyzes different types of expectation formation and learning in that model. In Section 3 we discuss the design of the experiment. The experimental results are presented and analyzed in Section 4. Finally, Section 5 summarizes the results and adds some concluding remarks. Details of the experimental procedure and the instructions for the participants (including an English translation) can be found in the appendix.

2 Expectations and Learning in an Overlapping Generations Model

2.1 The overlapping generations economy

We consider a standard overlapping generations exchange economy with two generations, a single consumption good and a government creating money at a constant rate.[1] Each generation consists of a representative household, whose preferences are represented by a utility function $u(c_0, c_1)$, where c_0 and c_1 denote consumption of the good when the household is "young" and "old", respectively. The representative household is endowed with w_0 (w_1) units of the consumption good when young (old).

Let p_t be the price of the good in the current period t and p^e_{t+1} the expectation of the price in period $t+1$. The young generation in period t maximizes its lifetime utility $u(c_0, c_1)$ given the perceived lifetime budget constraint $p_t c_0 + p^e_{t+1} c_1 \leq p_t w_0 + p^e_{t+1} w_1$. Assuming that $u(c_0, c_1)$ is continuous, strictly increasing and strictly quasi-concave a unique solution to the utility maximization problem, $(c^*_0, c^*_1) = \left(c_0\left(\frac{p^e_{t+1}}{p_t}\right), c_1\left(\frac{p^e_{t+1}}{p_t}\right)\right) \geq 0$, exists.[2] The savings function $S(\cdot)$ is defined as

$$S\left(\frac{p^e_{t+1}}{p_t}\right) = w_0 - c_0\left(\frac{p^e_{t+1}}{p_t}\right),$$

which we will assume to be positive and downward-sloping, $S(\cdot) > 0$ and $S'(\cdot) < 0$, for all values of $\frac{p^e_{t+1}}{p_t}$.

In equilibrium aggregate savings equal real money balances

$$\frac{M_t}{p_t} = S\left(\frac{p^e_{t+1}}{p_t}\right). \quad (1)$$

The total money stock M_t is exogenously determined by the money creation policy of the government, which is given by

$$M_t = \vartheta M_{t-1}, \quad \vartheta > 1, \quad (2)$$

where ϑ is the rate of growth of the money stock, used by the government to finance

[1] The general structure of the OLG model is identical to that of the model studied by Bullard (1994), Schönhofer (1999) and Tuinstra and Wagener (2007).

[2] Note that in period $t+1$ the old generation will use all of its income to buy consumption goods against the market price. Actual consumption of the old generation is therefore determined by the realized price in period $t+1$ and hence will typically differ from planned consumption $c_1\left(\frac{p^e_{t+1}}{p_t}\right)$ when $p_{t+1} \neq p^e_{t+1}$.

Published by De Gruyter, 2012

expenditures. Solving (1) for M_t, substituting in both sides of (2) and rearranging gives the *actual law of motion*:

$$\pi_t = \vartheta \frac{S\left(\pi_t^e\right)}{S\left(\pi_{t+1}^e\right)}, \tag{3}$$

where $\pi_t = \frac{p_{t+1}}{p_t}$ is the gross inflation rate. Note that the *monetary steady state* is given by $\pi_t^e = \pi_t = \vartheta$, for all t. At this monetary steady state inflation is constant and fully determined by the government's monetary policy.[3]

The model needs to be closed by specifying how agents form expectations. Under *rational expectations* or *perfect foresight*, agents correctly anticipate future inflation rates, that is, $\pi_t^e = \pi_t$ for all t. Given our assumptions on the savings function the only feasible equilibrium trajectory satisfies $\pi_t = \vartheta$ for all t, every other trajectory diverges to infinity or becomes infeasible in finite time.[4] In recent years it has been recognized that the rational expectations hypothesis may be too demanding. It requires that agents have full knowledge of the underlying structure of the model and can use this to exactly determine the future development of the inflation rates. In the next subsection we consider some alternative models of expectation formation and learning that have been considered in dynamic economic models in recent years.

2.2 Expectations and learning

Probably the simplest expectation rule is the *naive expectations* rule, where agents expect the next inflation rate to be equal to the last observed inflation rate,[5] that is

$$\pi_{t+1}^e = \pi_{t-1}. \tag{4}$$

[3]Observe that the *autarkic steady state equilibrium*, where $\pi_t^e = \pi^a$ for all t, with π^a such that $S\left(\pi^a\right) = 0$, does not exist in our model since we require $S\left(\pi\right)$ to be positive for all $\pi \geq 0$.

[4]See e.g. Proposition 1 from Tuinstra (2003). Note, however, that the assumptions on the savings function are quite stringent. If an autarkic steady state exists many equilibrium trajectories will converge to that steady state. Moreover, if income effects are sufficiently strong the savings function will be non-monotonic which may generate rational expectations equilibria with a complicated cyclical character, see e.g. Grandmont (1985). One implication for our model is that any volatility or cyclical behavior is due to deviations from full rationality.

[5]Note that, since π_t depends upon both π_t^e and π_{t+1}^e(see equation (3)), the agent cannot observe π_t when having to forecast π_{t+1}^e. Agents therefore have to predict two periods ahead, and the last observed inflation rate when predicting the inflation rate for period $t+1$ is indeed π_{t-1}.

Inserting (4) into (3) shows that the evolution of inflation rates is now governed by the following nonlinear second order difference equation

$$\pi_t = \vartheta \frac{S(\pi_{t-1})}{S(\pi_t)}. \tag{5}$$

Stability of the monetary steady state ϑ under naive expectations turns out to depend crucially upon the sensitivity of the savings function with respect to changes in the inflation rate, which can be measured by the *inflation elasticity of savings*

$$a(\pi) = -\frac{\pi}{S(\pi)} \frac{\partial S(\pi)}{\partial \pi}. \tag{6}$$

Since we assumed the savings function to be positive and decreasing, we have $a(\pi) \geq 0$ for all $\pi \geq 0$.

The monetary steady state ϑ is a locally stable steady state of (5) if and only if $a(\vartheta) < 1$. For $a(\vartheta) = 1$ the dynamics undergoes a *Neimark-Sacker* bifurcation, in which an invariant circle is created, along which inflation rates fluctuate perpetually. Similar dynamic behavior may emerge for other types of expectation schemes. For example, the monetary steady state ϑ is locally stable under *adaptive expectations*

$$\pi_{t+1}^e = \pi_t^e + \alpha \left(\pi_{t-1} - \pi_t^e\right), \quad 0 < \alpha \leq 1, \tag{7}$$

if and only if $\alpha a(\vartheta) < 1$.[6]

A third type of expectation rule predicts the inflation rate by taking the average of the last k observed inflation rates

$$\pi_{t+1}^e = \frac{1}{k} \sum_{s=1}^{k} \pi_{t-s}. \tag{8}$$

We will denote expectation rule (8) by *average expectations*.[7] The monetary steady state ϑ is locally stable under average expectations if and only if $a(\vartheta) < a_k^*$, where k is the number of lags used. Straightforward computations show that $a_2^* = \sqrt{2}$, $a_3^* \approx 1.854$, $a_4^* \approx 2.309$, and so on (obviously $a_1^* = 1$).

The three types of expectation rules discussed above are applied mechanistically and are not adapted to the particular economic environment in which they

[6] See e.g. Proposition 2 from Tuinstra (2003). Moreover, the dynamics under naive or adaptive expectations may exhibit erratic and endogenous fluctuations when the monetary steady state is unstable.

[7] Note that both adaptive expectations and average expectations are generalizations of naive expectations (the latter corresponding to $\alpha = 1$ and $k = 1$, respectively).

are used. Consequently, there will typically be some autocorrelation structure in the resulting time series of forecast errors which could be exploited by more sophisticated expectation mechanisms. To deal with this, models of *learning* assume agents adapt their beliefs on the basis of the forecast errors, and in fact behave more like an econometrician would.[8] A substantial part of this learning literature indeed assumes that agents try to learn about the underlying economy by running least squares regressions (for an overview, see Evans and Honkapohja, 2001).

Bullard (1994), for example, assumes that agents in the overlapping generations economy introduced in the previous subsection have the following *perceived law of motion*

$$p_t = \beta p_{t-1}. \tag{9}$$

Here β is an unknown constant which is estimated every period on the basis of a least squares regression on past prices. The least square estimate for β, which is equivalent to the expected inflation rate, is given by

$$\beta_t = \frac{\sum_{s=1}^{t-1} p_{s-1} p_s}{\sum_{s=1}^{t-1} p_{s-1}^2}.$$

This estimate can be written recursively as

$$\beta_{t+1} = \beta_t + g_t \left(\pi_{t-1} - \beta_t\right), \tag{10}$$

where $g_t = p_{t-1}^2 \left[\sum_{s=1}^{t} p_{s-1}^2\right]^{-1}$ is a weight factor. Note the similarity between (10) and the adaptive expectations rule (7). In both cases the new estimate is a correction of the old estimate in the direction of the last observed inflation rate. The difference between the two lies in the correction factors α and g_t, which is constant in one case, and depends upon the data in the other.[9]

The monetary steady state ϑ of the overlapping generations economy is locally stable under least squares learning on prices (10) if and only if (see Bullard, 1994)

$$\left(1 - \vartheta^{-2}\right) a\left(\vartheta\right) < 1. \tag{11}$$

At first sight the potential nonconvergence of least squares estimates is surprising, since the agents weigh every price observation equally. Therefore one would expect

[8] Sargent (1993, p.22), for example, writes that “We can interpret the idea of bounded rationality broadly as a research program to build models populated by agents who behave like working economists or econometricians”.

[9] Adaptive expectations of the type (7) are therefore sometimes referred to as *constant gain* algorithms (or *perpetual learning*), since the gain of a new observation is always the same, independent of the length of time that has passed (see e.g. Cho et al., 2002, Orphanides and Williams, 2005).

the estimate β_t to converge eventually. However, estimates may fail to convergence due to the nonstationary character of prices. In fact, for $\vartheta > 1$, prices will grow exponentially. As a consequence, the steady state value of g_t, $g^* = 1 - \vartheta^{-2}$, is strictly positive, implying that the estimate β_t may keep on fluctuating.[10]

Tuinstra and Wagener (2007) argue that the endogenous fluctuations in the least squares learning model introduced in Bullard (1994), so-called *learning equilibria*, are therefore due to the estimation procedure. They suggest the perceived law of motion $\pi_t = \beta$ which, like (9), represents the belief that the inflation rate is constant. A least squares regression on the stationary series of past inflation rates will then give as an estimate

$$\beta_{t+1} = \sum_{s=0}^{t-1} \pi_{t-s} = \beta_t + \frac{1}{t}\left(\pi_{t-1} - \beta_t\right). \tag{12}$$

This least squares learning model converges for almost all savings functions.[11] Note that (12) is an example of a *decreasing gain* algorithm, where the gain of new observations, represented by the weight $\frac{1}{t}$, decreases as time passes.

The discussion in this subsection shows that, depending upon the way in which expectations and learning is modelled, many different types of dynamic behavior are possible in our overlapping generations economy.[12]

2.3 Numerical simulation of expectation and learning rules

We will now present numerical simulations with some of the expectation and learning rules reviewed in the previous subsection, for a specified version of the model. Consider the following savings function[13]

$$S(\pi) = \delta + (1-\delta)\frac{w_0}{1 + (\nu\pi)^{\frac{\rho}{1-\rho}}}, \tag{13}$$

[10]In fact, local stability conditions for (7) and (10) are exactly the same for $\alpha = g^*$.

[11]See Tuinstra and Wagener (2007) for a discussion. Wenzelburger (2002) presents another learning algorithm that converges for almost all savings functions.

[12]Many other learning models are possible. Tuinstra (2003), for example, considers perceived laws of motion of the form $\pi_t = \beta\pi_{t-1}$ and $\pi_t = \alpha + \beta\pi_{t-1}$, where the belief parameters α and β are estimated by a least squares regression. These belief parameters converge, but the corresponding inflation rates may still exhibit endogenous fluctuations.

