Should Macroeconomists Consider Restricted Perceptions Equilibria? Evidence from the Experimental Laboratory

Klaus Adam*
CEPR, London and European Central Bank
Postfach 16 03 19
60066Frankfurt a.M.
Germany
email: klaus.adam@ecb.int

May 1, 2004, preliminary

Abstract

This paper studies a simple model of output and inflation in the experimental laboratory. While the Rational Expectations Equilibrium (REE) predicts output and inflation to be white noise processes, output and inflation in experimental sessions display stable cyclical patterns. For about 50 model periods agents' expectations, which are the sole source of these patterns, are described extremely well by a Restricted Perceptions Equilibrium (RPE). In this equilibrium agents use the univariate forecast function which generates the lowest mean squared forecast error at the 1-step forecast horizon and iterate these forecasts to derive multi-step predictions. After about 50 model periods agents seem to learn that their simple univariate forecast function is misspecified and start to employ different forecast models for different prediction horizons. The data suggests that the different models are again optimal univariate forecast functions and evidence in favor of convergence towards the REE remains weak, even after more than 100 model periods. However, for model parameterizations where an RPE does not exist, agents' expectations are captured relatively well by the REE.

KEYWORDS: Experiments, Equilibrium Selection, Restricted Perceptions Equilibrium, Univariate Forecast Functions
JEL-Class.No.: E32, E37, C91

*I would like to thank the Carisal Foundation in Salerno, Italy, and the Center for Financial Studies in Frankfurt for financial support. Thanks go to seminar participants at the University of Helsinki and the University of Frankfurt.

1 Introduction

A number of recent contributions has considered economic models where agents can choose their forecast functions only subject to constraints (e.g. Evans and Ramey (1992), Evans and Honkapohja (1993),(2001), Sargent (1999), Adam (2001)). The presence of forecasting constraints may generate so-called Restricted Perceptions Equilibria where agents' expectations are only constrained rational since binding forecasting constraints prevent full rationality.[1],[2]

Model behavior in a Restricted Perceptions Equilibrium typically differs considerably from behavior in a Rational Expectations Equilibrium. As a result, models whose rational expectations performance is rather poor may perform much better when considering their performance under restricted perceptions (e.g. Evans and Ramey (1992), Ball (2000), Adam (2001)).

The aim of this paper is to provide an assessment of the empirical relevance of Restricted Perceptions Equilibria versus Rational Expectations Equilibria using laboratory experiments. Since Rational Expectations Equilibria can be interpreted as Restricted Perceptions Equilibria where forecast restrictions are completely relaxed, part of this assessment consists of identifying empirically plausible forecasting restrictions.

Resorting to laboratory experiments is justified on the grounds that it is rather difficult to identify empirically plausible constraints using field data. Since expectations in the field remain largely unobservable, empirical tests are always joint tests of the underlying economic model and the constraints imposed on the forecasting schemes. Consequently, to identify forecasting constraints with field data one would need a 'true' economic model, which is something only few economists could potentially agree about. Without such a true model imposing restrictions on agents' forecasting schemes must be considered a rather dubious endeavor, since one might suspect finding almost always a restriction that makes a given model consistent with the data.

Relying on laboratory experiments has the paramount advantage that one can disentangle issues related to the plausibility of the model from those regarding how well expectations might be described through optimal expectations that are subject to constraints: in laboratory experiments the economic model is true by definition and agents' expectations can be made directly observable.

[1] Even though constraints are binding in such an equilibrium this does not imply that the forecasting constraints have been chosen such that they prevent agents from being fully rational in any case. This subtle but important difference arises due to the self-referential nature of forecast rationality, see Adam (2001).

[2] Since deviations from complete rationality is the result of optimal forecasting behavior that is subject to (forecasting) constraints, these models are robust to the Lucas critique since expectations functions remain to be determined inside the economic model (see also Evans and Ramey (2001)).

To study the plausibility of Restricted Perceptions Equilibria this paper will implement the cash-in-advance model of Adam (2001) in the experimental laboratory. This model is particularly suited to test whether univariate forecast functions constitute an empirically plausible restriction capable of explaining deviations from full forecast rationality. This is the case because the model possesses a Restricted Perceptions Equilibrium (RPE) besides the standard Rational Expectations Equilibrium (REE) if and only if agents' expectations functions are restricted to univariate forecast functions and when the elasticity of labor supply is sufficiently high.[3] The RPE does not exist and the REE remains the unique equilibrium if either agents can use forecast functions with more than one variable or the elasticity of labor supply is sufficiently low.

This feature allows to design two experimental treatments: a *high-elasticity* treatment where an RPE coexists with an REE if univariate forecast restrictions are in place, and a *low-elasticity* treatments where the REE is the only equilibrium, even if agents used only univariate forecast functions. While the emergence of a RPE in high-elasticity treatments would point towards the existence of univariate forecast restrictions, the low-elasticity treatments can be used to assess whether there is anything particular about the REE that might prevent agents to coordinate on it.

The output and inflation series in the experimental laboratory were generated as follows. The temporary equilibrium equation of the underlying economic model determines the current values of output and inflation as a function of lagged values and agents' expectations of the 1-step and 2-step ahead inflation rate.[4] Subjects participating in the experiments could observe these lagged values of output and inflation and were then asked to forecast inflation rates for the next 2 periods. The new output level and inflation rate were then computed by substituting inflation expectations in the temporary equilibrium with the average forecasts entered by subjects.

Once the new inflation rate was announced, agents received the rewards for the past forecasts of this rate and the process repeated itself. Overall, the experiments generated data for 530 model periods based on 5300 individual inflation forecasts.

The results from the high-elasticity treatments can be summarized as follows. In the four baseline treatments the RPE explains agents' actual expectations extremely well. The REE performs significantly worse than the RPE and also performs rather poorly in absolute terms. Therefore, the baseline treatments provide strong support for the RPE and the existence of univariate forecast restrictions.

As a first robustness check agents from two of the baseline treatments were subjected to the a second high-elasticity treatment.

[3] The *REE* also exists under these restrictions because forecast functions in the REE are also univariate.

[4] See Grandmont (1988) for a description of the concept of temporary equilibrium.

This check revealed that the RPE, as formulated in this paper and in Adam (2001), has an important deficiency when it comes to explaining agents' actual expectations: it assumes that agents' 2-step inflation forecasts are derived by iterating forward the (optimal univariate) 1-step forecast model.

Although, this describes agents' actual forecast behavior in most of the baseline treatments, it does not generate an optimal univariate 2-step forecast since the 1-step forecast model is misspecified.[5]

Agents seem to have recognized this fact during the additional high-elasticity treatment and have substituted their iterated 2-step forecasts by a *separate* 2-step forecast model that conditions on a different variable than the 1-step forecast model.

The evidence suggests that agents' new 2-step forecast model is described surprisingly well by the univariate 2-step forecast function that is optimal if the economy is in a RPE. Therefore, univariate forecast restrictions also seem to capture the change in agents' actual forecast behavior.

Obviously, substitution of the 2-step forecast model undermines the RPE. In particular, theory suggests that it should cause a change in the variable used for the 1-step models and ultimately result into convergence to the REE. However, there is no evidence that such substitution in 1-step forecast model is taking place within the 110 model periods that have been generated. Nevertheless, I suspect that more time to learn would eventually lead to such a substitution and finally result into convergence to the REE.

A second robustness check was performed by subjecting agents first to a low-elasticity treatment (where the REE is the only equilibrium) and then to a high-elasticity treatment (where RPE and REE coexist).

In the low-elasticity treatment the REE describes the data reasonably well and far better than in the baseline sessions. The performance of the REE in the second treatment is also improved compared to the baseline treatment but the REE does not outperform the RPE significantly.

All this shows that for a considerable amount of time Restricted Perceptions Equilibria can describe the behavior of dynamic economies better than Rational Expectations Equilibria. Moreover, univariate forecast restrictions seem to be able to capture the restrictions faced by relatively unexperienced forecasters such as the ones participating in economic laboratory experiments.[6] The results also suggest that the failure of the RPE to be stable over longer timer periods is related to the fact that forecast optimality has not been applied to each forecast horizon separately. This should be an essential feature of a RPE.

This paper has been inspired by and is related to the contributions of Marimon and Sunder ((1993), (1994)) who studied experimental overlapping generations economies to select between multiple Rational Expectations Equilibria.

[5] See Bhansali (2002 forthcoming) for a discussion of this point.

[6] Obviously, such restrictions would not apply to specialists familiar with sophisticated econometric forecasting techniques.

These authors studied the relationship between the stability and instability of REE under adaptive learning schemes and the observed laboratory outcomes. The equilibria considered in the present paper are both stable under adaptive learning rules and the focus of this paper is on testing for the existence of forecast restrictions.[7] To my knowledge there exists no experimental results studying the *equilibrium implications* of such forecasting restrictions.

The remaining part of the paper is structured as follows. Section 2 briefly introduces the temporary equilibrium underlying the laboratory economy. The details of the experimental setup are described in detail in section 3 and section 4 derives the Restricted Perceptions Equilibria and Rational Expectations Equilibria for the different experimental treatments. The results of the experiments are analyzed and confronted with theory in section 5. The instructions given to subjects participating in the experiments and technical details can be found in the appendix.

2 The Model and its Equilibria

This section introduces the temporary equilibrium equations underlying the experimental economies.

We consider a two-variable temporary equilibrium equation in output y_t and inflation Π_t where the current values of these variables are given by:

$$\begin{pmatrix} \Pi_t \\ y_t \end{pmatrix} = a_0 + a_1\, y_{t-1} + A \begin{pmatrix} {}_{t-1}\Pi^e_t \\ {}_{t-1}\Pi^e_{t+1} \end{pmatrix} + b\, v_t \tag{1}$$

where ${}_{t-1}\Pi^e_t$ and ${}_{t-1}\Pi^e_{t+1}$ denote the (potentially non-rational) $t-1$ expectations of inflation in period t and $t+1$, respectively, and v_t a white noise demand shock. Equation (1) implies that current inflation and current output are a function of lagged output, expectations of the future inflation rates, and the demand shock.

Adam (2001) derives equation (1) by linearizing a cash-in-advance model with monopolistic competition and a one period price stickiness around its deterministic zero inflation steady state. However, for the purpose of this paper the details of the underlying economic model are not essential. Therefore, readers interested in the underlying model should consult the reference cited above.

The underlying cash-in-advance model implies that the vectors a_0, a_2 and b and the matrix A in equation (1) are given by

$$a_0 = \begin{pmatrix} -\Pi \\ \left(1+\frac{1}{\Pi}\right) y \end{pmatrix} \qquad a_1 = \begin{pmatrix} \frac{1}{y\varepsilon} \\ \frac{1}{\Pi} - \frac{1}{\Pi^2\varepsilon} \end{pmatrix}$$
$$A = \begin{pmatrix} 1-\frac{1}{\Pi\varepsilon} & 1 \\ -\frac{y}{\Pi^2}\left(1-\frac{1}{\Pi\varepsilon}\right) & -\frac{y}{\Pi^2} \end{pmatrix} \quad b = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

[7] Learnabilits of the REE and RPE is shown in Adam (2001)

where $\Pi \approx 1$ and y denote steady state inflation and output, respectively, and $\varepsilon > 0$ is a parameter denoting the real wage elasticity of the labor supply function.

2.1 Agents' Forecast Models and Resulting Equilibria

The temporary equilibrium equation (1) determines current output and inflation as a function of lagged output and expectations of future inflation rates.

We now suppose that agents use univariate models to forecast inflation. With the economy being described by two state variables, output and inflation, this implies that forecasts are either a function of lagged output or lagged inflation, i.e. agents use one of the following models to predict inflation:

$$\begin{aligned}
\text{Model Y}: \quad & \Pi_t = \alpha_y + \beta_y y_{t-1} \\
\text{Model } \Pi: \quad & \Pi_t = \alpha_\Pi + \beta_\Pi \Pi_{t-1}
\end{aligned}$$

While Model Y supposes that output is the driving force of inflation, Model Π supposes that inflation is mainly determined by lagged inflation.

Such a restriction to univariate inflation forecasts can be given several economic interpretations. Firstly, it might simply describe the restriction imposed by the prediction technology available to agents. Secondly, it may be interpreted as the result of an optimal choice of a class of forecasting models that trades off the forecasting performance with the cost of considering smaller or larger classes of forecast models. The restriction is then an artefact of existing calculation costs. Finally, the restriction can be interpreted as a temporary phenomenon due to agents who perform a specification search for suitable forecast models and start out by considering a certain class of models. Unsatisfactory prediction performance may then lead to an enlargement of the class.

Given the above restriction on the available forecast models one can now define a Restricted Perceptions Equilibrium (RPE). Intuitively, a RPE is a situation where forecasts are required to be optimal only in the considered class of forecast models. This differs from the common notion of a Rational Expectations Equilibrium (REE) where forecasts are optimal in the class of all conceivable forecast functions. Formally,

Definition 1 *A Restricted Perceptions Equilibrium (RPE) is a stochastic process for output and inflation generated by equation (1) where agents*

- use least squares to estimate the coefficients (α_y, β_y) and (α_Π, β_Π),
- produce 1-step forecasts using the forecast model that generated the lowest 1-step mean squared forecast error in the past,
- produce 2-step forecasts by iterating the 1-step forecast model, and where

- the parameter estimates (α_y, β_y) and (α_Π, β_Π) are stable over time.

The definition implies that forecast models are chosen on the basis of their 1-step mean squared forecast error. This is justified on the grounds that squared forecast errors represent a quadratic approximation to the correct utility-based choice criterion. 2-step forecasts are then derived by iterating the 1-step model, which is the optimal procedure if agents believe that their 1-step model is an accurate description of the underlying inflation process.[8]

As is easily seen, a REE can also be a RPE if the implied forecast function falls into the class of considered models and if the coefficients of the forecast functions are stable under the least-squares estimation procedure. The latter requirement rules out REE which are unstable under least squares learning; see Evans and Honkapohja (2001) for an extensive treatment.

The definition of an RPE given above is very similar to the one in Evans and Honkapohja (2001) and to Anderson's and Sonnenschein's (1985) *rational expectations equilibria with econometric models*. The main difference is that in the present model not all available regressors can enter the forecast function, as agents are assumed to consider only univariate forecasts.

Suppose agents use Model Y to forecast. Equation (1) then implies that the actual law of motion of inflation is a function of lagged output only.[9] Therefore, a RPE where agents use Model Y must be a REE. As shown in Adam (2001), there is a unique REE for which the parameters of the forecast functions are stable under least-squares learning. In this REE inflation depends on lagged output with output itself being a white noise process:

$$\Pi_t = \frac{\Pi}{y} y_{t-1} \tag{2a}$$

$$y_t = y + v_t \tag{2b}$$

Equilibrium expectations are then given by

$$_{t-1}\Pi_t^e = \frac{\Pi}{y} y_{t-1} \tag{3a}$$

$$_{t-1}\Pi_{t+1}^e = \Pi \tag{3b}$$

Thus, the restriction to univariate forecasts *per se* does not rule out that the economy is in a REE.

[8] While forward iteration for Model Π is straightforward and leads to $_{t-1}\Pi_{t+1}^e = \alpha_\Pi + \beta_\Pi(\alpha_\Pi + \beta_\pi \Pi_{t-1})$ iteration based on Model Y requires a forceast of y_t, since $_{t-1}\Pi_{t+1}^e = \alpha_y + \beta_y \, _{t-1}y_t^e$. Taking $t-1$ expectations of a (linearized) accounting identity of the model delivers $_{t-1}y_t^e = y + \frac{1}{\Pi} y_{t-1} - \frac{y}{\Pi^2} \, _{t-1}\Pi_t^e + v_t$. Subsituting this into the previous equation and using the Model Y value for $_{t-1}\Pi_t^e$ delivers the iterated forecast.

[9] See the previous footnote for how to express $_{t-1}\Pi_{t+1}^e$ as a function of y_{t-1}.

To simplify terminology I will refer to equilibrium (2) as the model's REE and reserve the term RPE to equilibria where the constraint to univariate forecast functions is strictly binding.

Now suppose agents use Model Π to forecast inflation. Equation (1) then implies that inflation depends on lagged inflation *and* on lagged output. As a result, both forecast models will be misspecified, which generates the possibility that Model Π delivers superior predictions.

Substituting the expectations generated by Model Π

$$ {}_{t-1}\Pi_t^e = \alpha_\Pi + \beta_\Pi \Pi_{t-1} \tag{4a} $$

$$ {}_{t-1}\Pi_{t+1}^e = \alpha_\Pi + \alpha_\Pi \beta_\Pi + \beta_\Pi^2 \Pi_{t-1} \tag{4b} $$

into equation (1) delivers

$$ \begin{aligned} \begin{pmatrix} \Pi_t \\ y_t \end{pmatrix} &= \begin{pmatrix} \alpha_\Pi(2+\beta_\Pi - \frac{1}{\Pi\varepsilon}) - \Pi \\ y(1+\frac{1}{\Pi}) - \frac{y}{\Pi^2}\alpha_\Pi(2+\beta_\Pi - \frac{1}{\Pi\varepsilon}) \end{pmatrix} \\ &+ \begin{pmatrix} \beta_\Pi(1-\frac{1}{\Pi\varepsilon}+\beta_\Pi) & \frac{1}{y\varepsilon} \\ -\frac{y}{\Pi^2}\beta_\Pi(1-\frac{1}{\Pi\varepsilon}+\beta_\Pi) & \frac{1}{\Pi} - \frac{1}{\Pi^2\varepsilon} \end{pmatrix} \begin{pmatrix} \Pi_{t-1} \\ y_{t-1} \end{pmatrix} + \begin{pmatrix} 0 \\ v_t \end{pmatrix} \end{aligned} \tag{5} $$

Since inflation is a function of lagged output and lagged inflation, a situation where agents use Model Π cannot be a REE but must be a RPE.

In the RPE, (α_Π, β_Π) are given by the least squares estimates obtained from fitting Model Π to process (5), which is itself a function of (α_Π, β_Π). Therefore, determining the RPE involves solving a fixed point problem as it is the case when determining the REE. Appendix 7.1 shows how the fixed point values $(\alpha_\Pi^*, \beta_\Pi^*)$ can be determined.

Process (5) with (α_Π, β_Π) given by the fixed point values is a RPE whenever forecast Model Π delivers a better forecast than Model Y and when the fixed point is stable under least-squares learning. As shown in Adam (2001), this is the case if the elasticity of labor supply ε is larger than 1.75.

Importantly, the RPE described above does not emerge if agents can handle forecast models with two (or more) regressors. Equation (1) implies that the actual law of motion of the economy will not be more complicated than the admitted forecast models. As a result, a rest point of the least squares estimation process must be a REE. Thus, all alternative hypotheses that consider larger classes of forecast models lead to the same equilibrium prediction, namely that output and inflation are described by the REE. This causes the present model to be particularly suitable to test for univariate restrictions on agents' forecast functions.

3 Experiment Setup and Implementation

3.1 Experiment Setup

Six experimental sessions were conducted during 6 days. Five subjects participated in each session with no subject taking part in more than one session. Experiments took place at the University of Salerno, Italy and at the University of Frankfurt, Germany. Most subjects were undergraduate business and engineering majors and only one of them had an economics major.

There were two kinds of experimental treatments: **low-elasticity** treatments where the elasticity of labor supply was set to $\varepsilon = 1$, and **high-elasticity** treatments where the elasticity of labor supply was set to $\varepsilon = 2$.

In low-elasticity economies the REE is the only equilibrium outcome, even if agents restrict consideration to univariate forecast models.[10] In high-elasticity economies an RPE coexists with the REE if agents consider only univariate forecast models; if agents consider also forecast functions with more variables then the REE is again the unique equilibrium.

Since the other model parameters do not affect the existence of the various equilibria, their values were kept constant across all sessions and treatments: the steady state inflation rate was set equal to 4%, the steady state output was equal to 100, and demand shocks v_t were independently drawn from a uniform distribution with support on $[-1, +1]$.

Table 1 lists details of the sessions and treatments. The first treatment of Sessions 1 to 4 constitute the baseline case for assessing how well the RPE and REE explain agent's inflation forecasts.

Subjects participating in Sessions 3 and 4 experienced a second high-elasticity treatment to check for the stability of the results obtained in the baseline case.

Subjects participating in Sessions 5 and 6 were first exposed to a low-elasticity treatment where only an REE exists and then to a high-elasticity treatment where REE and RPE coexist again.

High-elasticity economies lasted for 55 periods and low-elasticity for 45 periods. This implies that overall 530 model periods were generated.[11] With 5 subjects participating in every session and each subject making 2 forecasts per period, 5300 individual forecasts were collected.

[10] If agents' used Model Π in a low-elasticity treatment, the resulting inflation process (5) could be better forecasted using Model Y.

[11] The average length of a treatment was close to 2 hours, which made it unwise to choose a much higher number of periods.

3.2 Experiment Implementation

Experiments have been implemented using *MacroLab*, which is an experimental software package designed to analyze dynamic interactive decision settings as they appear in macroeconomics. Adam and Marimon (2001) provide a description of the software, which can be freely downloaded from the Internet. There is also an experiment database that can be accessed and that allows to replicate the experiments of this paper (possibly involving different parameterizations).[12]

At the beginning of the experimental session subjects received written instructions. These are reproduced in appendix 7.3. To introduce subjects to the *MacroLab* software a trial session lasting for a few periods was started. Subjects were told that they could not learn anything from the trials apart from how to handle the software.

The labor supply elasticity in trial sessions was set to $\varepsilon = 0.5$. For this value only a REE exists even if univariate forecasting constraints are present. Therefore, trials would have biased results in favor of the REE, although this seems unlikely.

At the start of each treatment agents observed a single data point for output and inflation. Subjects did neither know the steady state values of output and inflation nor the elasticity value of the labor supply or any other feature of the underlying economy. All they were told was that they had to forecast inflation and that they would observe a longer and longer history of output levels and inflation rates as time proceeds. The length of the treatment was also unknown to the subjects.[13]

In each model period t subjects were asked to forecast the inflation rate for periods $t+1$ and $t+2$. With these forecasts a new output level and inflation rate was calculated for $t+1$ using equation (1) and by substituting expectations by the average of the individual inflation forecasts.[14] Averaging of forecasts is justified on the grounds that it represents a first-order approximation to the exact (non-linear) aggregation of heterogeneous expectations.

The new output level and inflation rate were then announced to agents who received points for their past forecasts of the newly announced inflation rate. The number of points received depended on the absolute value of the forecast error. Points for each forecast were calculating according to the following formula

$$\frac{400}{1+f} - 100$$

[12] See www.wiwi.uni-frankfurt.de/professoren/wieland/adam.

[13] With a small positive probability that the economy continues for another period it is a strictly dominant strategy for agents to report their best forecasts in each period.

[14] Obviously, forecast were averaged separately for each forecast horizon.

where f denotes the absolute forecast error (expressed in percentage terms). Subjects received a maximum of 300 points per forecast and could lose up to 100 points depending on forecast accuracy.

A strictly positively sloped reward schedule was chosen because it provides an incentive to improve the precision of forecasts at each average level of forecast precision. The convexity of the schedule avoids punishments in the form of a large amount of negative points, which is not credible. The convexity also implies that it is particularly important to get "the last bit" of the forecast right because the reward schedule is steepest at $f = 0$. [15] This is important if it is mentally more demanding to follow small fluctuations in the data.

The points received during the treatments were added up and converted into cash payments according to a conversion rate announced at the end of the session. Conversion rates were calculated to make an average payment of 30 Euros for subjects participating in single treatment sessions and 50 Euro (60 Euros) for subjects participating in double treatment sessions in Italy (Germany). The average hourly payment per session was always larger than 8 Euros (12 Euros) in Italy (Germany) and went up 15 Euros in one session. Given the alternative wages available to undergraduate students, pecuniary incentives should have motivated subjects sufficiently well.

4 REE and RPE in the High and Low Elasticity Treatments

This section briefly determines output and inflation behavior and equilibrium expectations for the REE and RPE in the high-elasticity and low-elasticity treatments.

The REE exists in both treatments. Moreover, output and inflation dynamics are independent of the elasticity of labor supply. Given the parameterization described in section 3.1, output and inflation in the experimental REE are given by

$$\begin{aligned} y_t &= 100 + v_t \\ \Pi_t &= .0104 \cdot y_{t-1} \end{aligned}$$

Consequently, REE inflation expectations are given by

$$ {}_t\Pi_{t+1}^{REE} = 0.0104 \cdot y_t \tag{6a}$$

$$ {}_t\Pi_{t+2}^{REE} = 1.04 \tag{6b}$$

[15] This was the main reason why I did not choose a quadratic loss function as suggested by the underlying theory.

Next, consider the RPE which exists only in high elasticity treatments. For the chosen parameterization, the fixed point values of (α_Π, β_Π) are given by:[16]

$$\alpha_\Pi^* \approx .248842$$
$$\beta_\Pi^* \approx .760728$$

Equations (5) then imply that in a RPE output and inflation are described by

$$y_t \approx 143.699 + .499260 \cdot y_{t-1} - 90.0242 \cdot \Pi_{t-1} + v_t$$
$$\Pi_t \approx -.472649 + 0.005000 \cdot y_{t-1} + .973701 \cdot \Pi_{t-1}$$

Likewise, equations (4a) and (4b) imply that RPE inflation expectations are given by

$$_t\Pi_{t+1}^{RPE} = .248842 + .760728 \cdot \Pi_t \tag{7a}$$
$$_t\Pi_{t+2}^{RPE} = .438143 + .578708 \cdot \Pi_t \tag{7b}$$

5 Experimental Results

This section analyzes the data generated in the experimental sessions. The yardstick chosen for assessing how well the REE and RPE explain the experimental data is the ability of these equilibria to match the representative agents' forecast function, i.e. the average forecast function of the participating subjects.

Matching average forecasts is justified on the grounds that it is average forecasts that drive output and inflation dynamics in the economy (see section 3). Therefore, the equilibrium notion capturing the representative agent's forecast function also captures the behavior of output and inflation.

5.1 The Unconditional Inflation Forecasts

I first consider the average inflation forecasts. Average forecasts can be interpreted as an estimate of agents' unconditional inflation forecasts.

REE and RPE both predict unconditional 1-step and 2-step forecasts to be equal to the steady state inflation rate, which was set to 4% in all treatments. Unconditional forecasts, therefore, do not allow to discriminate between the REE and RPE. Nevertheless, it is important to analyze unconditional forecasts because a failure of the REE and RPE to explain the first moment of expectations would constitute a failure of first order importance.

Table 3 lists the actual average 1-step and 2-step inflation forecasts (the average is taken across periods and agents). Actual average forecasts are rather close to the predicted value. Of the 20 values reported, 16 are within 2 standard deviations and 19 within 3 standard deviations of the steady state value. The

[16] See appendix 7.1 for how to calculate these.

value for the 2-step forecast in Treatment 1 of Session 6, which is more than 3 standard deviations lower than the predicted value, is driven by a large negative forecast that one subject entered in period 12 of the experiment.

Given the results of Table 3, first moments of the experimental data seem consistent with the REE and RPE predictions. The subsequent sections will consider the ability of the REE and RPE to explain agents' conditional inflation forecasts. Since conditional forecasts differ across equilibria, this will allow to discriminate between the two competing equilibrium explanations.