[13]It is worth noting that for $\delta = 0$ (13) corresponds to the savings function derived from the well-known CES utility function $u(c_0, c_1) = ((\nu c_0)^\rho + c_1^\rho)^{\frac{1}{\rho}}$, with endowments given by $w_0 > 0$ and $w_1 = 0$. For $\delta > 0$ the savings function (13) corresponds to an economy where a proportion δ of the young generation receives an endowment of $w_0 = 1$ and saves this completely, while the remaining proportion $1 - \delta$ has preferences represented by the above CES utility function. The advantage of

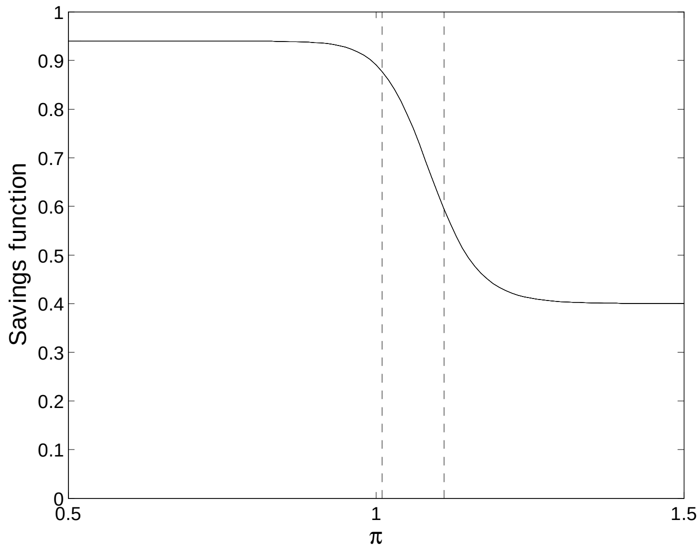

Figure 1: The savings function (13) as a function of π. Parameter values are $\delta = 0.4$, $\rho = 0.965$, $\nu = 0.92$ and $w_0 = 0.9$. The two vertical lines at $\pi = 1.01$ and $\pi = 1.11$ correspond to the two benchmark steady state inflation rates ϑ_1 and ϑ_2, respectively.

with $\delta \in [0,1]$, $\nu > 0$ and $0 \neq \rho < 1$. Note that $S(\pi) > 0$ and $S'(\pi) < 0$ for all values of π, provided $\rho > 0$. We choose $w_0 = 0.9$, $\nu = 0.92$, $\rho = 0.965$ and $\delta = 0.4$, which implies that $\frac{20}{47} < \frac{\pi_t}{\vartheta} < \frac{47}{20}$. The savings function (13) for these parameter values is shown in Figure 1. As benchmark monetary growth rules we choose $\vartheta_1 = 1.01$ and $\vartheta_2 = 1.11$, which are shown by the vertical lines in Figure 1. Note that, because of the high value of ρ, savings decrease very quickly for inflation rates a little above 1.

As we have seen above, the dynamics under different expectations schemes depends critically upon the inflation elasticity of savings at the monetary steady state, $a(\vartheta)$. Figure 2 shows $a(\vartheta)$ for $1 \leq \vartheta \leq 1.2$, together with critical values for $a(\vartheta)$ under least squares learning on prices (the higher dashed curve) and under average expectations (8) for $k = 1$, 2, 3 and 4, respectively (the horizontal dashed lines).[14] For our two benchmark economies, represented by the vertical lines in

(13) is that the inflation rate will always lie in an open interval, characterised by

$$\frac{\delta}{\delta + (1-\delta) w_0} < \frac{\pi_t}{\vartheta} < \frac{\delta + (1-\delta) w_0}{\delta}.$$

[14]Substituting (13) and its derivative in the definition of $a(\pi)$ gives the explicit form of the elas-

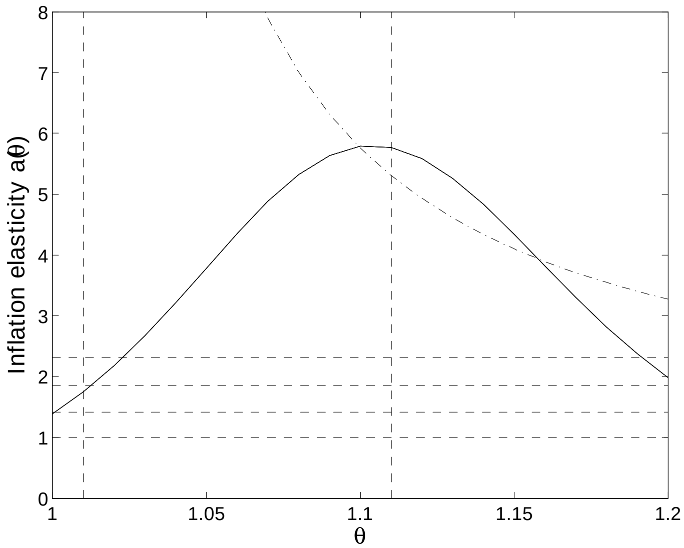

Figure 2: The inflation elasticity of saving (6) as a function of the steady state interest rate ϑ. Parameter values are $\delta = 0.4$, $\rho = 0.965$, $\nu = 0.92$ and $w_0 = 0.9$. The dashed horizontal lines correspond to the critical values of $a(\vartheta)$ for local stability of the steady state under average expectations (8) for (from low to high) $k = 1$, $k = 2$, $k = 3$ and $k = 4$, respectively. The higer dashed curve indicates the critical value of $a(\vartheta)$ for stability of the least squares learning model on prices (see (11)). The two vertical lines at $\vartheta_1 = 1.01$ and $\vartheta_2 = 1.11$ indicate the two benchmark steady state inflation rates.

Figure 2 we have $a(\vartheta_1) \approx 1.75$ and $a(\vartheta_2) \approx 5.77$. From Figure 2 and the discussion in the previous subsection we can immediately infer the following. First, when $\vartheta = \vartheta_1$, the monetary steady state is locally stable under adaptive expectations (7) with $\alpha < [a(\vartheta_1)]^{-1} \approx 0.57$ and under average expectations (8), provided $k \geq 3$. Furthermore, this steady state is also stable under least squares learning on prices (10) or inflation rates (12). Local stability properties of the monetary steady state are quite different when $\vartheta = \vartheta_2$, however. The steady state is then locally stable under adaptive expectations (7) only if $\alpha < [a(\vartheta_2)]^{-1} \approx 0.17$ and under average expectations (8) only if $k \geq 12$. Moreover, in this case least squares learning on

ticity curve as:

$$a(\pi) = \frac{\rho}{1-\rho} \frac{w_0 (\nu\pi)^{\frac{\rho}{1-\rho}}}{\left(\frac{\delta}{1-\delta} + \frac{w_0}{1+(\nu\pi)^{\frac{\rho}{1-\rho}}}\right)\left(1 + (\nu\pi)^{\frac{\rho}{1-\rho}}\right)^2},$$

part of which is shown in Figure 2 after subsituting the parameter values.

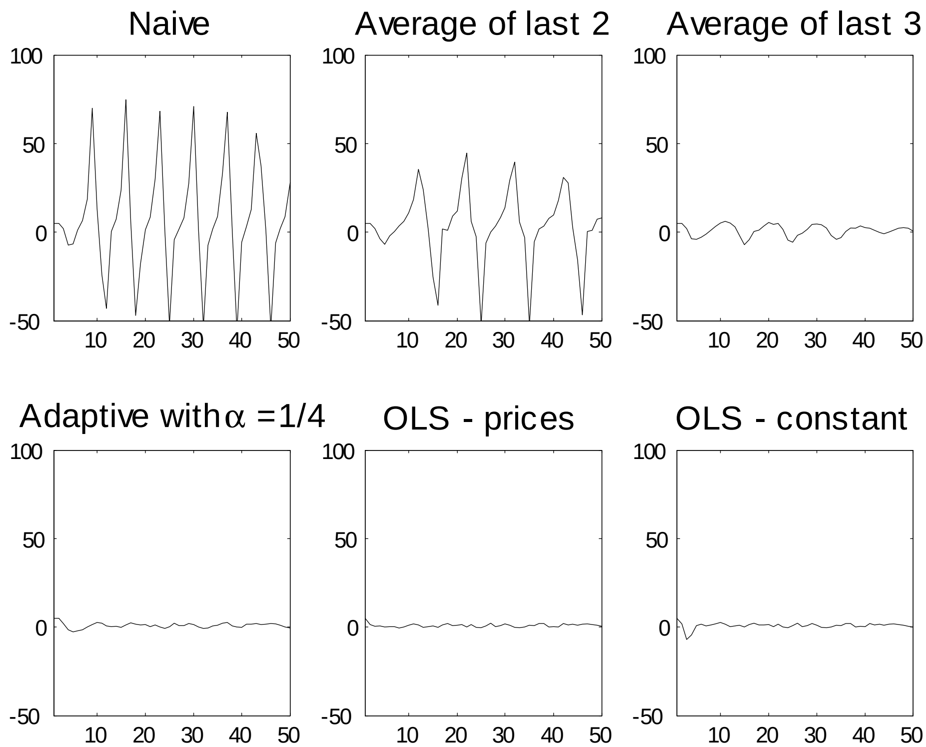

Figure 3: Simulations of the overlapping generations model (14) with $\vartheta = 1.01$ for different expectation and learning models. Top row: naive expectations (4), average expectations (8) with $k = 2$ and $k = 3$ respectively. Bottom row: adaptive expectations (7) with $\alpha = \frac{1}{4}$, learning on prices (10), and learning on inflation rates (12). Inflation rates are represented in percentages.

prices (10) is unstable since $a(\vartheta_2) > \left(1 - \vartheta_2^{-2}\right)^{-1} \approx 5.31$.

Figures 3 and 4 show simulations of the two benchmark cases ϑ_1 and ϑ_2, with six different expectation rules: averaging expectations (8) with $k = 1$ (i.e. naive expectations), $k = 2$ and $k = 3$, adaptive expectations (7) with $\alpha = \frac{1}{4}$, and least squares learning on prices (10) and inflation rates (12), respectively. For these simulations (and for the laboratory experiment that we will discuss in the next sections) a slightly different law of motion was used, namely

$$\pi_t = \vartheta \varepsilon_t \frac{S\left(\pi_t^e\right)}{S\left(\pi_{t+1}^e\right)}, \tag{14}$$

with ε_t uniformly distributed on the interval $[0.975, 1.025]$. This implies that in the constant rational expectations equilibrium $\pi_t^e = \vartheta$ holds for all t, while the actual

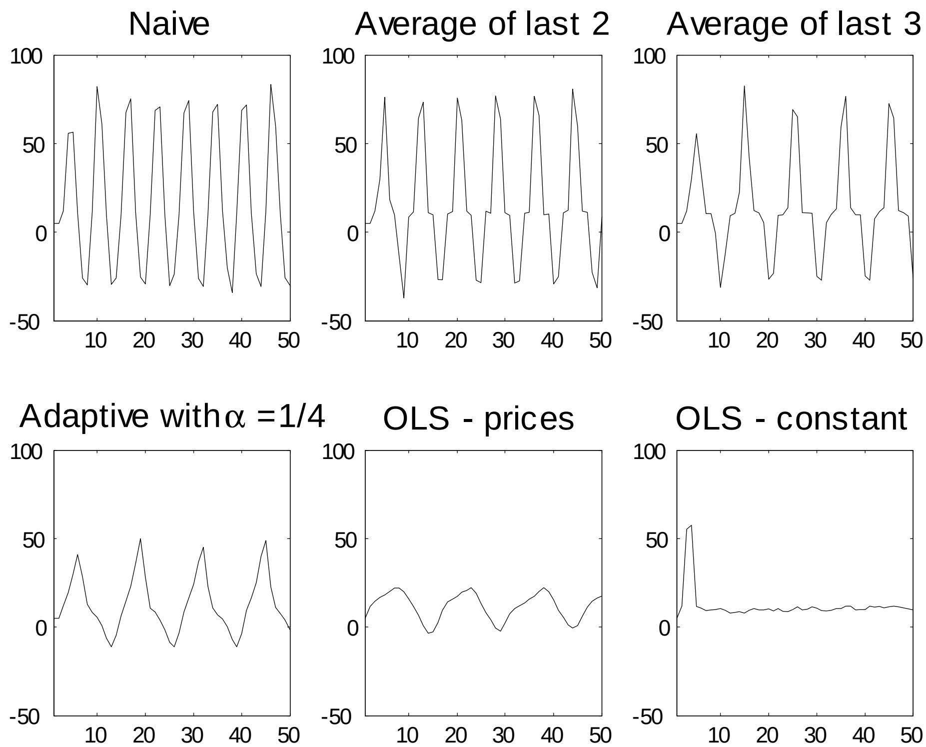

Figure 4: Simulations of the overlapping generations model (14) with $\vartheta = 1.11$ for different expectation and learning models. Top row: naive expectations (4), average expectations (8) with $k = 2$ and $k = 3$ respectively. Bottom row: adaptive expectations (7) with $\alpha = \frac{1}{4}$, learning on prices (10), and learning on inflation rates (12). Inflation rates are represented in percentages.

inflation rate fluctuates around the monetary steady state with an amplitude of at most 2.5 percent of that steady state.[15]

For both benchmark cases *myopic* rules that rely heavily on very recent information on the inflation rate, such as naive expectations, average expectations with $k = 2$, or adaptive expectations with an adjustment parameter α that is relatively close to 1, lead to volatile and persistent fluctuations in interest rates. On the other hand, the *decreasing gain* algorithm (12) stabilizes the dynamics for both values of the money growth rate ϑ. The difference between the two benchmark

[15] The random disturbances represent fluctuations in the expansion of the money supply by the central bank. In the experiment these disturbances prevent trivial developments in the realized inflation rates. The same realization of shocks ε_t is used in each simulation. This realization of shocks is also used in the experimental treatments discussed in the following sections. As initial conditions for the simulations we used $\pi_1 = \pi_2 = 1.05$ for both benchmarks.

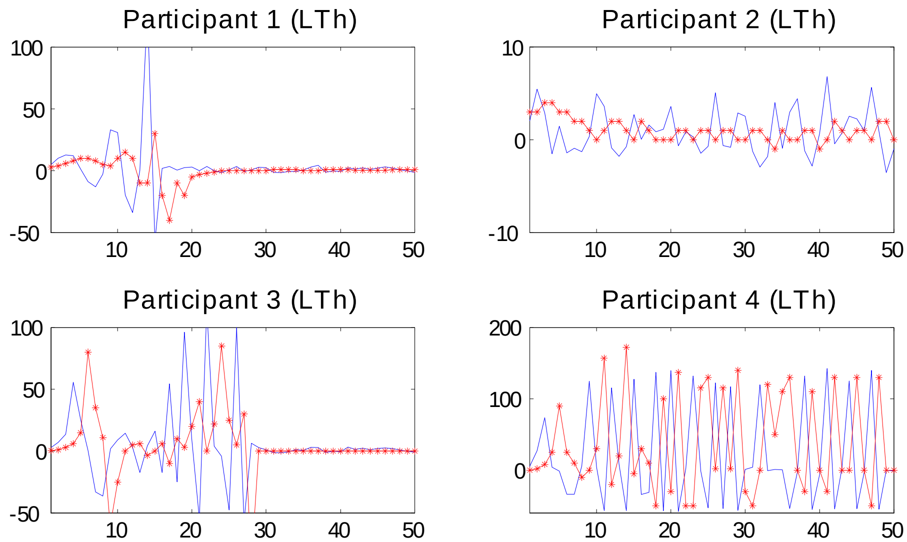

Figure 5: Predictions (red crosses) and inflation rates (blue line) for participants $1-4$ in treatment LTh. Note that the range on the vertical axis differs across participants.

cases occurs with expectation rules that mainly base their prediction on relatively recent information, for example, adaptive expectations with a moderate value of α, average expectations using at least $k=3$ lags, or least squares learning on prices (10). Under these rules the monetary steady state is locally stable for $\vartheta=\vartheta_1$, but unstable for $\vartheta=\vartheta_2$.

3 Experimental Design

The experiment was conducted on April 25, 2006 in the CREED experimental laboratory of the University of Amsterdam in two sessions of about 90 minutes each. Participants had to predict future inflation rates. The resulting realized inflation rate is determined by (14), with the savings function given by (13), for parameter values equal to $\delta=0.4$, $\nu=0.92$, $\rho=0.965$ and $w_0=0.9$. The two treatments correspond to the two benchmark cases studied in the previous section, $\vartheta_1=1.01$, denoted *Low Theta* (LTh), and $\vartheta_2=1.11$, denoted *High Theta* (HTh), respectively. Both treatments are single-agent treatments, i.e. each participant provides the inflation estimates necessary to determine the realized inflation in (14), generating a series of inflation rates that is not influenced by the actions of other participants.[16] Both treatments consist of 16 participants, each generating a series of predictions and associated realized inflation rates.[17]

In the experiment, participants submit inflation predictions that result in actual inflation rates. As soon as these inflation rates are realized, they are made available to participants for possible use in the formation of further predictions. Since both inflation predictions π_t^e and π_{t+1}^e are required to determine the actual inflation rate π_t, a participant predicts the inflation rate two periods ahead. The experiment lasts a total of 50 periods, meaning that each participant submitted 51 predictions producing a series of 50 realized inflation rates, with the last prediction only serving to determine the 50^{th} inflation. In an arbitrary period t, during which a participant i is asked to submit an inflation prediction π_{t+1}^e, the information set available to him is $I_{i,t} = \left\{ \left\{ \pi_{i,1}^e, \pi_{i,2}^e, \ldots, \pi_{i,t}^e \right\}, \left\{ \pi_1, \pi_2, \ldots, \pi_{t-1} \right\} \right\}$. This information is provided to the participant in a table as well as in a graph.

Subjects are rewarded for the accuracy of their predictions. The number of points subject i earns for his/her prediction $\pi_{i,t}^e$ for period t, given that the realized inflation rate for that period is π_t, is given by

$$P_i \left(\pi_{i,t}^e, \pi_t \right) = \max \left\{ 100 - 400 \left| \pi_{i,t}^e - \pi_t \right|, 0 \right\}. \tag{15}$$

Participants therefore earn a reward for each prediction with an error of no more than 0.25. For each participant all rewards per period are rounded and then added at the end of the experiment to be converted into euros at the rate of 200 points for 1 euro. Subjects could therefore earn a maximum of 25 euros for the whole experiment.[18]

[16]This experimental set-up is analogous to the use of a single representative agent in most literature on expectation formation and learning in macroeconomic models.

[17]Ideally, one would like to investigate how the same individual behaves in the two different treatments in order to better understand the way in which this individual forms expectations. However, in that case experience of this individual with one treatment would influence his behavior in the other treatment. No subject therefore participates in both treatments.

[18]In earlier 'Learning to Forecast' experiments (e.g. Hommes et al., 2005, 2007 and Heemeijer et al., 2009) participants earnings are determined by quadratic instead of absolute forecast errors. The drawback of using quadratic forecast errors is that small forecast errors are hardly penalized because the quadratic function is rather flat close to its minimum. On the other hand, Sonnemans and Tuinstra (2010) show that outcomes of 'Learning to Forecast' experiments are relatively robust to changes in the incentive structure.

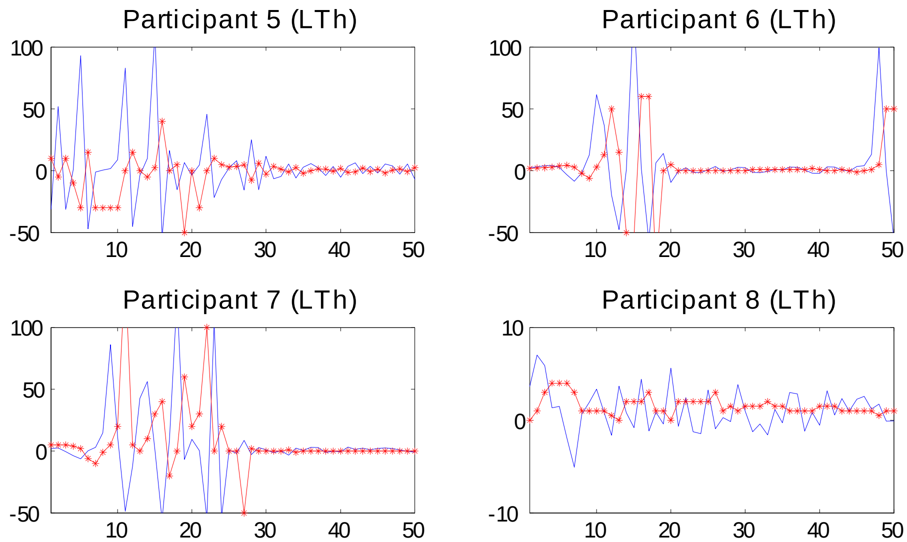

Figure 6: Predictions (red crosses) and inflation rates (blue line) for participants $5-8$ in treatment LTh.

4 Experimental Results

The experiment is designed in such a way that it allows us to discriminate between the different types of learning behavior participants may exhibit. In particular, if participants use myopic short memory learning procedures, inflation rates will be unstable in both treatments, as is illustrated in the first two panels of Figures 3 and 4. On the other hand, if participants use decreasing gain learning algorithms, inflation rates will be stable in both treatments (see e.g. the last panel of Figures 3 and 4). Finally, if participants use constant gain learning algorithms, for example adaptive expectations with a relatively low adaptation parameter, typically inflation rates would be stable in treatment LTh but unstable in treatment HTh, as is illustrated by the remaining panels of Figures 3 and 4.

In the following subsections we will present and analyze the experimental results to see whether we can infer which type of expectation rules participants use.

4.1 A classification of the results

Since the experiment has only single-agent treatments, results for each of the 32 participants can be studied separately. The experimental results of each participant are given by a series of inflation predictions and realized inflation values, of length

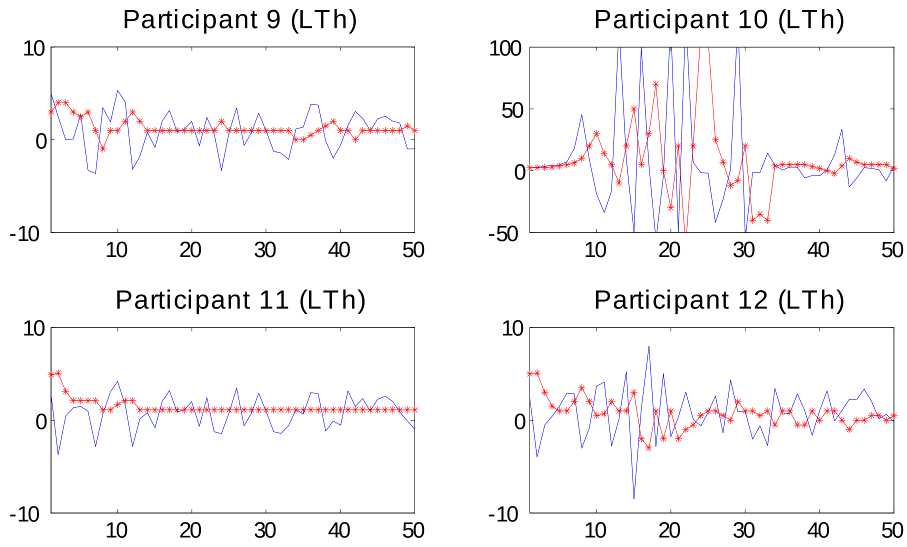

Figure 7: Predictions (red crosses) and inflation rates (blue line) for participants $9 - 12$ in treatment LTh.

51 and 50 respectively. These series are shown for the 16 participants of treatment LTh in Figures 5–8 and for the 16 participants of treatment HTh in Figures 9–12. Note that the range on the vertical axis may differ across participants in these figures. Although this may complicate direct comparison between some participants, it does facilitate a better understanding of the forecasting strategies that individual participants use.

Let us start out with investigating the participants of treatment LTh. The question is whether the inflation rate tends to converge to the monetary steady state of 1%. The results in Figures 5–8 are helpful in answering this question: there is a clear separation between participants who manage to stabilize inflation around the monetary steady state and those who do not.[19]

The experimental results of the participants fall roughly into one of three qualitative categories that we denote "stable" (category *S*), "unstable to stable" (*US*)

[19]The term "convergence" is used in this subsection in the sense of approaching relatively closely the monetary steady state, taking into account the exogenous disturbance term and the fact that inflation fluctuations can easily have an amplitude of 40 (see e.g. participants 4, 10 and 13 from treatment LTh). Note that the chosen savings function imposes that the inflation rate is always between -58.1 and 143.3 for treatment LTh, and between -54.0 and 167.4 for treatment HTh, see the discussion following equation (13). In Subsection 4.2 a statistical test of convergence is performed on the inflation series of both treatments to determine unambiguously whether convergence is achieved or not.