5.2 The Baseline High-Elasticity Treatments

This section considers the results of the baseline treatments, i.e. the first treatment of Sessions 1 to 4. In these treatments an REE and a RPE coexist if agents use univariate forecast functions.

Throughout the remaining part of the paper the main strategy for assessing how well the REE and RPE explain the experimental data is to report OLS-estimates of the parameter β for the following regression

$$_t\Pi_{t+i}^{actual} = \alpha + \beta \cdot {}_t\Pi_{t+i}^{RPE} + (1-\beta) \cdot {}_t\Pi_{t+i}^{REE}$$

where $_t\Pi_{t+i}^{actual}$ denotes the actual time t forecast of the $t+i$ inflation rate $(i = 1, 2)$ and $_t\Pi_{t+i}^{REE}$ and $_t\Pi_{t+i}^{RPE}$ the corresponding equilibrium forecasts in a REE and RPE, respectively, as given by equations (6) and (7). The estimate of β can be interpreted as the share of agents using the RPE-forecasts. An estimate of β close to 1 indicates that the RPE explains the forecast functions well, while a value close to zero indicates that the REE offers a superior description of the forecast function.[17]

Figures 1 to 4 show the actual forecasts, the REE-forecasts, and the RPE-forecasts for the four baseline treatments.

The figures show that the RPE-forecasts track actual forecasts extremely well while the REE-forecasts perform rather poorly. This is true for all 1-step forecasts and the 2-step forecasts in Sessions 2 and 4. For the 2-step forecasts of Sessions 1 and 3 this seems to hold only for about the first half of the treatment; in the second half the 2-step REE and RPE-forecasts seem to perform equally bad.

The evidence shown in Figures 1 to 4 provide strong support in favor of the RPE. The close fit between RPE-forecasts and actual forecasts is remarkable since the RPE-forecast is calculated *without* reference to agents' actual forecasts, i.e. the two curves have not been fitted to each other! The equilibrium forecasts shown in these figures rely only on information that is available to agents at the

[17]Sometimes this interpretation is not appropriate, e.g. because the R^2 is very low or even negative. When this is the case it will be explicitely mentioned.

time they formulate their forecasts and then assumes that the economy is in an REE or RPE, respectively.

The visual impression from Figures 1 to 4 is confirmed by a more formal analysis. The first panel of Table 3 reports the estimated share of RPE-forecasters β for the 1-step inflation forecasts. Estimates are reported for the entire treatment and the last 20 periods to assess whether there is some variation over time due to learning processes taking place.

The point estimates in the upper panel of Table 3 are relatively close to 1 and imply that in each treatment more than 85% of agents use RPE-forecasts. The shares are estimated rather precisely and there are only weak signs (in Session 1 and 3) that they are significantly lower in the last 20 periods of the treatments.

The second panel of Table 3 reports the share of RPE-forecasters for the 2-step forecasts. Here the situation is different across the four sessions. This should hardly be surprising given the evidence provided in the lower panels of Figures 1 to 4. The RPE-forecasts clearly dominate in Sessions 2 and 4. Also, in these sessions the dominance of RPE-forecasts appears to be stable over time. RPE-forecasts also perform well in the first part of Sessions 1 and 3. Yet, the estimates for the last 20 periods suggest that towards the end of treatment the REE-forecasts dominate in these sessions.

Overall, the baseline treatments provide overwhelming evidence in favor of the RPE. In none of the session does the REE offer a good description of agents' inflation expectations.

Figure 5 provides additional support for this claim by depicting actual inflation rates together with the rates forecasted by agents one or two periods before. Due to space constraints, only evidence for Session 4 is reported but the graphs for the other sessions look very similar.

If forecasts were rational, then the difference between the actual and forecasted inflation series would be white noise processes. However, as is easy to spot in Figure 5, inflation forecasts lag actual inflation. While 1-step forecasts seem to lag by one period, 2-step forecasts seem to lag by 2 periods.[18] As a result, forecast errors are strongly positively auto-correlated, a feature that is consistent with the RPE but not with the REE.

The next subsection analyzes the data from Sessions 1 and 3 in greater detail since these sessions seemed to offer evidence of an improved performance of the REE 2-step forecasts.

[18] The figure depicts the actual inflation rate and past forecasts of this inflation rate at the same point of the x-axis, so this feature is not due to a problem of representations.

5.2.1 Baseline Treatments: What Happened in Sessions 1 and 3 ?

Results from Table 3 suggest that the REE 2-step forecasts start to perform better than the RPE forecasts towards the end of Sessions 1 and 3. Yet, Figures 1 and 3 clearly reveal that actual forecasts still display regular cyclical patterns with relatively large amplitudes. Obviously, these cannot be captured by the REE 2-step forecasts.

This suggests that the low β estimate for the end of Sessions 1 and 3 is due to a deterioration of the fit of the RPE forecast rather than to an improved fit of the REE forecasts.

Figure 6 depicts actual 2-step forecasts from Sessions 1 and 3 together with the following output-based forecast function:

$$ {}_t\Pi^e_{t+2} \approx 0.4172 + 0.0062 y_t \tag{8} $$

Towards the end of the considered sessions, agents' 2-step forecasts seem to be captured rather well by equation (8). Moreover, forecasts (8) start to perform well precisely when the performance of the RPE-forecasts starts to deteriorate, see Figures (1) and (3).

This suggests that agents participating in Sessions 1 and 3 have substituted their RPE 2-step forecast function (7b) with the output-based forecast function (8).

Forecast function (8) is the optimal (in a mean squared error sense) univariate 2-step inflation forecast function for an economy that is in a RPE, as the data suggest to be the case for the first half of the considered sessions.

It might come as a surprise that the optimal 2-step forecast function differs from the RPE 2-step forecast function (7b). This suboptimality (even in the class of univariate forecast functions!) arises because 2-step forecasts have been assumed to be obtained by iterating forward the (optimal univariate) 1-step forecast equation.

Although such an iteration is the standard procedure in econometrics to derive a multi-step prediction from a linear econometric model, it is suboptimal here since the 1-step forecast function is misspecified, see Bhansali (2002 forthcoming) for details.

As a result, a superior univariate 2-step forecast can be derived by regressing inflation directly on twice lagged output or twice lagged inflation. Doing so one finds that function (8) delivers the univariate prediction with the lowest mean-squared forecast error in the RPE.

The somewhat informal discussion above is support by a more formal analysis. The lower panel of Table 3 lists the OLS-estimate of β obtained from the regression

$$ {}_t\Pi^{actual}_{t+2} = \alpha + \beta \cdot {}_t\Pi^{RPE}_{t+2} + (1-\beta) \cdot {}_t\Pi^{Output}_{t+2} $$

where ${}_t\Pi^{Output}_{t+2}$ denotes the forecast given in equation (8).

While Sessions 2 and 4 do not show any signs of an increased share of output-based forecasters, there is a significant increase in the share of output-based 2-step forecasters in Sessions 1 and 3: Point estimates imply that the large majority of agents used output-based forecasts during the last 20 periods of the treatment.

The previous findings suggest that agents' forecasts still seem to be described by optimal univariate forecast models. However, agents seem to have become aware that their simple forecast models are misspecified and have started to use different models for different forecast horizons.

5.3 Additional High-Elasticity Treatments

To check for the stability of the results, subjects participating in Sessions 3 and 4 were subjected to a second high-elasticity treatment, which is analyzed in this section.

Given that in some sessions agents switched to output-based 2-step forecasts, one has to ask for the potential rest points of a learning process where agents condition 1-step forecasts on inflation and 2-step forecasts on output. Such a rest point could be expected to emerge in these additional treatments.

In appendix 7.2 it is shown that there exists a unique stationary rest point where agents use an optimally parameterized inflation-based 1-step forecast model and an optimally parameterized output-based 2-step forecast model.[19] This rest point will be referred to as the *mixed-forecast situation* subsequently.

As shown in appendix 7.2, optimally parameterized forecast functions for the mixed-forecast situation are given by

$$ {}_t\Pi^e_{t+1} \approx 0.6887 + 0.3378 \cdot \Pi_t \tag{9a} $$
$$ {}_t\Pi^e_{t+2} \approx 0.7373 + 0.003027 \cdot y_t \tag{9b} $$

It is important to note that equations (9) do *not* describe an equilibrium situation where agents use optimal univariate forecast functions for each forecast horizon. If agents used equations (9) to forecast, a univariate output-based prediction for the 1-step forecast would dominate the inflation-based forecast function shown above. Thus, the mixed forecast situation is only a rest-point when *taking as given* the variables that enter the respective forecast functions.

Despite this suboptimality, one would expect that equations (9) at least initially describe actual forecast behavior. Once agents substitute the 1-step forecast function (9a) by an output-based forecast function, one would expect

[19] Optimality is again defined in terms of mean-squared errors.

the emergence of the REE.[20] For this reason I let the mixed-forecast situation compete against the REE in the subsequent analysis.

Note that the coefficient on lagged inflation in equation (9a) has the same sign as in the RPE. This might explain why agents' actual 1-step forecasts remain to be captured rather well by the RPE 1-step forecasts in all of the baseline sessions even though 2-step RPE-forecasts have been replaced by output-based forecasts in some of these sessions.

Furthermore, the coefficient on lagged output in equation (9b) has the same sign as in equation (8). This suggests that once agents switch to an output-based forecast function the learning process will only lead to a decrease in the reaction coefficient on the output term but not to a change of sign. Therefore, the learning process will not cause major problems to the analysis.

Figures 7 and 8 graph actual forecasts, REE forecasts, and mixed forecasts for the second treatments of Sessions 3 and 4. Interpretation of the data from Session 3 is somewhat difficult because one subject experimented with large negative inflation forecasts in period 16-22 to learn about the economy's reaction to these forecasts.[21] Experimentation during these periods caused output levels to increase, which caused output-based forecasts to be off track for some time. However, this should not be interpreted as a genuine failure of output-based forecasts.

The figures suggest that in both sessions 1-step forecasts are still captured far better by the (inflation-based) mixed-forecast than by the (output-based) REE-forecast. This is confirmed by the quantitative evidence presented in Table 4. The share of agents using the (inflation-based) mixed forecast function is estimated to be close to one and is not significantly lower in the last 20 periods of the treatments. Thus, regarding 1-step forecasts there is no evidence in favor of a convergence process towards the REE.[22] At the same time the mixed-forecast situation captures the forecasts rather well.

Figures 7 and 8 also suggests that 2-step forecasts seem to be more in line with the mixed-forecasts than with the REE-forecasts, in particular towards the end of the treatments.

The visual impression is confirmed by the quantitative results reported in the second and third panel of Table 4. The second panel shows that the mixed-forecast function dominates in Session 4 and seems to gain weight in Session

[20] Recall that the REE is the unique stationary equilibrium where the parameters in agents' output-based forecast functions are optimal given the process of output and inflation that they generate.

[21] The subject mentioned to me that he experimented after the end of the experiment. He also mentioned that he had abandoned experimentation after a while as it became to costly and did not generate a lot of information.

[22] There is some evidence that actual forecasts initially fluctuate more than the mixed 1-step forecasts (9a) suggest. This is likely to be the case because agents still use the RPE-forecast, which has a larger coefficient on the lagged inflation term.

3, where interpretation is hampered by the fact that one subject experimented with large negative forecasts. Furthermore, the third panel shows that almost all agents seem to use output-based forecast rules. Subjects from Session 4, who used inflation-based 2-step forecasts throughout the first treatment, now also seem to use output-based 2-step forecasts.

Overall, the additional high-elasticity treatments strongly suggest that agents' forecast functions moved into the direction of the mixed forecast situation (9). The mixed forecast situation offers a better prediction of agents' expectations than the REE, especially towards the end of the treatments. Thus, after more than 110 model periods the REE does not yet emerge as the dominant explanation of the data. Of course, this does not exclude that additional treatments would eventually cause it to become the dominant explanation. Given the logic according to which agents seem to substitute their forecast functions, the mixed forecast situation can be expected to be transient.

To assess whether there are situations in which the REE offers a good explanation of the data early on, the next section considers a parameterization of the economy where an RPE does not exist.

5.4 The Low-High Treatment Combination

This section discusses the results from Sessions 5 and 6. In these sessions subjects first experienced a low-elasticity treatment before being subjected to a high-elasticity treatment.

In low-elasticity treatments an RPE does not exist even when agents restrict attention to univariate forecast functions. Therefore, one would expect the REE to be able to explain the data better than in the baseline sessions. Moreover, if agents learned the REE in the first treatment, their experience should facilitate coordination on the REE in the subsequent high-elasticity treatment.

5.4.1 Low-Elasticity Treatments

Figure 9 depicts actual forecasts and REE-forecasts for the low-elasticity treatment of Session 5.

While in the first 25 periods the REE 1-step forecasts tend to peak 1 period before actual forecasts, the REE and actual forecasts move in a highly synchronized fashion towards the end of the treatment when the fit between the two forecasts seems to be much better than after the 110 periods of high-elasticity treatment applied in Sessions 3 and 4.

Actual 2-step forecasts also seem to be roughly in line with the REE prediction. Although these forecasts still display some cyclical variation, the standard deviation of 2-step forecasts is now rather small.

Figure 10 displays information for the low-elasticity treatment of Session 6. Interpretation of the data is complicated by the fact that in period 12 one subject entered a 1-step inflation forecast of -28%.[23] This caused inflation to be very low and output to be very high. The high output level subsequently caused a strong rise in inflation. This strong cyclical 'swing' in the data might well have influenced agents' forecast functions.

Correspondingly, the signs in favor of the REE are much weaker for this session. 1-step forecasts peak well before actual forecasts even towards the end of the treatment. Also, 2-step forecasts do not move as closely to the steady state value as in Session 5 although variability seems to decrease over time.

The visual impression from Figures 9 and 10 is confirmed by a more formal analysis. Since an RPE does not exist in low-elasticity treatments one has to compare the ability of REE-forecasts to explain actual forecasts with the ability to do so in high elasticity treatments.