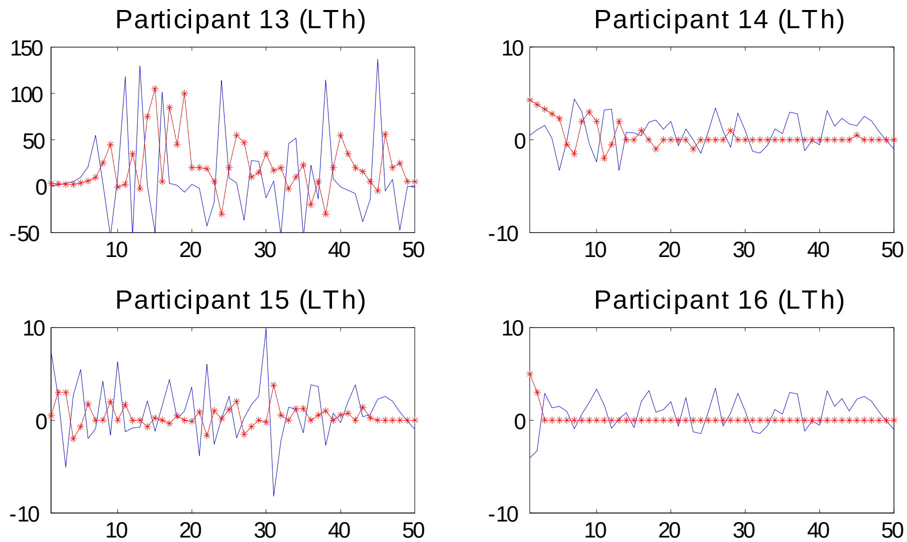

Figure 8: Predictions (red crosses) and inflation rates (blue line) for participants $13-16$ in treatment LTh.

and "unstable" (U). The first category consists of subjects that succeed from the beginning of the experiment in keeping the inflation rate close to the monetary steady state. From Figures 5–8 we see that participants 2, 8, 9, 11, 12, 14, 15 and 16 from treatment LTh can be put in this category. This corresponds to exactly half the number of subjects in this treatment. The second category contains participants who manage to stabilize inflations around the monetary steady state, but only after going through an initial phase of substantial instability. Participants 1, 3, 5, 6 and 7 from treatment LTh fit this description, representing almost a third of the first treatment.[20] The third and last category contains subjects unable to stop inflation from fluctuating wildly. In treatment LTh only participants 4, 10 and 13 display this sort of movement, making permanent instability a relatively uncommon phenomenon.[21]

The difference between treatments LTh and HTh lies in the money growth rate ϑ which, next to fixing the steady state level of inflation, influences the magnitude of the random disturbances in each period (see equation (14)). However,

[20]Participant 6 is somewhat special: inflation rates do nearly converge from period 21 onwards, although this convergence breaks down in the last couple of periods.

[21]Note that participant 10 corresponds to a borderline case. An initial phase of high volatility is followed up around period 30 by moderately successful attempts to control the inflation. These do not result in convergence though within the length of the experiment, making it difficult to classify the last 20 periods as either converging or unstable.

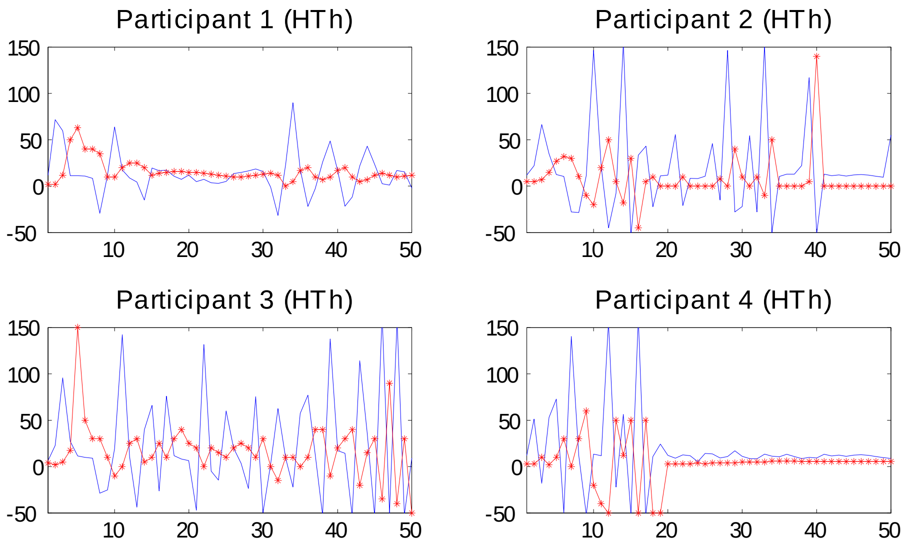

Figure 9: Predictions (red crosses) and inflation rates (blue line) for participants $1-4$ in treatment HTh.

increasing ϑ from 1% to 11% has an effect on the evolution of inflation rates that goes far beyond the shift in the steady state equilibrium and the increase in the amplitude of random disturbances. Considering the 16 time series in Figures 9–12, corresponding with each of the participants in treatment HTh, it is clear that the volatility of inflation on the whole increases dramatically. This increase in volatility with respect to treatment LTh clearly reflects a diminished ability of participants to stabilize the inflation dynamics.

Classifying subjects' predictions in the same three categories as before yields a very different distribution. First, none of the participants from treatment HTh manages to keep the inflation rate at or around the monetary steady state rate from the beginning of the experiment, clearly indicating that the higher equilibrium rate of inflation made it harder for subjects to stabilize the system (recall that half of the subjects in treatment LTh were able to do this). Furthermore, five participants in treatment HTh (participants 4, 5, 6, 10 and 11) can be classified in category *US*, whereas the remaining 11 participants experience large volatility in inflation rates until the end of the experiment.[22]

[22]There are some borderline cases again. Participant 2 for example, classified in category *U*, is somewhat successful in stabilizing the inflation rate by consistently predicting an inflation rate of 0 between periods 41 and 50. Similarly, participants 7, 9 and 15 seem to be mildly successful in stabilizing the inflation dynamics towards the end of the experiment. On the other hand, participants

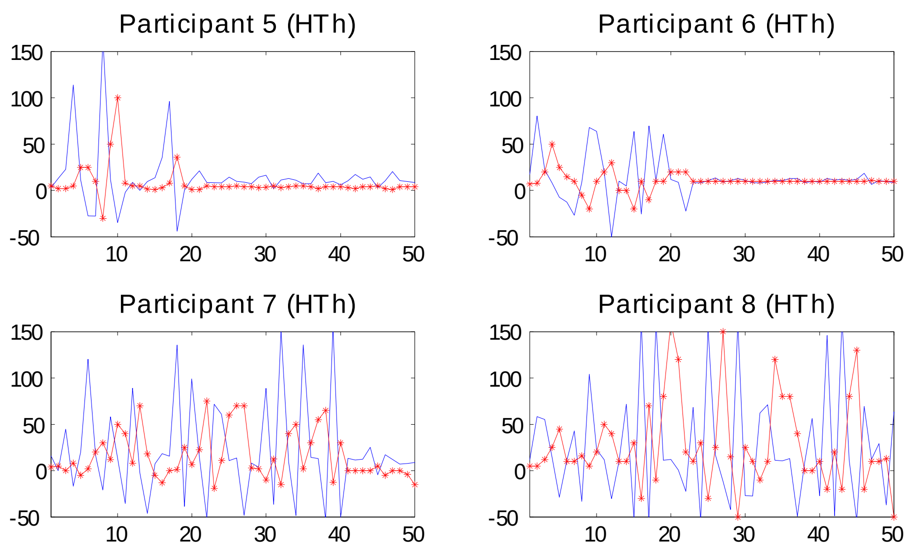

Figure 10: Predictions (red crosses) and inflation rates (blue line) for participants $5-8$ in treatment HTh.

Category	LTh	participants	HTh	participants
S	**8**	$2,8,9,11,12,14-16$	–	–
US	**5**	$1,3,5-7$	**5**	$4-6,10,11$
U	**3**	$4,10,13$	**11**	$1-3,7-9,12-16$

Table 1: Classification of participants in different categories.

Table 1 summarizes the classification of participants in the different categories. From the table we see that the number of participants having a uniformly high inflation volatility throughout the experiment increases from three to eleven when moving from the first to the second treatment. Moreover, the number of participants able to stabilize the inflation rate from the beginning of the experiment decreases from eight in the first treatment to zero in the second treatment.[23] This

4 and 5, who are classified in category *US*, stabilize inflation rates by consistently submitting predictions around 4 for the last 30 periods. One could also argue that these participants belong to category *U*, since their predictions are not converging to a small neighborhood of the monetary steady state inflation rate of 11%. Classifying them in category *U* would make the difference between the two treatments even more pronounced.

[23]Moreover, average earnings in the second treatment (8.16 euros) are distinctly below those of

suggests that indeed there is a strong effect of the treatment on stability.[24] From this we conclude that neither myopic expectations, nor decreasing gain algorithms provide a convincing general explanation of the experimental results.

4.2 Learning the monetary steady state

We have seen that predictions of some of the participants, in particular from treatment LTh, converge to the monetary steady state. This suggests that the Rational Expectations Hypothesis (REH) might provide a reasonable description of the prediction behavior for these participants, at least in the long run. In this subsection we formally test this conjecture. Note that at the unique rational expectations (RE) equilibrium the inflation rate fluctuates around the monetary steady state with low amplitude such that the individual only makes prediction errors due to the random shocks in the economy. None of the subjects in the experiment acted exactly in accordance with this equilibrium; first because the instructions did not give them sufficient information to derive it theoretically and second because no market information was available during the first two periods, making it virtually impossible to guess the right inflation rate from the start.

Figure 13 displays the median absolute forecast, i.e. the averages between the eighth and ninth ranked errors, for all periods and both treatments. The dashed line at 40 draws the threshold above which predictions errors are so large that earnings for that period are zero (see equation (15)). In treatment LTh the accuracy of predictions tends to increase during the experiment, as is shown by the decreasing trend in the median errors in Figure 13. By the end of the treatment a majority of all participants satisfy the REH (in the last four periods at least 10 subjects do so; in the last 10 periods the average number of REH-compatible errors is 9.7)[25] and

the first (17.19 euros), as is verified by a one-sided t test, giving $H_0 : \mu_1 = \mu_2$, $t \approx 4.159$ (30 df.), $p < 0.001$. The greater difficulty subjects had in stabilizing inflation development under the high money growth rate ϑ_2 is therefore also reflected in the significantly lower average earnings in the second treatment. Within treatments also major differences exist, with the top-ranked participants earning more than five times as much as the bottom-ranked one. We will get back to this in Subsection 4.3.

[24]What might also contribute to the treatment effect on stability is that inflation rates are more likely to start out in the vicinity of the monetary steady state in treatment LTh than in treatment HTh, because participants tend to submit initial inflation predictions that are positive but not too large (the median predictions for the first two inflation rates are 3% and 2% in treatment LTh and 3% and 3% in treatment HTh). This may partially explain the higher number of stable inflation series in treatment LTh.

[25]Note that, from (14) and the distribution of ε_t, it follows that in the unique RE steady state the distribution of inflation rates is uniform on the domain $[0.975\vartheta, 1.025\vartheta]$. A participant using rational expectations would have absolute errors uniformly distributed on $[0, 0.025\vartheta]$. The latter distribution easily translates into a test of the RE equilibrium in any single period, since, at a significance level of

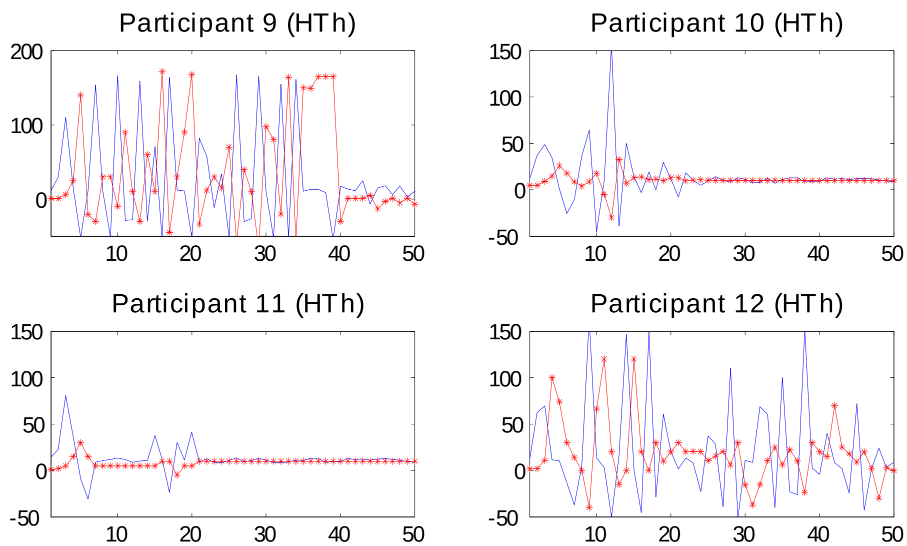

Figure 11: Predictions (red crosses) and inflation rates (blue line) for participants $9 - 12$ in treatment LTh.

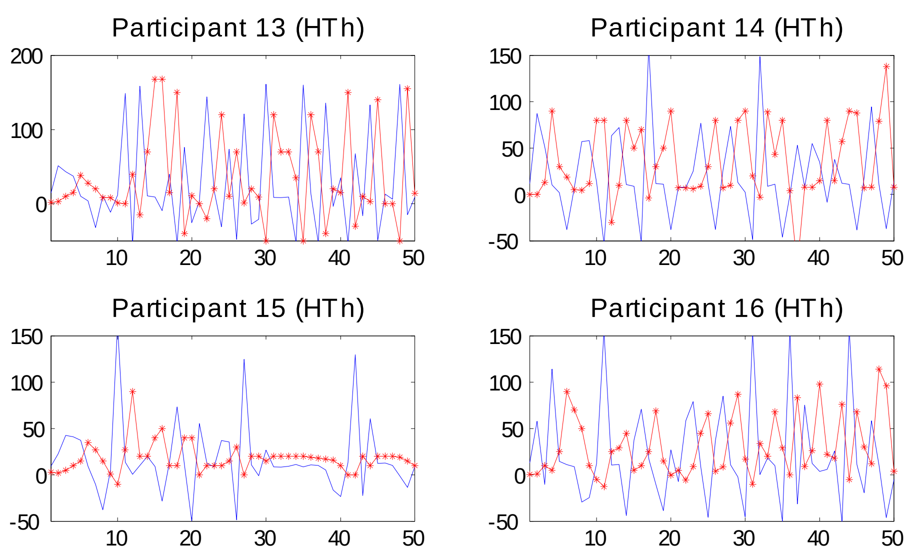

Figure 12: Predictions (red crosses) and inflation rates (blue line) for participants $13 - 16$ in treatment HTh.

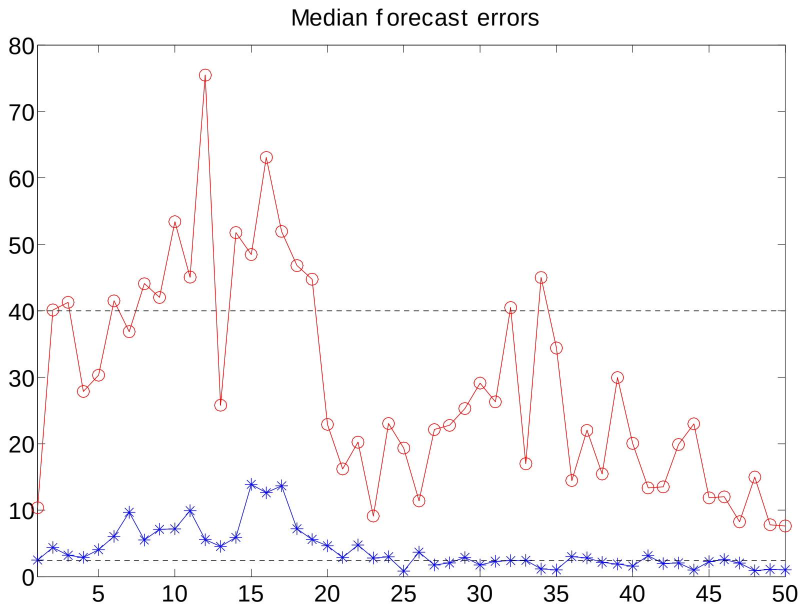

Figure 13: Median forecast errors in the two treatment (blue crosses correspond to treatment LTh, red circles correspond to treatment HTh). Above the dashed line at 40 earnings are zero, the lower dashed line at 2.4 corresponds to the critical value for RE (in treatment LTh).

virtually all participants manage to keep their prediction error below 40. A test at 5% significance level of the RE equilibrium on all 16 subjects yields poor results however. Given that the critical value should not be exceeded in 95% of the cases, at least 13 out of 16 predictions should fall within the RE confidence interval. This is only the case in periods 43 and 50. It can therefore not be said that the steady state equilibrium is a good description of individual behavior for the entire first treatment.

It is apparent from Figure 13 that absolute errors in treatment HTh are on the whole much higher than in treatment LTh. In fact, for a considerable number of periods up to nearly period 20, the median forecast error surpasses the zero earnings boundary of 40, meaning that at least half of the subjects earned nothing during those periods. At the same time the median error series clearly shows that the accu-

5%, the rational individual would make an absolute prediction error of no more than $\frac{19}{20} * 0.025\vartheta \approx 0.0238\vartheta$ in 95% of all periods. This means that the one-sided critical values at 5% for the absolute prediction error under the RE equilibrium are 2.4% and 2.6% approximately for treatments LTh and HTh respectively.

racy of predictions tends to rise during the experiment. At the end of the experiment however the median absolute error still is more than double the 95% critical value under the RE equilibrium. In fact, the number of REH-compatible predictions, at least within the length of the experiment, never exceeds six. Obviously this means that the RE equilibrium cannot be assumed to hold for the second treatment as a whole, or for any subset of periods within the experiment.

Prediction errors in both treatments tend to decrease during the experiment and, especially in treatment HTh, errors in the beginning of the experiment are at a level motivating many subjects to change their prediction rules. In fact, the thought process of the participants seems to involve two steps of reasoning. The first step is realizing that their own expectations drive the inflation dynamics, in the second step they use this knowledge in order to try to stabilize the dynamics. The results suggest that it is much harder to dampen oscillations in inflation rates in treatment HTh than in treatment LTh, even for participants understanding that their own behavior drives those oscillations. Taking these two steps of reasoning into account we exclude an initial phase of learning for the further analysis of convergence to the monetary steady state. Since the median absolute prediction error in both treatments makes a substantial drop in the periods leading up to period 20 and has local minima in period 21, the first 20 periods of the experiment are designated "learning phase", for all participants of both treatments.

In order to determine whether the inflation series generated by participants satisfy the REHs, the inflation means and standard deviations are compared to their approximate equilibrium distributions. The sample sizes are limited to 30 by excluding the learning phase determined above. Table 2 lists inflation averages and standard deviations for all participants from both treatments and compares these with the critical values under the RE equilibrium.

Comparing the inflation means for treatment LTh with the critical values under the REH gives mixed results. For 10 out of 16 subjects the null hypothesis of rational behavior is not rejected, while the remaining subjects all have inflation means significantly higher than the equilibrium rate. When restricting attention to the inflation mean, the majority of subjects is therefore in accordance with the REH, though this majority is smaller than the number of 14 required under the null hypothesis that the REH holds for the treatment as a whole.[26] From Table 2 we see that the REH fares slightly worse when testing the standard deviation in the first treatment. Six of the subjects, all of whom have an inflation mean consistent with the REH, have an inflation volatility low enough not to reject the possibility of rational behavior, while the others exceed the upper critical value. Similar to the

[26] Assuming the REH is true, the number of non-rejections is binomially distributed with success rate 95% and 16 trials. The critical value at 5% significance level for a one-sided test is 14.

LTh ($\vartheta = 1.01$)			**HTh** ($\vartheta = 1.11$)		
participant	$\overline{\pi}$	σ_π	participant	$\overline{\pi}$	σ_π
1	**0.962**	**1.721**	1	11.99	22.64
2	**0.789**	2.633	2	20.30	48.57
3	2.867	33.49	3	24.63	66.40
4	17.83	72.12	4	**11.11**	2.259
5	1.601	12.18	5	**11.14**	4.606
6	2.711	21.17	6	9.66	6.450
7	**0.774**	23.78	7	19.21	55.34
8	**0.787**	**1.618**	8	24.26	65.13
9	**0.826**	1.892	9	24.24	62.90
10	3.879	36.87	10	10.31	4.275
11	**0.820**	**1.585**	11	**10.81**	**1.744**
12	**0.949**	**1.731**	12	16.40	47.66
13	7.465	46.97	13	30.40	71.63
14	**0.820**	**1.514**	14	18.27	42.56
15	**0.812**	3.233	15	16.84	36.84
16	**0.820**	**1.585**	16	22.31	58.19
H_0 **not** rejected	**10**	**6**	H_0 **not** rejected	**3**	**1**

Table 2: Testing the Rational Expectations Hypothesis. Bold numbers correspond to cases where the null hypothesis of RE is not rejected. The critical values are $[0.478; 1.522]$ for the mean in treatment LTh, $[10.43; 11.57]$ for the mean in treatment HTh, $[1.084; 1.830]$ for the standard deviation in treatment LTh and $[1.192; 2.012]$ for the standard deviation in treatment HTh.

mean, the REH is unfit as a description of inflation volatility in treatment LTh as a whole, but for 6 out of 16 subjects the REH as a post-learning phase description of realized inflation rates cannot be rejected.

In the second treatment the REH fails almost completely as a description of inflation development. Only in three cases does the inflation mean fall inside the critical region, while the standard deviation with a single non-rejection does even worse. Participant 11, who uses a focal rule at 10% starting exactly after the end of the learning phase, is the only subject having both an inflation mean and standard deviation in accordance with the REH.[27] Summarizing, it is reasonable to say that

[27]Participants 6 and 10, with inflation means significantly below the equilibrium value, experienced very low inflation values directly following the learning phase, after which they both com-

the REH is a good description of almost a third of the inflation series when looking at their means. At the same time, the REH cannot explain the standard deviations, with 15 out of 16 inflation series exhibiting significant excess volatility. Obviously, this means that the REH is not suitable as a description of inflation development for treatment HTh as a whole.[28]

4.3 Individual expectation rules and switching

Above we argued that neither constant gain learning algorithms, nor myopic expectation rules give a satisfactory general description of our experimental results. Of course, it is possible that these models give an appropriate description of individual behavior of some of the participants. In this subsection we try to investigate in more depth the type of prediction strategies participants used in the experiment.

Looking at the inflation predictions of the participants in treatment LTh, grouped by the three categories according to Table 1, certain regularities stand out. Participants in category *S* tend to use adaptive predictions with a small adjustment parameter in the beginning of the experiment, while often fixing their predictions to specific numerical values or "focal points" after some time. Participants 9 and 11 clearly fit this description, with predictions slowly moving in the direction of recent realized inflations in the initial periods and sticking to specific values somewhere between periods 10 and 20. These values are very close but not identical: participant 9 fixes his predictions exactly at the monetary steady state of 1%, but deviates from it several times during the experiment. Participant 11, on the other hand, fixes his predictions at 1.11% from period 16 onwards.[29]

Participants 4, 10 and 13, constituting category *U*, all start out the experiment with expectations that are approximately naive. Under naive expectations the inflation rate diverges from the monetary steady state (as was established in Section

mitted to near-focal strategies. Extending the learning phase for both individuals by two periods to 22 in total, the averages are increased to 10.81 and 10.65 for participants 6 and 10, respectively. Adjusting for the change in the number of remaining observations the null hypothesis is no longer rejected for both participants. Standard deviations would decrease to 2.433 and 2.233, respectively, but still substantially exceed the upper critical value.

[28] Also earnings for both treatments are significantly below earnings under rational expectations (which are 23.78 euros for treatment LTh and 23.66 euros for treatment HTh). This is confirmed by one-sided t tests, giving resp. $H_0 : \mu_1 = \mu_1^{RE}$, $t \approx -3.91$ (15 df.), $p < 0.001$ and $H_0 : \mu_2 = \mu_2^{RE}$, $t \approx -11.2$ (15 df.), $p < 0.001$.

[29] Within category *S* the most frequently submitted predictions are at or around the monetary steady state. The two most frequently submitted values are 0% and 1%, making up respectively 31.4% and 22.1% of all predictions within category *S*. The maximum at 0% instead of the monetary steady state seems odd, but is caused mainly by the fact that participant 16 chooses this value in 49 periods. The percentages fall to 27.5% and 14.0% when all participants in treatment LTh are included.

2), leading in all cases to a strong upward inflationary movement in the beginning of the experiment. In none of the cases the participant subsequently is able to discover the key to stabilizing inflation, namely rigidity in the development of predictions. Until the end that is, since participants 4 and 13 keep their predictions constant in the last three periods, catching just a glimpse of its stabilizing effect on inflation. Participant 10 seems to switch to adaptive expectations with a low adjustment parameter in the last 17 periods.