Table 5 reports results of such a comparison for the 1-step forecasts. For all sessions and treatments considered thus far the table reports the coefficient β obtained from estimating equation

$$ {}_t\Pi_{t+1}^{actual} = \alpha + \beta \cdot {}_t\Pi_{t+1}^{REE} \tag{10} $$

using ordinary least squares.

While for the low-elasticity treatments β is significant and positive, it is either insignificant or negative for the high-elasticity treatments. This together with the fact that β is very high in the last 20 periods of Session 5 suggests that in low-elasticity treatments the REE-forecasts offers a much better explanation of actual 1-step forecasts than in high-elasticity treatments.

Table 6 reports evidence on 2-step forecasts. The table presents results from regressing the squared deviation of actual 2-step forecasts from REE 2-step forecasts on a constant.

For Session 5 the estimated constant is significantly lower than in all other high elasticity sessions. This holds true for the whole sample and the last 20 periods of the treatment. For Session 6 the picture is somewhat mixed, which is most likely due to the large negative forecast mentioned above. Nevertheless, in the last 20 periods of Session 6 the squared deviation is still significantly lower than in all but one high-elasticity treatment.

The previous evidence shows that in low-elasticity treatments the REE performs significantly better than in high-elasticity sessions. Especially evidence from Session 5 seems largely consistent with the REE towards the end of the treatment. This has not been the case in any of the high-elasticity treatments.

The next section considers whether the experience of the low-elasticity treatment facilitates coordination on the REE in the subsequent high-elasticity treatment.

[23] It is not entirely clear whether the forecast was an attempt to obtain information about the economy's reaction to such forecasts or whether this was a simple omission of the decimal point.

5.4.2 High-Elasticity Treatments

Figures 11 and 12 depict actual forecasts, REE-forecasts, and RPE-forecasts for the high-elasticity treatment of Sessions 5 and 6.

The match between RPE 1-step forecasts and actual forecasts is much weaker than in the baseline high-elasticity treatments. Indeed, actual 1-step forecasts seem to be in between the REE-forecasts and the RPE-forecasts. 2-step forecasts display some cyclical variation but they are not matched particularly well by RPE-forecasts. Also, deviations from the REE-forecasts are rather small (Session 5) or do become so over time (Session 6).

Table 7 presents estimates of the share of REE and RPE forecasters. The estimated share of RPE 1-step forecasters is still relatively high. Yet, it is significantly lower than in Sessions 1 to 4. The same holds for the 2-step forecasts where the data now favors REE-forecasts over RPE-forecasts.[24]

The first panel of Table 8 reports the value of β obtained from estimating equation (10) via least squares. The REE 1-step forecasts obtain a significantly positive coefficient for the whole of Session 5 and also for the last 20 periods of Session 6.

Combining this with evidence from Table 5 suggests that the present high-elasticity treatments look more like the previous low-elasticity treatment and much less like the high-elasticity treatments of Sessions 1 to 4.

The second panel of Table 8 shows evidence on 2-step forecasts. Reported are the results from regressing the squared difference of actual forecasts and REE-forecasts on a constant. These squared forecast errors are found to be very small and significantly lower than in any of the high-elasticity treatments of Sessions 1 to 4 (see Table 6). Moreover, they are of about the same order as in the low-elasticity treatment of Session 5, which was the session best described by an REE.

All this suggests that agents have carried over some of their experience from the low-elasticity treatment. However, neither the REE nor the RPE emerges as the dominant explanation of agents' expectations.

6 Conclusions

The main results obtained from the experimental sessions of this paper can be summarized as follows.

[24] There is no evidence that output-based forecasts would dominate the RPE forecasts, as was the case in the first treatment of Sessions 1 to 4 when the share of RPE-forecasters was found to be low. These results are not reported in the table. The estimated share of RPE-forecasters vs. Ouput based forecasters is close to 0.5 in Session 5. In Session 6 it is somewhat higher but produces negative R^2 values which indicates that neither of the forecasts offers a good explanation.

Firstly, in the baseline treatments the Restricted Perceptions Equilibrium outperforms the Rational Expectations Equilibrium as the dominant explanation of the experimental data. This suggests that the forecasting technology employed by relatively unsophisticated forecasters is accurately captured by univariate forecast models.

Secondly, although agents use a simple forecasting technology, they seem to be aware that their forecast models are possibly misspecified. Therefore, after having gained experience with their environment agents start to use different forecast models for different forecast horizons. The new forecast models seem again to be simple univariate models.

Finally, for a parameterization where a Restricted Perceptions Equilibrium does not exist, the Rational Expectations Equilibrium offers a good description of the experimental data. There is also tentative evidence that experience from such treatments facilitates coordination on the Rational Expectations Equilibrium for parameterizations where a Restricted Perceptions Equilibrium coexists with the Rational Expectations Equilibrium.

There is considerable room left for future research. To assess the robustness of the present results it is of interest to learn whether Restricted Perceptions Equilibria with univariate models also describe expectations in other experimental economies. Furthermore, the results obtained so far suggest to consider an experimental economy possessing a Restricted Perceptions Equilibrium with univariate forecast functions where forecasts are optimal at each forecast horizon. Consideration of such an economy would allow to assess whether Restricted Perceptions Equilibria could be truly stable over time instead of being just transitory phenomena of an economy ultimately converging to a Rational Expectations Equilibrium.

7 Appendix

7.1 Calculating (α_Π, β_Π) in Model Π Equilibrium

In a stationary equilibrium the least squares estimates of Model Π are given by

$$\beta_\Pi = \frac{cov(\Pi_t, \Pi_{t-1})}{var(\Pi_{t-1})} \tag{11}$$

$$\alpha_\Pi = \Pi(1 - \beta_\Pi) \tag{12}$$

where Π is the steady state inflation rate and where Π_t evolves according to (5). Let B denote the AR-matrix in (5) and vec be the column-wise vectorization operator. Then taking variances on both sides of (5) and assuming stationarity implies that

$$vec(\Sigma) = (I - B \otimes B)^{-1} vec(\Omega) \tag{13}$$

where $\Omega = \begin{pmatrix} 0 & 0 \\ 0 & \sigma_v^2 \end{pmatrix}$. Equation (13) delivers the expression for the denominator in (11). From

$$\Gamma = \begin{pmatrix} cov(\Pi_t, \Pi_{t-1}) & cov(\Pi_t, y_{t-1}) \\ cov(y_t, \Pi_{t-1}) & cov(y_t, y_{t-1}) \end{pmatrix} = B\Sigma$$

one obtains the expression for the numerator in (11). Equation (11) thus implies that β_Π solves

$$\beta_\Pi = \frac{\beta_\Pi(1 - \frac{1}{\Pi\varepsilon} + \beta_\Pi) + \frac{1}{\Pi} - \frac{1}{\Pi^2\varepsilon}}{1 + \beta_\Pi(1 - \frac{1}{\Pi\varepsilon} + \beta_\Pi)\left(\frac{1}{\Pi} - \frac{1}{\Pi^2\varepsilon}\right) + \frac{1}{\Pi^2}\beta_\Pi(1 - \frac{1}{\Pi\varepsilon} + \beta_\Pi)\frac{1}{\varepsilon}} \tag{14}$$

Solving this equation (e.g. numerically), one obtains β_Π. The value for α_Π can then be obtained using equation (12).

7.2 The Mixed Forecast Situation

Suppose agents use the following 1-step and 2-step forecast model:

$$\Pi_t = \alpha + \beta\Pi_{t-1} \tag{15a}$$

$$\Pi_{t+1} = \gamma + \delta y_{t-1} \tag{15b}$$

where in a stationary equilibrium

$$\beta = \frac{cov(\Pi_t, \Pi_{t-1})}{var(\Pi_{t-1})} \tag{16}$$

$$\delta = \frac{cov(\Pi_t, y_{t-2})}{var(y_{t-2})} \tag{17}$$

and

$$
\begin{aligned}
\alpha &= (1-\beta)\Pi \\
\gamma &= \Pi - \delta y
\end{aligned}
$$

where variables without subscript denote steady state values.

To calculate the variances and covariances insert the expectations derived from (15) into (**??**). This delivers

$$
\begin{pmatrix} \Pi_t \\ y_t \end{pmatrix} = a + B \begin{pmatrix} \Pi_{t-1} \\ y_{t-1} \end{pmatrix} + \begin{pmatrix} 0 \\ v_t \end{pmatrix} \tag{18}
$$

where

$$
B = \begin{pmatrix} (1-\frac{1}{\Pi\varepsilon})\beta & \delta + \frac{1}{y\varepsilon} \\ -\frac{y}{\Pi^2}(1-\frac{1}{\Pi\varepsilon})\beta & -\frac{y}{\Pi^2}\delta + (\frac{1}{\Pi} - \frac{1}{\Pi^2\varepsilon}) \end{pmatrix}
$$

Since the constant a in (18) does not influence the covariances one can ignore it from now on. Taking variances on both sides of (18) delivers

$$
vec(\Sigma) = vec\left(VAR\begin{pmatrix} \Pi_t \\ y_t \end{pmatrix}\right) = (I - B \otimes B)^{-1} vec(\Omega)
$$

where $vec(\cdot)$ denotes the vectorization operator and Ω is variance covariance matrix of the shocks, i.e.

$$
\Omega = \begin{pmatrix} 0 & 0 \\ 0 & \sigma_v^2 \end{pmatrix}
$$

Multiplying (18) by the once and twice lagged (y_t, Π_t) row-vector and taking expectations delivers

$$
\begin{aligned}
\Gamma_1 &= B\Sigma \\
\Gamma_2 &= B^2\Gamma_1
\end{aligned}
$$

where Γ_i is the covariance matrix of output and inflation with i-times lagged output and inflation. Taking the variances and covariances from these expressions delivers

$$
\beta = \frac{cov(\Pi_t, \Pi_{t-1})}{var(\Pi_{t-1})} = \frac{-1 - \beta\Pi(1-\Pi\varepsilon) + \varepsilon(\Pi - \delta y)}{\varepsilon\Pi^2 - \beta(1-\Pi\varepsilon)}
$$

This equation for has two fixed points given by

$$
\begin{aligned}
\beta_1 &= \frac{-\Pi - \sqrt{\Pi^2 - 4(-1+\varepsilon\Pi)(1-\varepsilon\Pi+\delta\varepsilon y)}}{2(-1+\varepsilon\Pi)} \\
\beta_2 &= \frac{-\Pi + \sqrt{\Pi^2 - 4(-1+\varepsilon\Pi)(1-\varepsilon\Pi+\delta\varepsilon y)}}{2(-1+\varepsilon\Pi)}
\end{aligned}
$$

Substituting β_1 into (17) using the expressions for the covariances derived above and solving for δ with $\Pi = 1.04$ and $\varepsilon = 2$ delivers two real and two imaginary solutions for δ. However, both real solutions imply values for β_1 smaller than -1 which would contradict stationarity of the inflation rate.

Substituting β_2 into (17) and solving for δ with $\Pi = 1.04$ and $\varepsilon = 2$ delivers three real solutions for δ. Only one of these solutions implies a value of $|\beta_2| < 1$. This is the solution shown in equation (9).

7.3 Instructions for Subjects

General

Today you will participate in an experiment of economic decision making. Various research foundations have provided funds for the conduct of this research. Instructions are simple and if you follow them carefully you can earn a considerable amount of money. The average payment will be around 60.000 Lire but, depending on how well you do, you may well earn up to 120.000 Lire.

You are assigned the role of a private agent whose task is to forecast the rate of inflation in the economy. In each experimental period t, you are asked to forecast the inflation rate for the next two periods, i.e. the inflation rate for period t+1 and for period t+2.

In period t when you make your inflation predictions for t+1 and t+2 you can observe the current and past data of the economy. This data consists of the current and past inflation rates and the current and past levels of real GDP, where real GDP is the quantity of goods that is produced in the economy.

At the beginning of an experiment when you start forecasting, there is just a single data point that consists of the current inflation rate and output level. After you have made your forecasts the experiment period will end and a new experiment period will start for which a new inflation rate and output level will be announced. Thus, as the experiment evolves you will have an increasing number of observations.

There will be various experiment 'sessions'. For each 'session' the economy will restart from period zero. Each session is unrelated to the previous session, in the sense that the level of inflation and output will be different across sessions. Also the relationships between inflation and output and past values of these variables is not necessarily the same from one session to another. The end of a session and the beginning of a new session will be clearly announced by the experimenter.

Earnings

During each period of an experiment session you will collect 'points' which at the end of the session will be transformed into Lira, as described below. The number of points that you get will depend on how close your inflation predictions are to the actual inflation rates. The details are explained now:

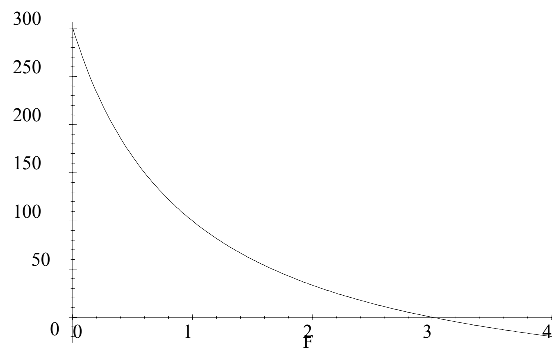

Each period t, the new inflation rate and the new output level are announced. You will have predicted the current inflation rate two different times, once 1 period ago and once 2 periods ago.

Let f denote the absolute value of your forecast error from one of these forecasts. The error is expressed in percentage points, i.e. if $f = 1.5$ your forecast was either 1.5% higher or lower than the actual inflation rate.

The points that you receive will depend on the errors f you make where larger errors will give you less points. In particular, points are calculated in the following way

$$\frac{400}{1+f} - 100$$

You can receive up to 300 points per forecast and may lose up to 100 points depending on the size of the forecast error. With a zero forecast error you would receive 300 points. However, if your forecast is 1% higher or lower than the actual inflation rate you will get only 100 points (400/2-100), likewise for a 3% forecast error you receive no points (300/3-100), and for even larger forecast errors points will be subtracted. The graph below shows the relationship between the forecast error and the points that your receive for your forecast.

You will receive points for each of the two forecast you made for the current inflation rate, i.e. you receive points for the forecast you made 1 period ago and points for the forecast you made 2 periods ago, where the number of points for each forecast depends on the forecast error as described above.

After the experimenter has announced the end of an experiment session, write down the total number of points that you received on a sheet of paper with your name on it. Briefly after the end of the session, the experimenter will announce a conversion rate that indicates the value of the points in terms of Lira.

Other Instructions

During the experiment sessions it is strictly forbidden to speak with other students that participate in the experiment. Doing so can lead to the exclusion from the experiment. In this case no payment will be made. If you have any questions or problems during the course of the experiment raise your hand and the experimenter will come to you.

At the start of each experiment session you will be asked to start the program that runs on your computer. Please carefully follow the instructions that you will receive from the experimenter.

If you have any questions please ask them now!

References

Adam, Klaus, "Adaptive Learning and Cyclical Behavior of Output and Inflation," *University of Frankfurt Mimeo, www.wiwi.uni-frankfurt.de/professoren/wieland/adam*, 2001.

___ **and Ramon Marimon**, "Recursive Experimental Economics with MacroLab: Presenting the Framework," 2001. University of Frankfurt Mimeo, available at www.wiwi.uni-frankfurt/professoren/wieland/adam.

Anderson, Robert M. and Hugo Sonnenschein, "Rational Expectations Equilibrium with Econometric Models," *Review of Economic Studies*, 1985, *52*, 359–369.

Ball, Laurence, "Near-Rationality and Inflation in Two Monetary Regimes," *Johns Hopkins University Mimeo*, 2000.

Bhansali, R. J., "Multi-Step Forecasting," in David Hendry, ed., *A Companion to Economic Forecasting*, Blackwell 2002 forthcoming.

Evans, George and Garey Ramey, "Adaptive Expectations, Underparameterization, and the Lucas Critique," 2001. University of Oregon, mimeo.

Evans, George W. and Garey Ramey, "Expectation Calculation and Macoeconomic Dynamics," *American Economic Review*, 1992, *82*, 207–224.

___ **and Seppo Honkapohja**, "Adaptive Forecasts, Hysteresis, and Endogenous Fluctuations," *Economic Review of the Federal Reserve Bank of San Fransico*, 1993, *1*, 3–13.

___ **and** ___ , *Learning and Expectations in Macroeconomics*, Princeton: Princeton University Press, 2001.

Grandmont, Jean Michel, *Temporary Equilibrium: Selected Readings* Economic Theory, Econometrics, and Mathematical Economics Series, London: Academic Press, 1988.

Marimon, Ramon and Shyam Sunder, "Indeterminacy of Equilibria in a Hyperinflationary World: Experimental Evidence," *Econometrica*, 1993, *61(5)*, 1073–1107.

___ **and** ___ , "Expectations and Learning under Alternative Monetary Regimes: An Experimental Approach," *Economic Theory*, 1994, *4*, 131–162.

Sargent, Thomas J., *The Conquest of American Inflation*, Princeton: Princeton Univ. Press, 1999.

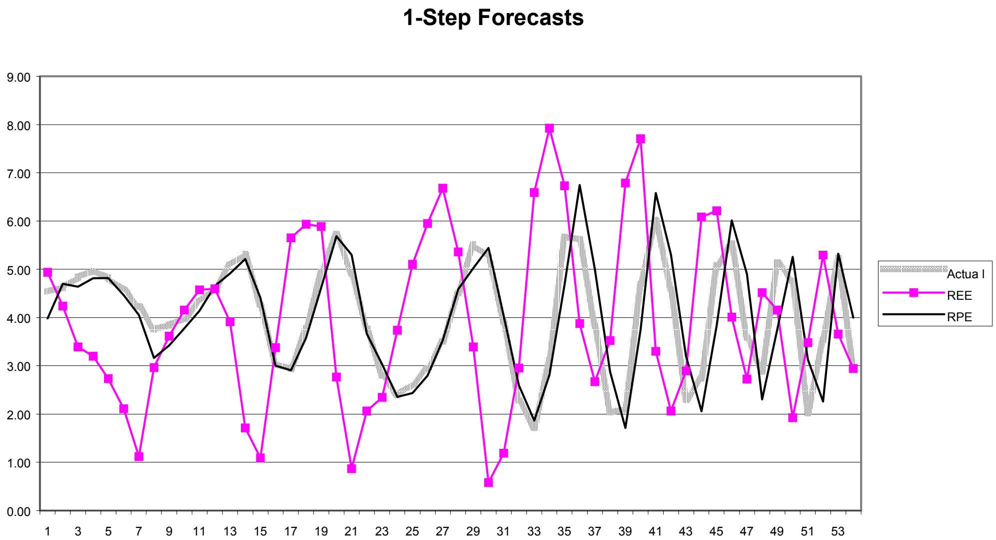

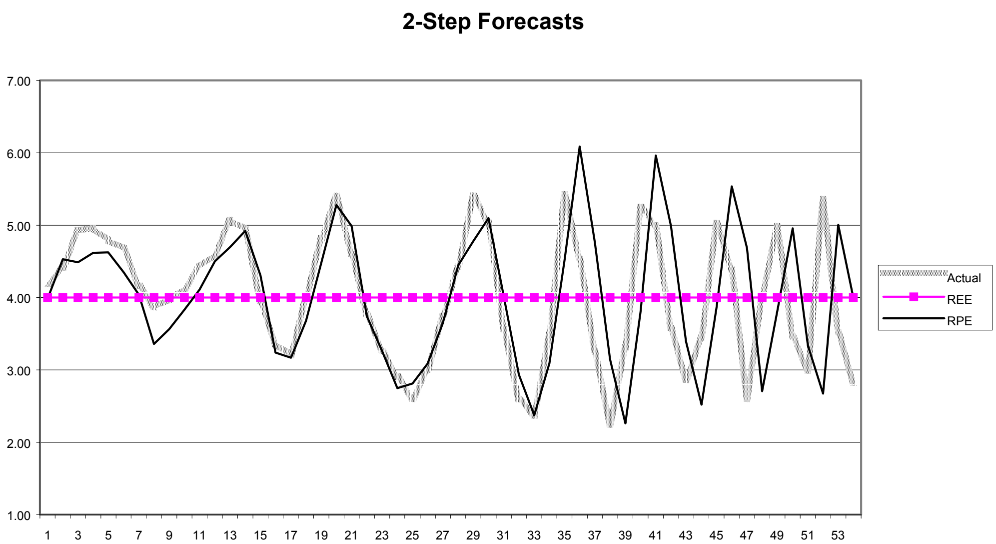

Figure 1: Session 1 (T1)

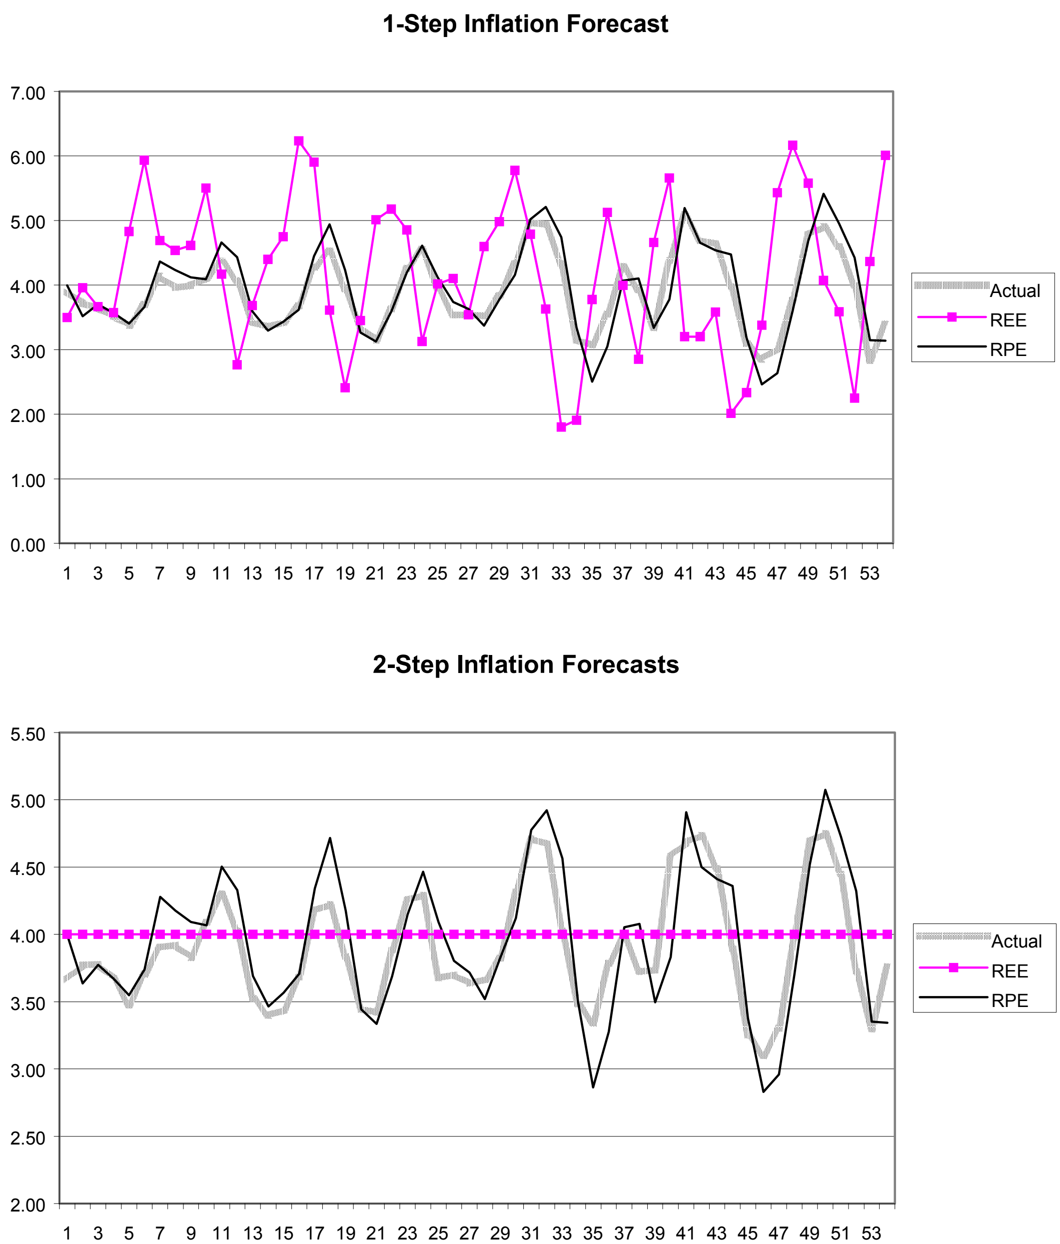

Figure 2: Session 2 (T1)

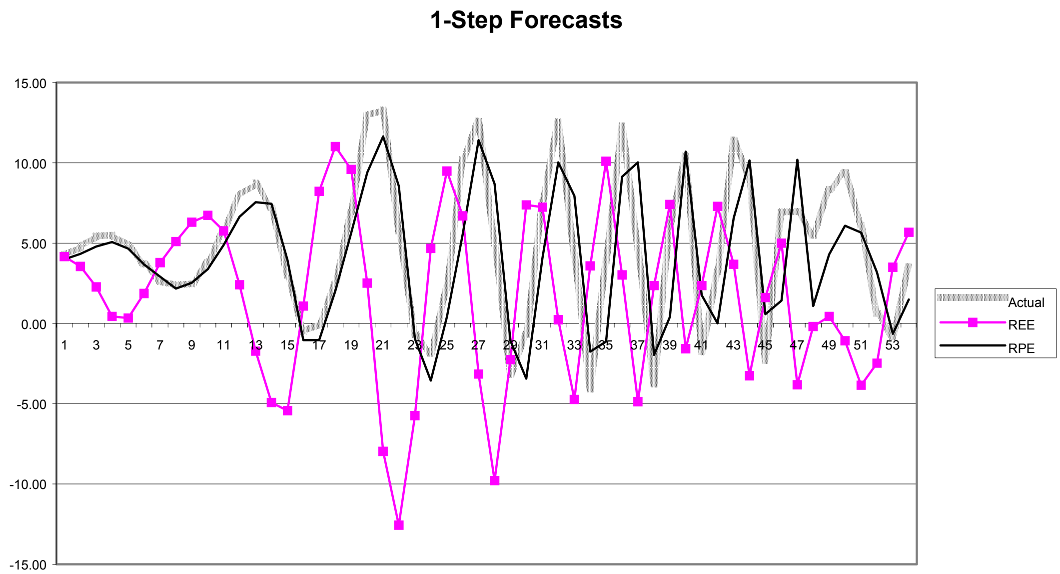

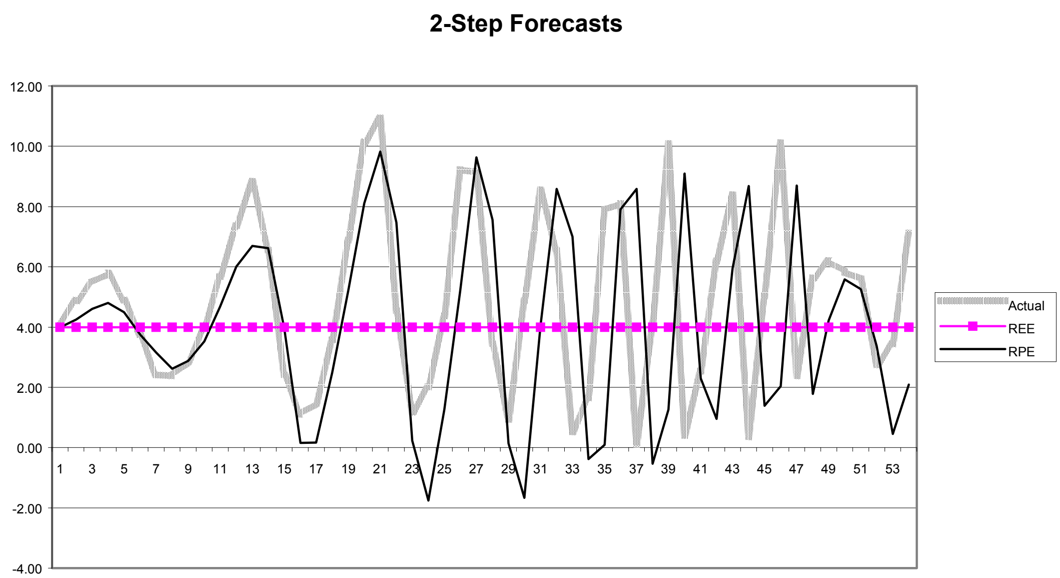

Figure 3: Session 3 (T1)