Predictions of subjects in category *US* have a combination of characteristics already observed in the other two categories. The initial phase of high volatility resembles the permanently high volatility of participants 4, 10 and 13, in that it is started by expectations which are approximately naive (as for example for participants 3 and 6) or adaptive with high adjustment (participant 7). The subsequent phase of relative stability is, like the prediction series in category *S*, dominated by fixation on focal values (e.g. 0 for participants 3 and 7, 0 and 1 for participants 1 and 6), in some cases arrived at through an adaptive rule with slow adjustment (participant 1). Of the five *US* series only the one by participant 5 seems unaffected by these regularities, showing no signs of a naive or adaptive rule and no fixation at a fixed prediction value.

The absence of early convergence in the treatment HTh is caused to a considerable degree by the fact that in the first two periods, when no previously realized inflation rates were available, all participants submitted predictions lower than the equilibrium rate. In the directly following periods many adopted a prediction rule close to naive or high-adjustment adaptive expectations (e.g. participants 1, 10 and 13), driving the inflation rate further away from its steady state value. For the participants in the *US* category of treatment HTh we again have that often approximately naive or high-adjustment adaptive rules are applied at the beginning of the experiment (see e.g. participants 6 and 10) and low-adjustment adaptive or focal rules in the course of it (see e.g. participants 6 and 11 who consistently predict 10 for the last half of the experiment, and participant 10 who predicts very close to 10 for the last 27 periods).[30] As in treatment LTh, little more can be said about category *U* then that its subjects tend to initially use adaptive prediction rules with high adjustment towards the most recent realized inflation rate (e.g. subject 12 uses an approximate naive rule in roughly the first 13 periods of the experiment), leading to volatile inflation dynamics, which they subsequently are not able to fully stabilize.[31]

[30] The prediction values most frequently chosen in the *US* category, including the 4 boundary cases of participants 1, 7, 9 and 15, are 10 and 0, making up 11.2% and 7.3% resp. of all predictions in the category (percentages drop to 10.4% and 6.3% for the whole treatment). Note that none of the subjects in treatment HTh repeatedly chooses the equilibrium value of 11.

[31] Four subjects, participants 1, 7, 9 and 15, are somewhat stabilizing the dynamics later in the experiment but have not completed the process at period 50 (participants 7 and 9 temporarily use

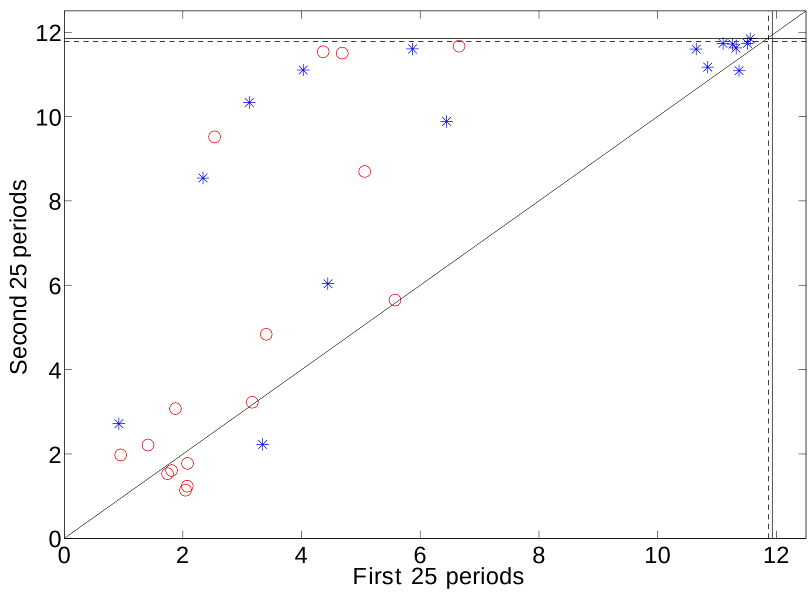

Figure 14: Payoffs in the first and second half of the experiment for both treatments (blue crosses correspond to participants in treatment LTh, red circles correspond to participants in treatment HTh). The solid horizontal and vertical lines correspond to payoffs under rational expectations in treatment LTh, and the dashed horizontal and vertical lines correspond to payoffs under rational expectation in treatment HTh.

The substantial number of participants experiencing high inflation volatility but managing to eliminate it during the experiment suggests that they may have switched prediction rules during the experiment to achieve better results.[32] Initially, many participants seem to use myopic expectations schemes, such as naive expectations or adaptive expectations with a high adjustment parameter, but subsequently, forced by the poor performance of these rules, change to predicting focal inflation rates, or use adaptive expectations with a low adjustment parameter. This switching behavior is, to a certain extent, confirmed by a closer inspection of participants' earnings during the experiment. Figure 14 shows average earnings of participants in terms of euros (see the linear scoring rule (15)), with the experiment divided in

focal prediction rules close to 0 after period 40; participants 1 and 15 use rules close to slow-adjustment adaptive expectations in a substantial part of the experiment but are not careful enough in applying them).

[32]Note however that a dramatic change in inflation development does not necessarily imply that a similarly distinct change occurred in the prediction rule. As explained in Section 2, small changes in parameters of prediction rules may cause a transition from instability to stability.

two equal halves of 25 periods. When a structural change in a participants prediction rule changes the inflation dynamics from unstable to stable, the participant will, largely independent of the exact period in which the change is implemented, earn substantially more in the second half of the experiment than in the first. Such a participant should show up in Figure 14 in the left upper part, well above the diagonal. In those cases in which the inflation dynamics are characterized by uniform instability or stability, earnings should be low or high respectively in both halves of the experiment, translating into positions around the far ends of the diagonal in Figure 14.

All of the above cases are indeed present in Figure 14. Close to the diagonal at the right upper corner lies a cluster of subjects which corresponds to the eight participants in category *S* of treatment LTh, some of which approach the average earnings under the rational expectations equilibrium in the second half of the experiment. Around the lower part of the diagonal fourteen participants (three from treatment LTh and eleven from treatment HTh) are depicted that earn not more than 6 euro in each half of the experiment. All of these participants were classified in category *U* of their respective treatment. The most interesting group of participants is located well above the diagonal and towards the left of Figure 14. These participants, five from each treatment, had poor earnings in the first half of the experiment, but much improved earnings in the second half of the experiment. This is consistent with the hypothesis that a substantial number of participants from both treatments changed their prediction rules during the experiment and as a result strongly improved inflationary stability and their own predictive accuracy.

5 Concluding Remarks

In this article we use laboratory experiments with incentivized human subjects to investigate the way in which individuals predict inflation rates in a standard overlapping generations model. We find that the ability to successfully predict inflation rates varies widely, both between subjects of the same treatment and between the treatments as a whole. Consequently, the Rational Expectations Hypothesis on the whole does a poor job of explaining the experimental results, even when allowing for a substantial learning phase. It does much better though in the first treatment than in the second and in both treatments the REH describes the inflation means much better than the inflation volatility. Regarding the difference between the treatments, it is clear that the inflation elasticity of the savings rate around the equilibrium rate of inflation is related inversely to the number of subjects succeeding in stabilizing inflation development within the 50 periods of the experiment. We identify three types of individual prediction behavior. Some participants learn to forecast accurately and stabilize the inflation rate almost from the outset of the experiment, whereas others only learn this during the experiment. The last subset of participants never learns to predict inflation rates to some degree of accuracy at all. The latter category contains much more participants from the second treatment, whereas the first category only contains participants from the first treatment.

We draw the following conclusions. First, the results are consistent with participants using constant gain algorithms, such as adaptive expectations with low adjustment rates, but not consistent with participants consistently using either myopic expectations, such as naive expectations, or decreasing gain algorithms, such as (ordinary) least squares learning. Second, although we designed the experiment in order to discriminate between different stationary prediction strategies (as the ones discussed in Section 2), the experimental results suggest that many participants actually switch from one prediction strategy to another, on the basis of forecasting performance. Therefore, although participants typically do not use least squares learning, a substantial number of them does try to improve the way in which they form expectations, allowing them to eventually end up close to the rational expectations steady state. Interestingly, models with evolutionary switching between different prediction strategies or heuristics have become more popular in macroeconomics in recent years, see e.g. Marcet and Nicolini (2003), Tuinstra and Wagener (2007), Branch and McGough (2010), Lines and Westerhoff (2010) and Anufriev et al. (2012) for some recent applications. Our experimental results support this type of approach to expectation formation in macroeconomics.[33]

The experiment discussed in this article uses a representative agent approach, where the predictions of a single individual determine the realized inflation rate. A future experiment could extend this to a group experiment, where inflation predictions of a number of individuals would determine the inflation rate. Whether multi-agent treatments on the whole produce more stable inflation series is a priori unclear. On the one hand there will be a smoothing effect on the inflation dynamics since the effect of individual outliers is limited and fluctuations in individual predictions cancel out against each other. On the other hand, however, it would be more difficult for individual participants to stabilize the dynamics by consistently choosing a focal inflation rate.

[33]In fact, Anufriev and Hommes (2012) use such a switching model to successfully explain the data from the earlier ‘Learning to Forecast’ asset pricing experiment presented in Hommes et al. (2005).

Appendix: Experimental Protocol and Instructions

The experiment was performed in two sessions, both taking place on April 25, 2006. In both sessions an experimental treatment was completed. The experiment was conducted in the CREED laboratory at the economics faculty of the University of Amsterdam.

During both sessions, when an excess of participants had left, a short welcoming message was read aloud from paper, after which participants were randomly assigned a place in the laboratory. Subjects were received in a room separate from the laboratory. The required amount of subjects was in both cases 16, making up one experimental treatment. In the case of an excess of subjects, volunteers were first asked to accept an immediate reward of 5 euros not to participate in the experiment; if the number of volunteers was not sufficient to remove the excess, people were randomly assigned to take 5 euros and leave until the required number was reached. The laboratory consists of several rows of cubicles equipped with computers, at the time of the experiment supplemented with pencil, paper and pocket calculator. The experiment was fully computerized, so no instructions were handed out in the laboratory.

When all participants were seated they were asked to begin reading the instruc-tions on their computer screens. Instructions were identical for both treatments. It had been made clear to the participants that they could at any time call one of the experimenters if they had a question. After everyone had finished reading the instructions, the experiment automatically started. When the experiment was finished, the participants were called to the reception room one by one to receive their earnings in cash. They left the laboratory immediately afterwards.

Below a translation follows of the instructions presented to the participants before the experiment started. A set of screenshots of the instructions in the original form is attached.

Translation of Experimental Instructions

Welcome to the economic laboratory

Experimental Team

Peter Heemeijer, Cars Hommes, Joep Sonnemans & Jan Tuinstra.

— Computer program written in *Mathematica* by Peter Heemeijer —

Structure of the experiment

You are part of an experiment about economic decision making. You will be rewarded based on the decisions you make during the experiment. The experiment will be preceded by a number of pages with instructions, which will explain to you how exactly it works. After the experiment you will be asked to answer a number of questions regarding its results.[34]

- The entire experiment, including the instructions and the questionnaire, is run on the computer. Therefore, you are not required to submit the sheet of paper on your desk, but you can use it to make notes.
- On your desk is a calculator. If necessary, you can use it during the experiment.
- If at any time you have a question, raise your hand, then someone will approach you for assistance.

General information about the experiment

You are a *statistical research bureau* that earns its income by making predictions of the *price level of consumption goods* in the economy. In particular you regularly make predictions concerning the change in the price level of consumption goods, i.e. the *inflation*. The experiment consists of a total of *50 periods*. In each period you are asked to predict the inflation in the price of consumption goods; your *reward at the end of the experiment* will be based on the accuracy of your predictions.
In the following instructions you will get more information about the *economy* in which you operate, about the *market for consumption goods* that your predictions are relevant to, and about the exact way in which *making predictions* works during the experiment. Also, the *computer program* used during the experiment will be explained.

Information about the economy

The economy you are part of splits into a *young population segment*, consisting of individuals of employable age (roughly 18 through 65 years old), and an *old population segment*, consisting of individuals who no longer work because of their age. Individuals from the young population segment receive an income consisting of a *fixed number of consumption goods*; individuals from the old segment, who no longer work, receive no income. In the economy the possibility exists for young people to *save* part of their income in order to obtain consumption goods when they have reached old age.