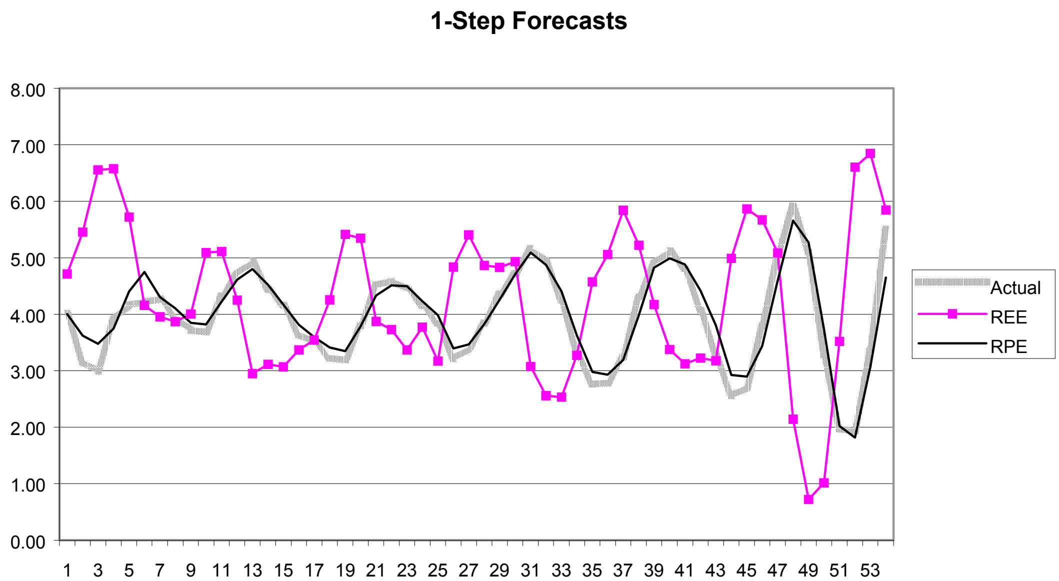

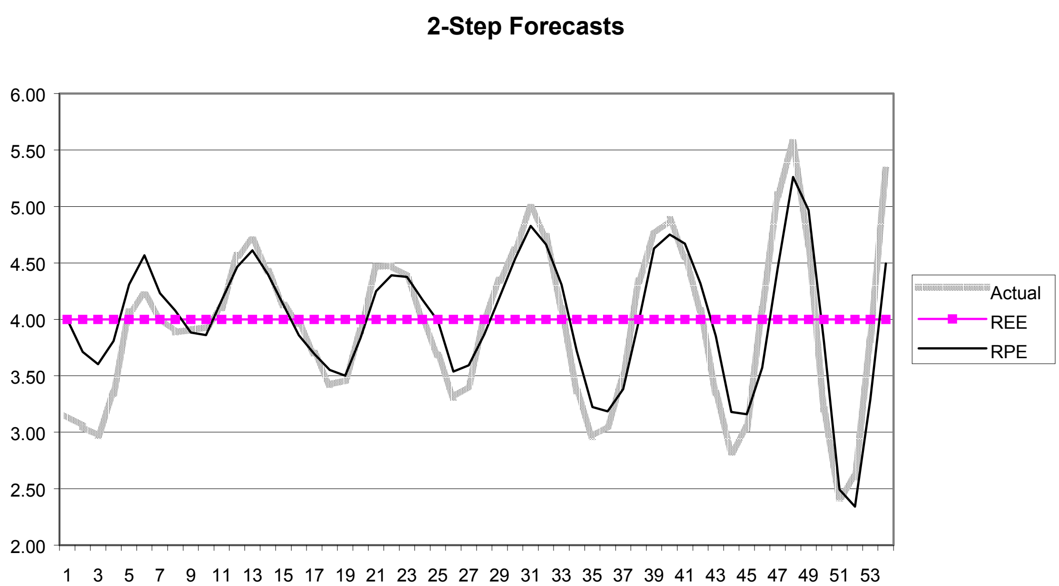

Figure 4: Session 4 (T1)

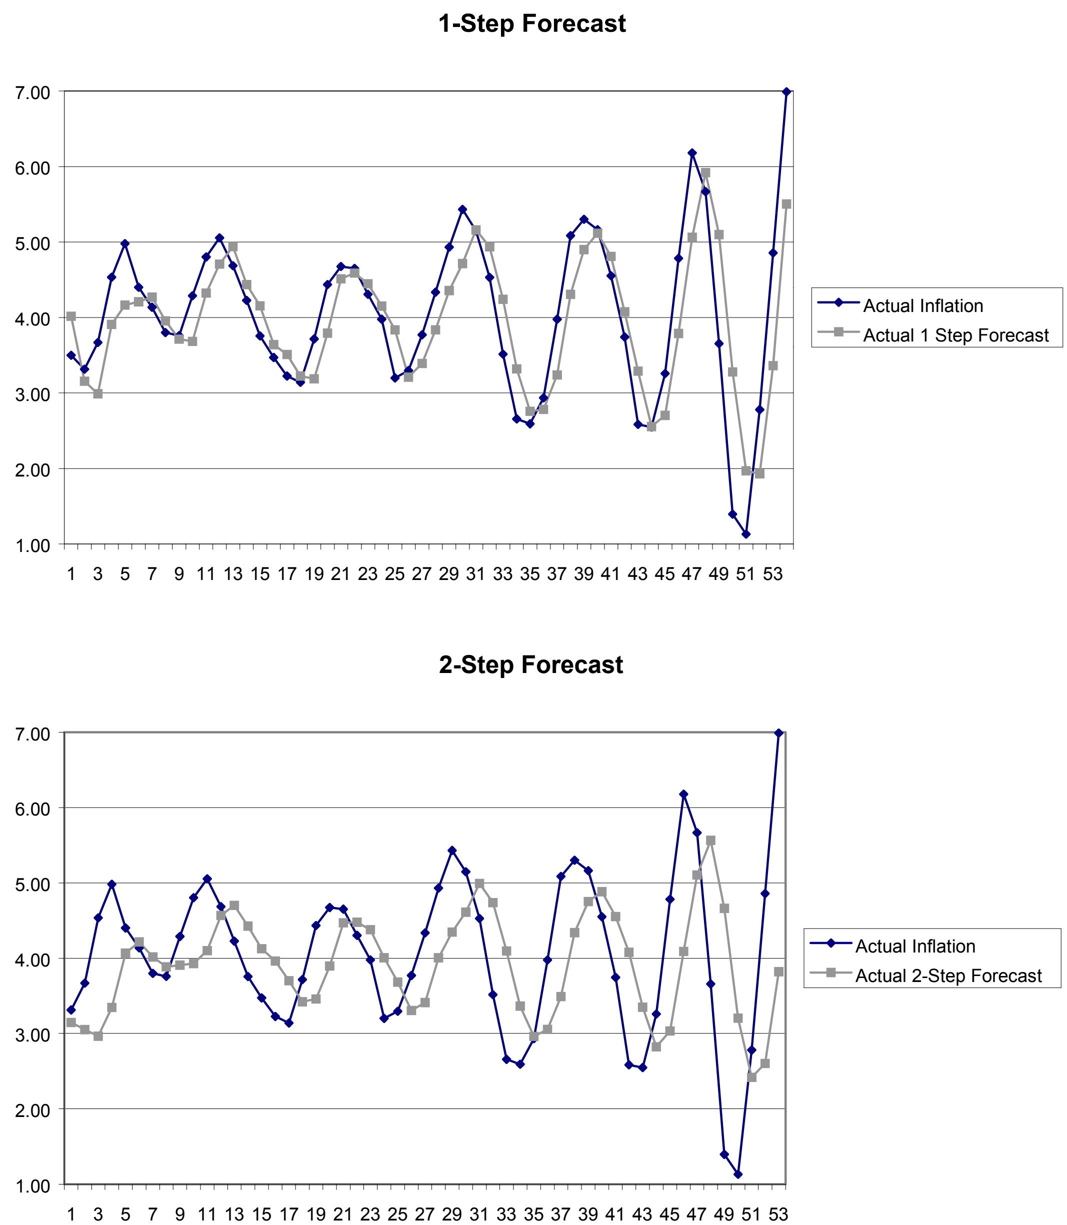

Figure 5: Actual and Predicted Inflation, Session 4 (T1)