Your predictions of the inflation are used by individuals from the young population segment to determine *which part* of their present income they will save

[34] In reality however subjects did not answer any questions at the end of the experiment. Originally, the intention was to present participants with a questionnaire, but due to a technical problem this was cancelled.

for the time when they themselves will belong to the old segment. The *money* which individuals use to keep their savings, is brought into circulation by a *central bank*.

Information about the market for consumption goods

In every period you make a prediction about the inflation in the price of consumption goods. Based on this young people determine which part of their income they will save by *selling* goods; the remaining goods they consume before retiring. Old people are not sensitive to your prediction, since they will always use their accumulated savings to *buy* consumption goods. The *true price level* on the market for consumption goods is in each period determined in such a way that with the savings of old people the available goods, offered by young people, can exactly be bought. Also, the price level is in each period slightly influenced in an unpredictable way by circumstances in the rest of the economy.

As stated earlier, a *central bank* brings money into circulation that individuals use to save. It is known that this bank has the tendency to let the *total money supply* increase slowly.

Information about making predictions

As stated earlier, the experiment consists of a total of *50 periods*. In each period you make a prediction about the inflation in the price of consumption goods. Because the true price level in every period is partly determined by the expected price level in the following period, which after all makes young people decide how much of their consumption goods to sell, your predictions of the inflation will be *one period ahead*.

Suppose for example that the experiment has progressed to period 12. You will then predict the inflation in the next period 13, i.e. the change in the price level between periods 12 and 13 [this is an error; given the inflation definition $\pi_t = \frac{p_{t+1}}{p_t}$ it should be “between periods 13 and 14”]. When making your prediction you can use the following *information* (which will be shown on your computer screen): the *inflations* up to and including the previous period 11, and your *predictions of the inflation* up to and including the present period 12. Notice that in period 12 the inflation prediction for that period (i.e. the change in price level between periods 11 and 12) [should be “between periods 12 and 13”] is already available to you, because you also predicted one period ahead in the previous period 11.

Information about your reward (part 1 of 2)

Your *reward after the experiment has ended* increases with the accuracy of your predictions. In the experiment this accuracy is measured by the *absolute error* between your predictions of the inflation and its true values. For each period this absolute error is calculated as soon as the true value of the inflation is known; you will then receive a Prediction Score that increases the smaller the absolute error becomes. The *table below* gives the relation between the absolute prediction error and the Prediction Score. If for a certain period you predict for example an inflation of 2%, and the true inflation turns out to be 7%, then you make an absolute error of $7\% - 2\% = 5\%$. You will therefore receive a Prediction Score of 80.

INSERT PAYOFF TABLE

Remark: The table serves as an illustration and contains only part of all possible prediction errors.

Information about your reward (part 2 of 2)

If you predict an inflation of 2%, and it turns out to be -3%, you also make a prediction error of $2\% - (-3\%) = 5\%$. You will therefore receive the same Prediction Score of 80. For a perfect prediction, with a prediction error of zero, you will receive a Prediction Score of 100. When you make a prediction error of 25% or more, your Prediction Score for the relevant period will be zero. Generally speaking, your Prediction Score decreases with four points if your prediction error increases with one percent. Notice that the inflation and your predictions of it in the experiment are expressed in *percentages*; naturally the inflation can become negative, because the price level might decrease.

Your *Final Score* at the end of the experiment consists simply of a summation of all Prediction Scores you have earned during the experiment. During the experiment your scores are shown on your computer screen; at the end of the experiment you will be shown an overview of your Prediction Scores, followed by the resulting Final Score. Your monetary reward consists of 1/2 *eurocent* for each point in your Final Score (so 200 points equals 1 euro).

Information about the computer program (part 1 of 3)

Directly below you see an example of the *left upper part* of the computer screen during the experiment.[35] It is a *graphical depiction* of the inflations in the price of consumption goods (red series) and your predictions of them (yellow series). On the horizontal axis are the time periods; the vertical axis has percentage as unit. In the *imaginary situation* depicted in the graph, the experiment is in period 28 and you are predicting the inflation in period 29 (the experiment lasts for 50 periods). Notice that the graph only shows results of the last 25 periods and that the next period is always at the far-right side.

Information about the computer program (part 2 of 3)

[35]The figures referred to are included in the original instructions directly following these translations.

To the side you see an example of the *right upper part* of the computer screen during the experiment. It consists of a *table* with information on the history of the experiment up to a maximum of 17 periods. This information is supplementary to the graph in the left upper part of the screen. The *first column* of the table shows the time period (the next period, in the example 29, is always on top). The *second and third columns* respectively show the inflations in the price of consumption goods and your predictions of them. Finally, the *fourth column* lists your Prediction Score for each period separately. Notice that you can use the sheet of paper on your desk to save data longer than 17 periods.

Information about the computer program (part 3 of 3)
Below you see an example of the bottom part of the computer screen during the experiment. In each period you will be asked to submit your prediction of the inflation in the next period (directly beneath *Submit your prediction*). In the *first two periods of the experiment* no information is yet known about the market (the true inflation in period 1 becomes known as soon as you have submitted your prediction for period 2); submit in these periods inflation predictions that seem reasonable at first glance. When entering your prediction, use the *decimal point* if necessary. Notice that your predictions and the true inflation in the experiment are rounded to two decimals. In addition Prediction Scores are rounded to whole numbers.

— *You are now ready to start the experiment* —

Please wait until all participants are ready to begin the experiment
Remark: When all participants are ready, the experiment will start immediately. There are consequently no practice rounds. When the experiment begins, it may take a moment for the experimental program to load.

CeNDEF

Center for Non-linear Dynamics in Economics and Finance

WELKOM IN HET ECONOMISCH LABORATORIUM

<u>Experimenteel Team</u>

Peter Heemeijer, Cars Hommes, Joep Sonnemans & Jan Tuinstra.

— Computerprogramma geschreven in *Mathematica* door Peter Heemeijer —

Opzet van het experiment

U doet mee aan een experiment naar economische besluitvorming. Op basis van de beslissingen die u tijdens het experiment neemt, zult u worden beloond. Het experiment zal voorafgegaan worden door een aantal pagina's aan instructies, waarin wordt uitgelegd hoe het in zijn werk gaat. Na afloop van het experiment zult u worden gevraagd om een aantal vragen over het verloop ervan te beantwoorden.

- Het gehele experiment, inclusief de instructies en de vragenlijst, verloopt via de computer. U hoeft het papier dat op uw bureau ligt dus niet in te leveren, maar u kunt het gebruiken om aantekeningen te maken.
- Op uw bureau ligt een rekenmachine. Die kunt u, zo nodig, tijdens het experiment gebruiken.
- Heeft u op enig moment een vraag, steek dan uw hand op, dan komt iemand u helpen.

Algemene informatie over het experiment

U bent een *statistisch onderzoeksbureau* dat zijn inkomsten verkrijgt door het doen van voor-spellingen over het *prijsniveau van consumptiegoederen* in de economie. In het bijzonder doet u regelmatig voorspellingen over de verandering

Published by De Gruyter, 2012

in het prijsniveau van consumptiegoederen, ofwel de *inflatie*. Dit experiment bestaat in totaal uit *50 perioden*. In iedere periode wordt u gevraagd om een voorspelling te doen van de inflatie in de prijs van consumptiegoederen; uw *beloning na afloop van het experiment* is gebaseerd op de nauwkeurigheid van uw voorspellingen. In de hierna volgende instructies krijgt u meer informatie over de *economie* waarin u zich bevindt, over de *markt voor consumptiegoederen* waarop uw voorspellingen betrekking hebben, en over de manier waarop *het doen van voorspellingen* tijdens het experiment in zijn werk gaat. Daarnaast zal het *computerprogramma* worden toegelicht dat tijdens het experiment wordt gebruikt.

Informatie over de economie

De economie waarin u zich bevindt, valt uiteen in een *jonge bevolkingsgroep*, bestaande uit personen van werkzame leeftijd (ruwweg 18 tot en met 65 jaar), en een *oude bevolkingsgroep*, bestaande uit personen die vanwege hun leeftijd niet langer werkzaam zijn. Personen uit de jonge bevolkingsgroep ontvangen een *inkomen* bestaande uit een *vast aantal consumptiegoederen*; personen uit de oude groep, die niet meer werken, ontvangen geen inkomen. In de economie bestaat de mogelijkheid voor jonge personen om een deel van hun inkomen te *sparen* om ook op hun oude dag over consumptiegoederen te kunnen beschikken. Uw voorspellingen van de inflatie worden gebruikt door personen uit de jonge bevolkingsgroep om te bepalen *welk deel* van hun huidige inkomen ze zullen sparen voor de tijd waarin ze zelf tot de oude bevolkingsgroep behoren. Het *geld* waarin personen hun besparingen aanhouden, wordt in omloop gebracht door een *centrale bank*.

Informatie over de markt voor consumptiegoederen

In iedere periode van het experiment doet u een voorspelling over de inflatie in de prijs van consumptiegoederen. Op basis hiervan bepalen jonge personen welk deel van hun inkomen ze zullen sparen door goederen te *verkopen*; de rest maken ze op voordat ze met pensioen gaan. Oude mensen zijn niet gevoelig voor uw voorspelling, omdat ze altijd hun eerder opgebouwde besparingen gebruiken om consumptiegoederen te *kopen*. Het *werkelijke prijsniveau* op de markt voor consump-tiegoederen krijgt steeds een hoogte zodanig dat oude personen met hun spaargeld precies de beschikbare goederen, aangeboden door jonge personen, kunnen aanschaffen. Daarnaast wordt het prijsniveau in iedere periode op onvoorspelbare wijze lichtelijk beïnvloed door omstandigheden in de rest van de economie. Zoals gezegd brengt een *centrale bank* het geld in omloop dat mensen gebruiken om te sparen. Het is bekend dat deze de gewoonte heeft om de *totale geldhoeveelheid* langzamerhand te laten toenemen.

Informatie over het doen van voorspellingen

Zoals gezegd bestaat het experiment in totaal uit *50 perioden*. In iedere periode doet u een voorspelling over de inflatie in de prijs van consumptiegoederen. Omdat het werkelijke prijsniveau in iedere periode mede bepaald wordt door het verwachte prijsniveau in de periode erna, dat jonge mensen immers doet beslissen hoeveel van hun consumptiegoederen ze verkopen, voorspelt u de inflatie steeds *één periode vooruit*. Stel bijvoorbeeld dat het experiment zich bevindt in periode 12. U voorspelt dan de inflatie in de volgende periode 13, ofwel de verandering in het prijsniveau tussen periode 12 en 13. Bij het doen van uw voorspelling kunt u gebruik maken van de volgende *informatie* (die op uw computerscherm zal worden afgebeeld): de *inflaties* tot en met de vorige periode 11, en uw *voorspellingen van de inflatie* tot en met de huidige periode 12. Merk op dat u in periode 12 reeds beschikt over uw inflatievoorspelling voor die periode (dat wil zeggen de verandering in prijsniveau van periode 11 naar 12), omdat u ook in de vorige periode 11 één periode vooruit voorspelde.

Informatie over uw beloning

(deel 1 van 2)

Uw *beloning na afloop van het experiment* neemt toe met de nauwkeurigheid van uw voorspellingen. In het experiment wordt deze nauwkeurigheid gemeten door de *absolute fout* tussen uw voorspellingen van de inflatie en de werkelijke waarden daarvan. Voor iedere periode wordt deze absolute fout berekend zodra de werkelijke waarde van de inflatie bekend is; u ontvangt vervolgens een *Voorspelscore* die toeneemt naar mate de absolute fout kleiner is. De *onderstaande tabel* geeft het verband weer tussen de absolute voorspelfout en de Voorspelscore. Als u voor een zekere periode bijvoorbeeld een inflatie van 2% voorspelt, en de werkelijke inflatie blijkt 7% te zijn, dan maakt u een absolute fout van 7% – 2% = 5%. U ontvangt dan een Voorspelscore van 80.

— *Alinea gaat verder op de volgende pagina* —

Absolute voorspelfout	0	5	10	15	20	≥ 25
Voorspelscore	100	80	60	40	20	0

Opmerking: De tabel dient ter illustratie en bevat slechts een aantal van de mogelijke voorspelfouten.