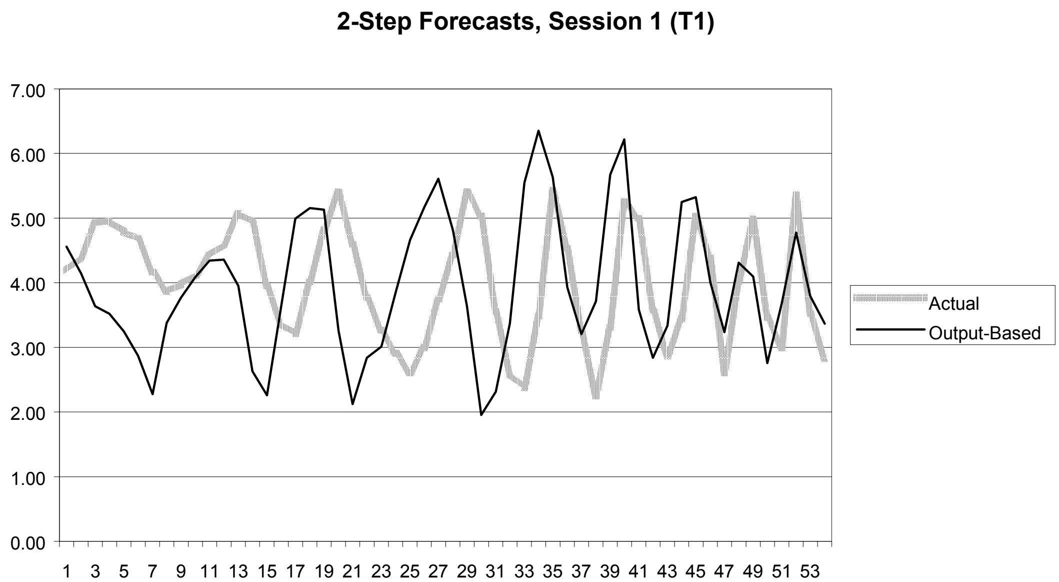

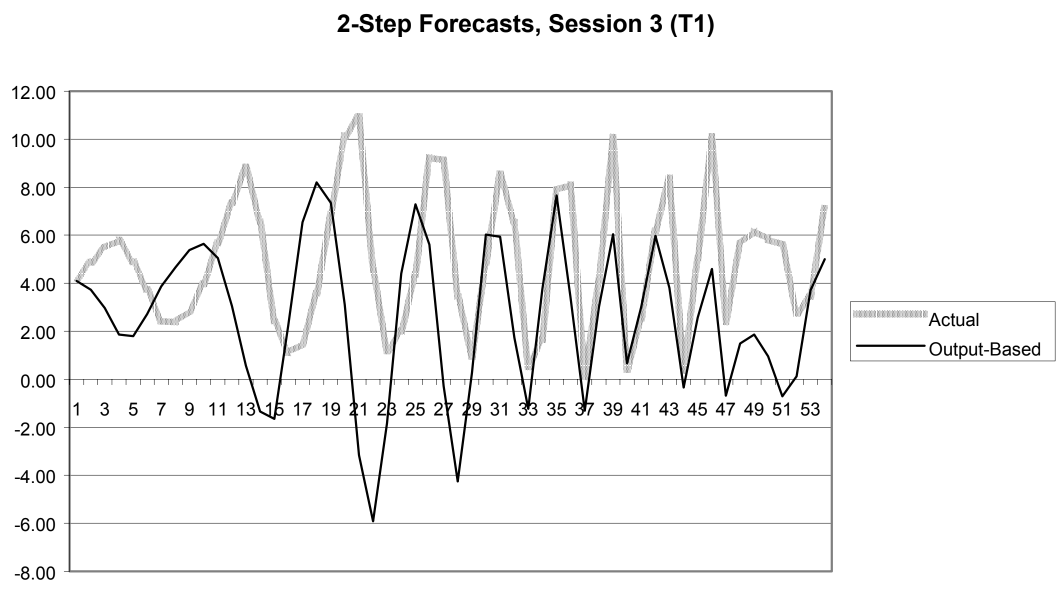

Figure 6: Output-Based 2-Step Forecasts

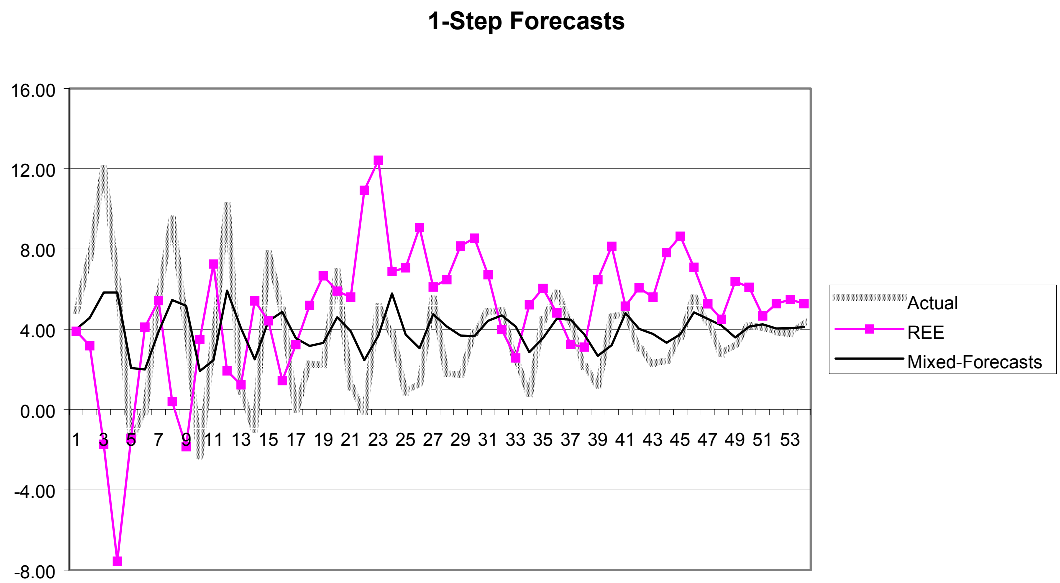

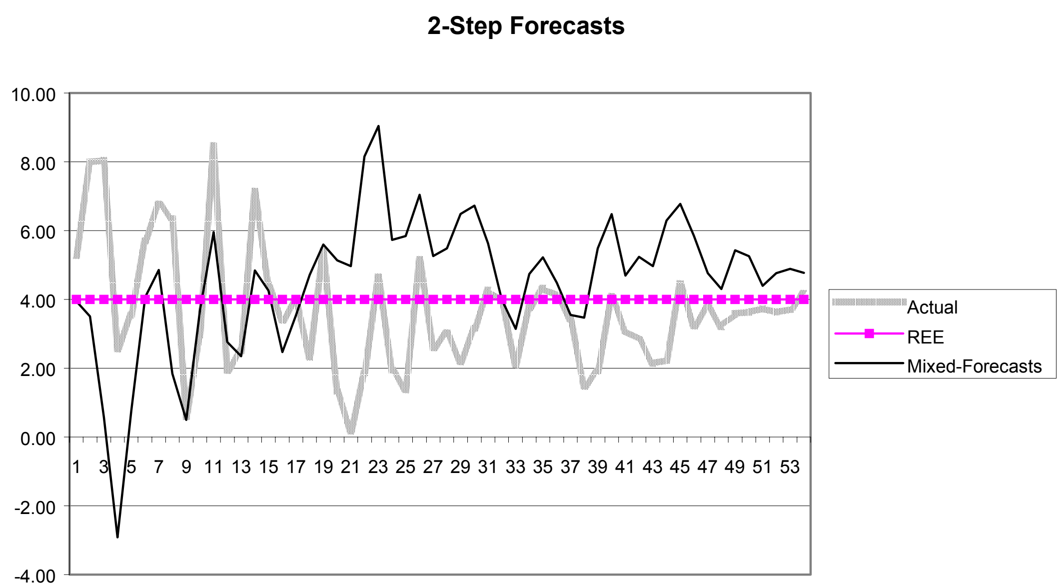

Figure 7: Session 3 (T2)

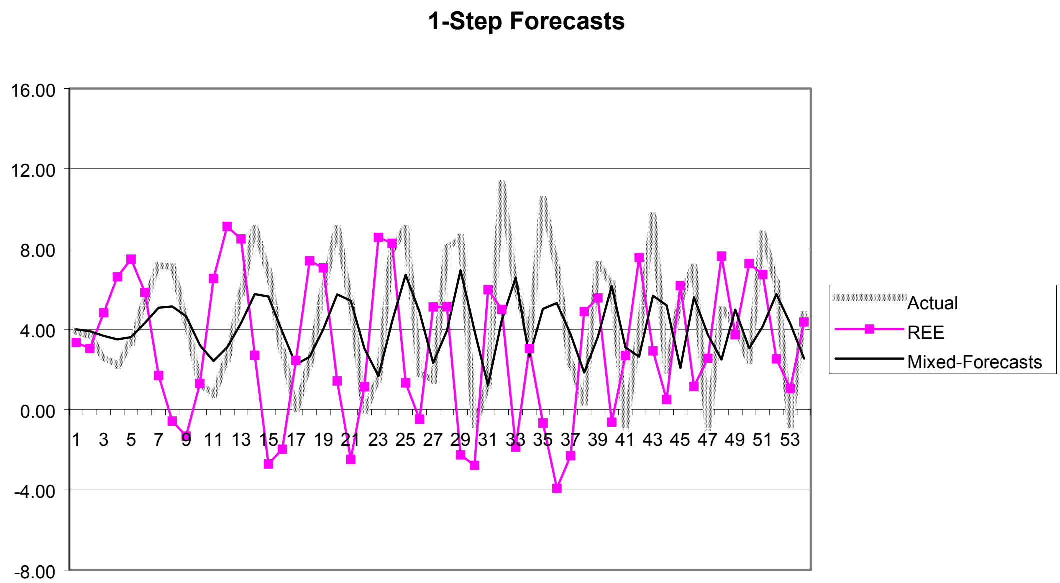

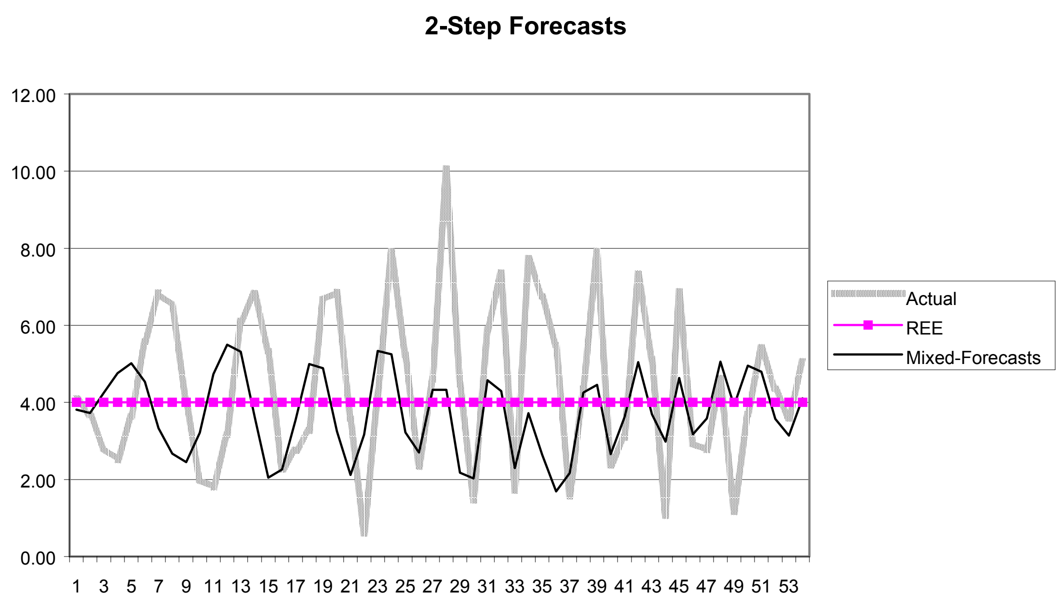

Figure 8: Session 4 (T2)

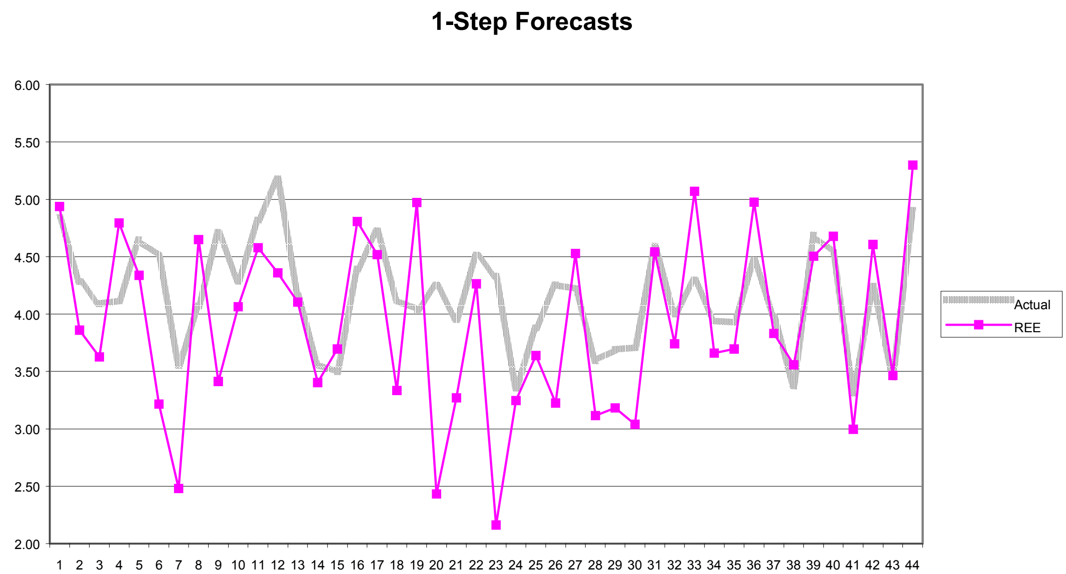

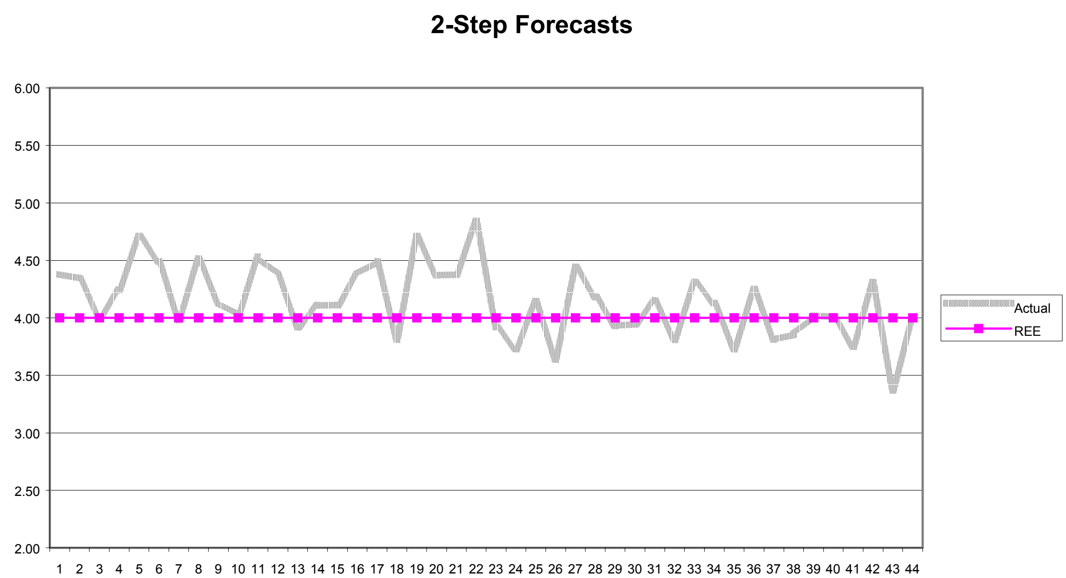

Figure 9: Session 5 (T1)

1-Step Forecast

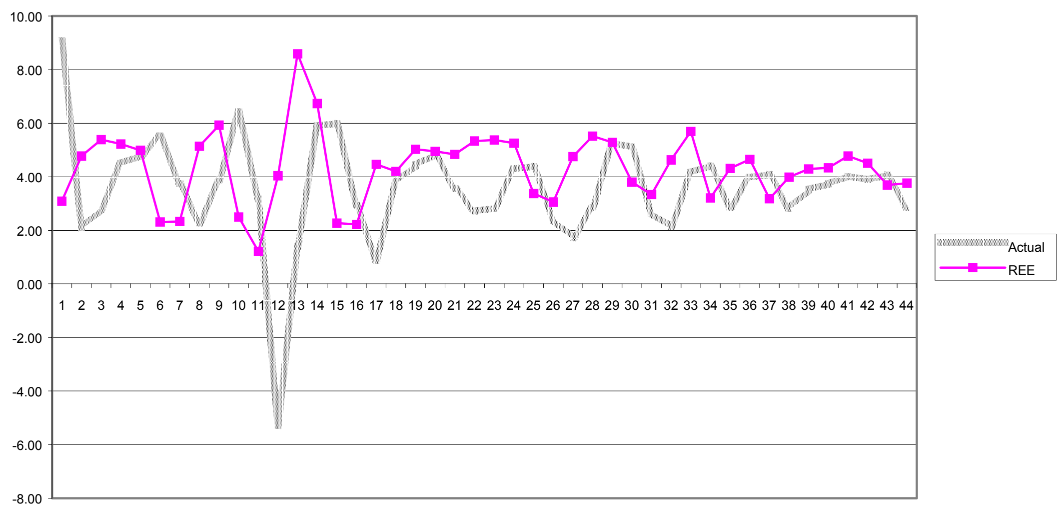

2-Step Forecasts

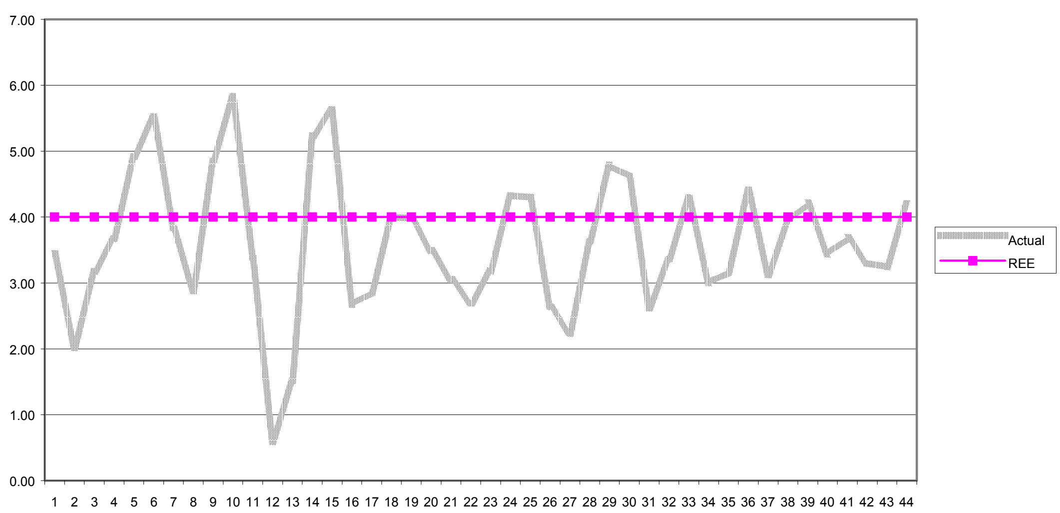

Figure 10: Session 6 (T1)