Informatie over uw beloning

(deel 2 van 2)

Als u een inflatie van 2% voorspelt, en deze blijkt –3% te zijn, maakt u eveneens een voorspelfout van 2% – (–3%) = 5%. U ontvangt daarom dezelfde Voorspelscore van 80. Voor een perfecte voorspelling, met een voorspelfout van nul, ontvangt u een Voorspelscore van 100; als u een voorspelfout maakt van 25% of meer, is uw Voorspelscore in de betreffende periode nul. In het algemeen daalt uw Voorspelscore met vier punten als uw voorspelfout met een procentpunt toeneemt. Merk op dat de inflatie en uw voorspellingen ervan in het experiment worden uitgedrukt in *procenten*; de inflatie kan vanzelfsprekend negatief worden, omdat het prijsniveau van consumptiegoederen kan dalen. Uw *Totaalscore* aan het einde van het experiment bestaat simpelweg uit een opsomming van alle Voorspelscores die u tijdens het experiment heeft behaald. Tijdens het experiment worden uw scores afgebeeld op uw computerscherm; na afloop van het experiment krijgt u een overzicht te zien van uw Voorspelscores, gevolgd door de resulterende Totaalscore. Uw uiteindelijke beloning bestaat uit *1/2 eurocent* voor iedere punt in uw Totaalscore (200 punten staat dus gelijk aan 1 euro).

Informatie over het computerprogramma

(deel 1 van 3)

Hieronder ziet u een voorbeeld van het *linkerbovendeel* van het computerscherm

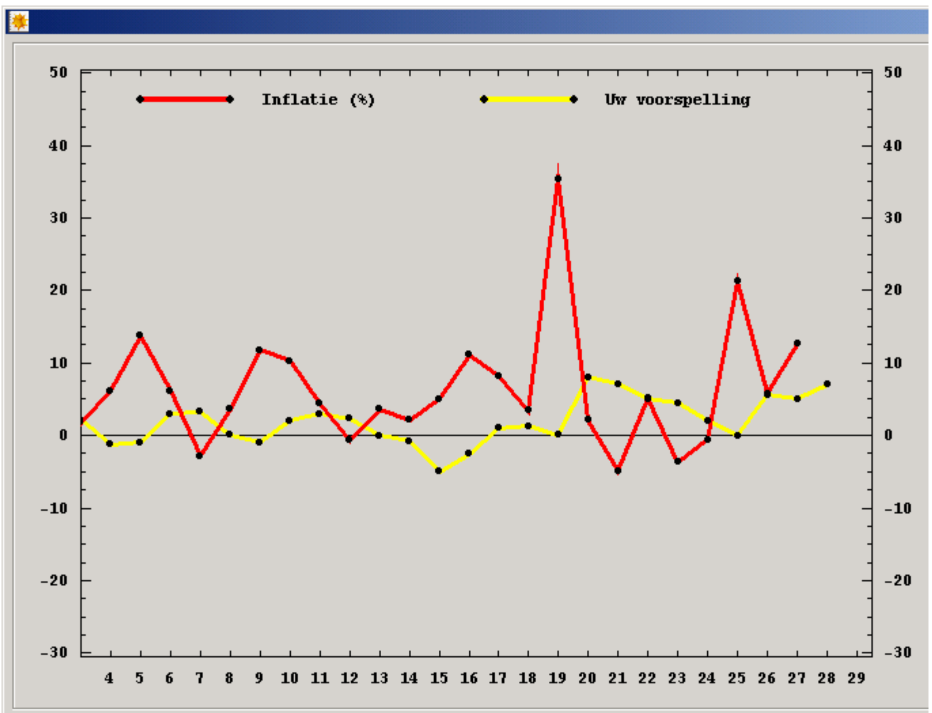

tijdens het experiment. Het bestaat uit een *grafische voorstelling* van de inflaties in de prijs van consumptiegoederen (rode reeks) en uw voorspellingen daarvan (gele reeks). Op de horizontale as staan de *tijdsperioden* vermeld; de verticale as is in procenten. In de *denkbeeldige situatie* die de grafiek toont, is het experiment in periode 28 en voorspelt u de inflatie in periode 29 (het experiment duurt 50 perioden). Merk op dat de grafiek slechts resultaten van ten hoogste de laatste 25 perioden toont en dat de eerstvolgende periode altijd helemaal rechts staat.

Informatie over het computerprogramma (deel 2 van 3)

Hiernaast ziet u een voorbeeld van het *rechterbovendeel* van het computerscherm tijdens het experiment. Het bestaat uit een *tabel* met informatie over het verloop van het experiment in ten hoogste de laatste 17 perioden. Deze informatie is een aan-vulling op de grafiek in het linkerbovendeel van het scherm. De *eerste kolom* van de tabel laat de tijdsperiode zien (de eerstvolgende periode, in het voorbeeld 29, staat altijd bovenaan). De *tweede en derde kolom* respectievelijk tonen de inflaties in de prijs van consumptiegoederen en uw voorspellingen daarvan. De *vierde*

Tijds-periode	**Inflatie (%)**	**Uw Voor-spelling**	***Voorspel-score***
29		???	
28		7.00	
27	12.67	5.00	69
26	5.83	5.56	99
25	21.28	0.00	15
24	-0.56	2.00	90
23	-3.64	4.50	67
22	5.16	5.00	99
21	-4.93	7.12	52
20	2.17	8.00	77
19	35.42	0.11	0
18	3.45	1.25	91
17	8.27	1.00	71
16	11.13	-2.56	45
15	5.03	-5.00	60
14	2.12	-0.75	89
13	3.61	0.00	86
12	-0.66	2.30	88
11	4.36	3.00	95

kolom ten slotte geeft uw Voorspelscore bij iedere afzonderlijke periode. Merk op dat u het kladvel op uw bureau kunt gebruiken om gegevens langer dan 17 perioden te bewaren.

Informatie over het computerprogramma (deel 3 van 3)

Hieronder ziet u een voorbeeld van het *onderste deel* van het computerscherm tijdens het experiment. In iedere periode wordt u gevraagd om uw voorspelling van de inflatie in de volgende periode te geven (onder *Vul uw voorspelling in*). In de *eerste twee perioden van het experiment* is er nog geen informatie bekend over

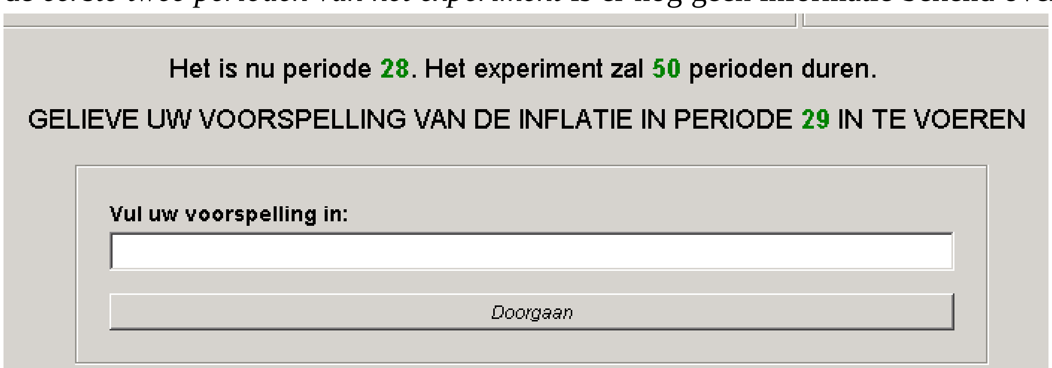

de markt (de werkelijke inflatie in periode 1 wordt bekend zodra u uw voorspelling voor periode 2 heeft opgegeven); doe in deze perioden inflatievoorspellingen die u op het eerste gezicht redelijk lijken. Maak bij het invoeren van uw voorspelling zo nodig gebruik van de *decimale punt*. Merk op dat uw voorspellingen en de werkelijke inflatie in het experiment worden afgerond op twee decimalen. Daarnaast worden Voorspelscores afgerond op gehele punten.

— U bent nu gereed om aan het experiment te beginnen —

GELIEVE TE WACHTEN TOT ALLE PARTICIPANTEN

GEREED ZIJN OM MET HET EXPERIMENT TE BEGINNEN

Opmerking: Als alle participanten gereed zijn, begint het experiment meteen.

Er zijn dus geen oefenrondes. Het kan straks even duren voordat het experimentele programma geladen is.

References

[1] Adam, K. (2007). Experimental evidence on the persistence of output and inflation. *Economic Journal*, 117, 603-636.

[2] Anufriev, M. and C. Hommes (2012). Evolution of market heuristics. *Knowledge Engineering Review*, 27, 255 - 271.

[3] Anufriev, M., T. Assenza, C. Hommes and D., Massaro (2012). Interest rate rules with heterogeneous expectations. *Macroeconomic Dynamics*, forthcoming.

[4] Assenza, T., P. Heemeijer, C.H. Hommes and D. Massaro (2011). Individual expectations and aggregate macro behavior. *CeNDEF Working paper 11-01* University of Amsterdam

[5] Branch, W.A., and B. McGough (2010). Dynamic predictor selection in a new Keynesian model with heterogeneous expectations. *Journal of Economic Dynamics & Control*, 34, 1492-1508.

[6] Bullard, J. (1994). Learning equilibria. *Journal of Economic Theory*, 64, 468-485.

[7] Cho, I.-K., N. Williams, and T.J. Sargent (2002). Escaping Nash inflation. *Review of Economic Studies*, 69, 1-40.

[8] Evans, G.W., and S. Honkaphja (2001). *Learning and Expectations in Macroeconomics.* Princeton: Princeton University Press.

[9] Grandmont, J. (1985). On endogenous competitive business cycles. *Econometrica*, 53, 995-1045.

[10] Heemeijer, P., C. Hommes, J. Sonnemans, and J. Tuinstra (2009). Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. *Journal of Economic Dynamics & Control*, 33, 1052-1072.

[11] Hommes, C. (2011). The heterogeneous expectations hypothesis: Some evidence from the lab. *Journal of Economic Dynamics & Control*, 35, 1-24.

[12] Hommes C, J. Sonnemans, J. Tuinstra, and H. van de Velden (2005). Coordination of expectations in asset pricing experiments. *Review of Financial Studies*, 18, 955-980.

Published by De Gruyter, 2012

[13] Hommes C, J. Sonnemans, J. Tuinstra, and H. van de Velden (2007). Learning in cobweb experiments. *Macroeconomic Dynamics*, 11, 8-33.

[14] Lucas, R.E. (1986). Adaptive behavior and economic theory. *Journal of Business*, 59, S401-S426.

[15] Lines, M., and F. Westerhoff (2010). Inflation expectations and macroeconomic dynamics: The case of rational versus extrapolative expectations. *Journal of Economic Dynamics & Control*, 34, 246-257.

[16] Marcet, A., and J.P. Nicolini (2003). Recurrent hyperinflations and learning. *American Economic Review*, 93, 1476-1498.

[17] Marcet, A,. and T.J. Sargent (1987). Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory*, 48, 337-368.

[18] Marimon, R., S.E. Spear, and S. Sunder (1993). Expectationally driven market volatility: An experimental study. *Journal of Economic Theory*, 61, 74-103.

[19] Marimon, R., and S. Sunder (1993). Indeterminacy of equilibria in a hyperinflationary world: Experimental evidence. *Econometrica*, 61, 1073-1107.

[20] Marimon, R., and S. Sunder (1994). Expectations and learning under alternative monetary regimes: An experimental approach. *Economic Theory*, 4, 131-162.

[21] Orphanides, A. and J.C. Williams (2005). The decline of activist stabilization policy: Natural rate misperceptions, learning and expectations. *Journal of Economic Dynamics & Control*, 29, 1927-1950.

[22] Pfajfar, D., and B. Žakelj (2010). Inflation expectations and monetary policy design: Evidence from the laboratory. Working paper University of Tilburg.

[23] Sargent, T.J. (1993). *Bounded Rationality in Macroeconomics*. Oxford: Clarendon Press.

[24] Schönhofer, M. (1999). Chaotic learning equilibria. *Journal of Economic Theory*, 89, 1-20.

[25] Sonnemans, J., and J. Tuinstra (2010). Positive expectations feedback experiments and number guessing games as models of financial markets. *Journal of Economic Psychology*, 31, 964-984.

[26] Tuinstra, J. (2003). Beliefs equilibria in an overlapping generations model. *Journal of Economic Behavior & Organization*, 50, 145-164.

[27] Tuinstra, J., and F.O.O. Wagener (2007). On learning equilibria. *Economic Theory*, 30, 493-513.

[28] Wenzelburger, J. (2002). Global convergence of adaptive learning in models of pure exchange. *Economic Theory*, 19, 649-672.