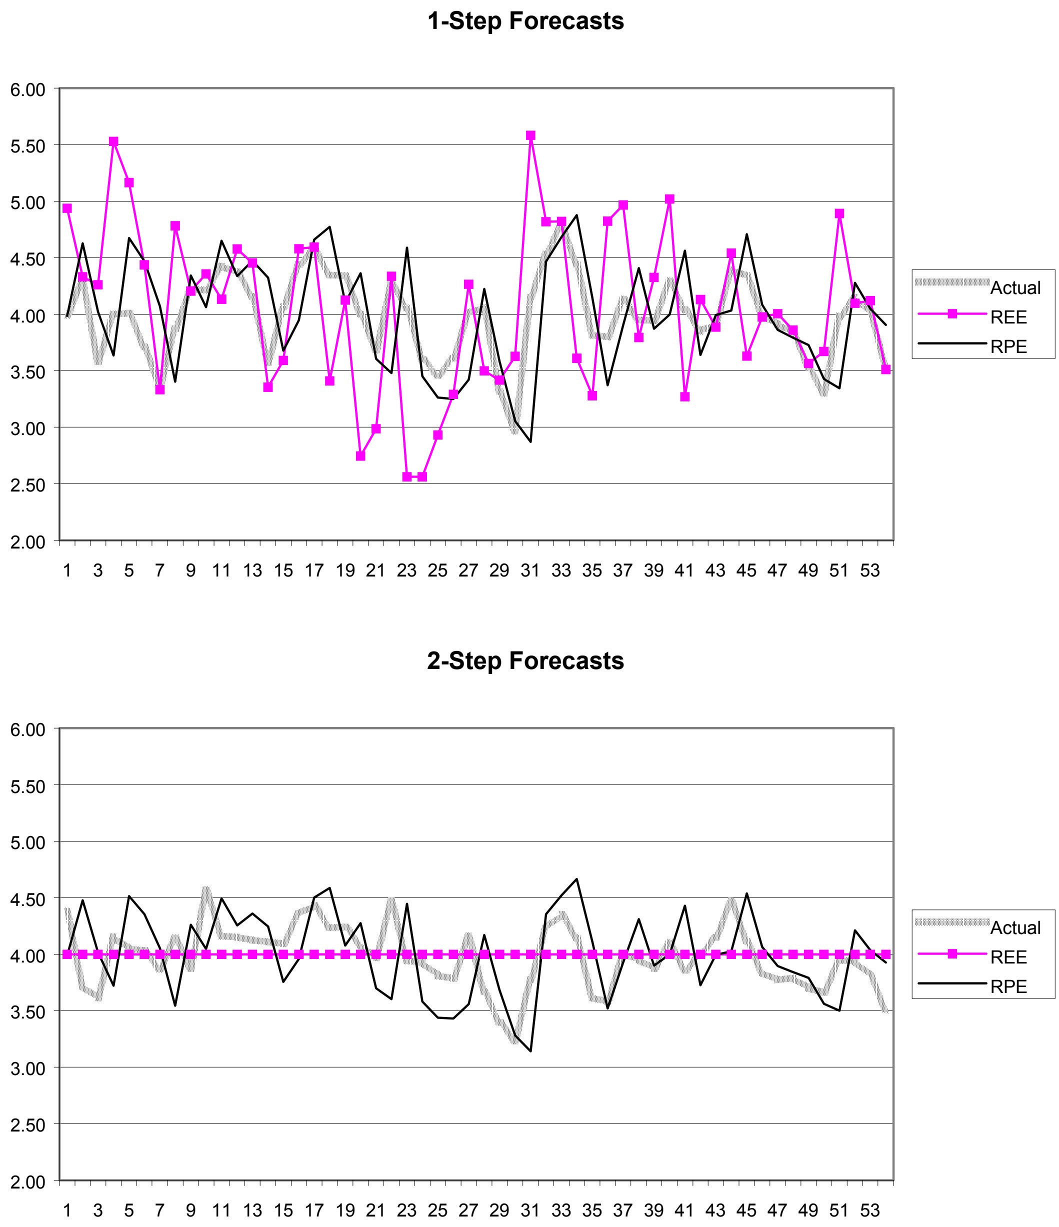

Figure 11: Session 5 (T2)

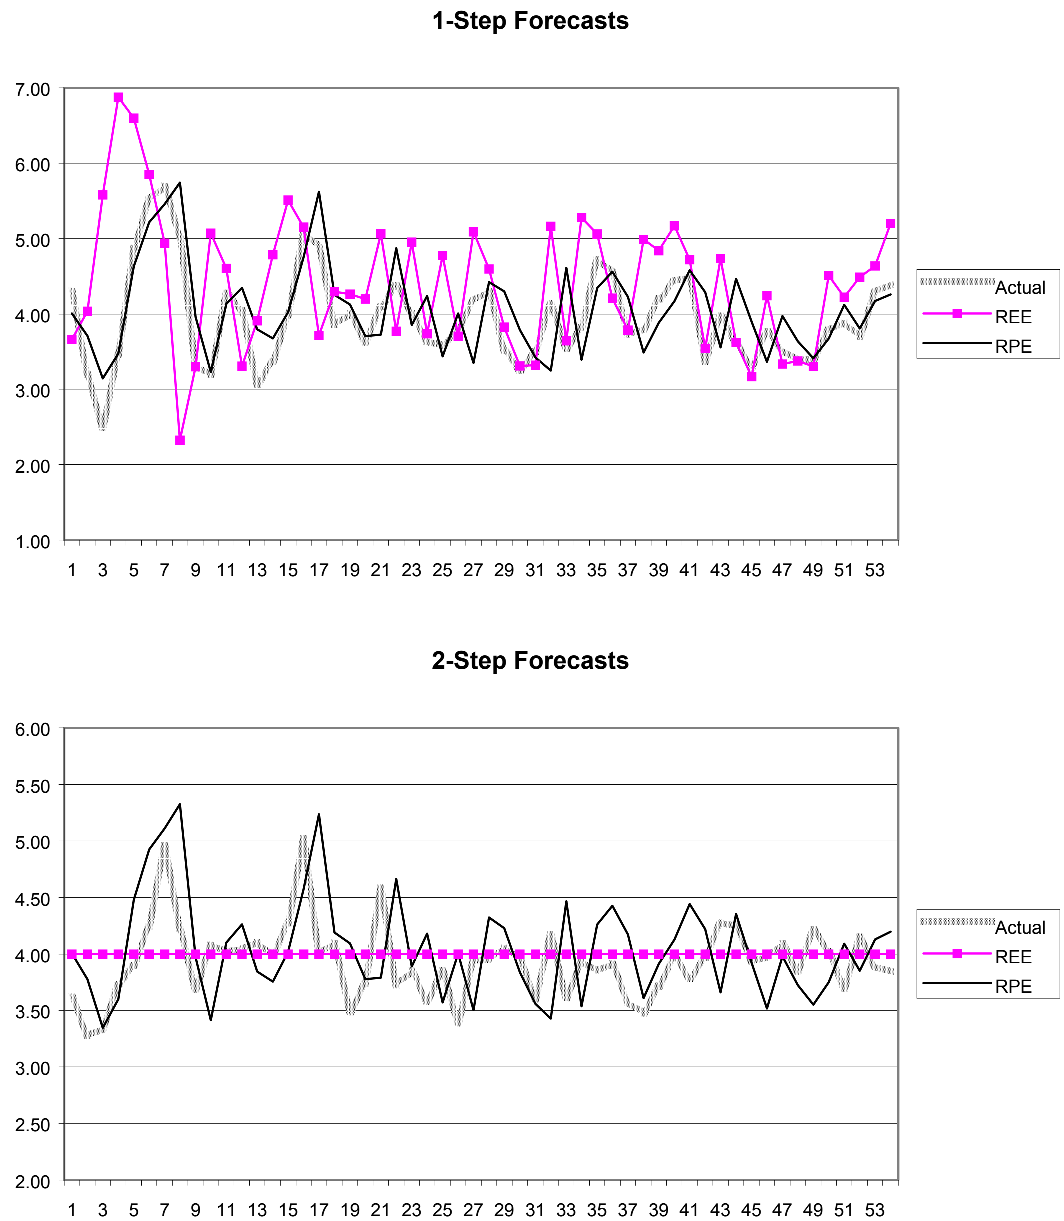

Figure 12: Session 6 (T2)

Table 1: Parameterization of the experimental treatments

	Treatment 1	Treatment 2 (T1)	Location (T2)	Date
Session 1	high	-	Salerno, Italy	July 16, 2001
Session 2	high	-	Salerno, Italy	July 23, 2001
Session 3	high	high	Frankfurt, Germany	June 3, 2002
Session 4	high	high	Frankfurt, Germany	June 7, 2002
Session 5	low	high	Salerno, Italy	July 2, 2001
Session 6	low	high	Salerno, Italy	July 24,2001

'low' indicates treatments where the elasticity of labor supply is given by $\varepsilon = 1.0$, 'high' indicates treatments where the elasticity is given by $\varepsilon = 2.0$.

Table 2: Average Inflation Forecasts (Across Agents and Periods)

	Average 1-Step Forecast	Average 2-Step Forecast
Session 1 (T1)	4.03	4.01
	(0.1584)	(0.1316)
Session 2 (T1)	3.89	3.90
	(0.0972)	(0.0771)
Session 3 (T1)	4.75	4.90
	(0.5138)	(0.3302)
Session 3 (T2)	3.64	3.65
	(0.2801)	(0.2642)
Session 4 (T1)	3.95	3.94
	(0.1481)	(0.1257)
Session 4 (T2)	4.48	4.43
	(0.3104)	(0.2202)
Session 5 (T1)	4.16	4.14
	(0.0605)	(0.0535)
Session 5 (T2)	3.98	3.97
	(0.0597)	(0.0499)
Session 6 (T1)	3.58	3.61
	(0.2284)	(0.1232)
Session 6 (T2)	3.95	3.93
	(0.1024)	(0.0559)

Newey-West standard errors (3 lags) in parentheses.

Table 3: Baseline Treatments

1-Step Forecast	Share of RPE-Forecasters (vs. REE) whole sample	last 20 periods
Session 1 (T1)	0.887	0.747
	(0.0324)	(0.0427)
Session 2 (T1)	0.881	0.855
	(0.0148)	(0.0265)
Session 3 (T1)	0.873	0.778
	(0.0319)	(0.0530)
Session 4 (T1)	0.989	0.969
	(0.0180)	(0.0228)

2-Step Forecast	Share of RPE-Forecasters (vs. REE) whole sample	last 20 periods
Session 1 (T1)	0.534	0.168
	(0.1477)	(0.1552)
Session 2 (T1)	0.694	0.687
	(0.0501)	(0.0584)
Session 3 (T1)	0.212	-0.436
	(0.1896)	(0.1435)
Session 4 (T1)	1.087	1.069
	(0.0704)	(0.0677)

2-Step Forecast	Share of RPE-Forecasters (vs. Ouput-Based) whole sample	last 20 periods
Session 1 (T1)	0.713	0.354
	(0.0892)	(0.1025)
Session 2 (T1)	0.792	0.744
	(0.0286)	(0.0407)
Session 3 (T1)	0.508	0.068
	(0.1197)	(0.0947)
Session 4 (T1)	0.988	0.928
	(0.0474)	(0.0334)

Newey-West standard errors (3 lags) in parentheses.

Table 4: Additional High Elasticity Treatments

1-Step Forecast	Share of Mixed-Forecasters (vs. REE) whole sample	last 20 periods
Session 3 (T2)	1.14	0.933
	(0.0646)	(0.0912)
Session 4 (T2)	0.950	0.890
	(0.0792)	(0.1663)

2-Step Forecasts	Share of Mixed-Forecasters (vs. REE) whole sample	last 20 periods
Session 3 (T2)	0.344	0.449
	(0.3561)	(0.3468)
Session 4 (T2)	0.609	0.812
	(0.2715)	(0.5870)

2-Step Forecast	Share of Mixed-Forecasters (vs. RPE) whole sample	last 20 periods
Session 3 (T2)	0.835	0.711
	(0.1119)	(0.1437)
Session 4 (T2)	0.795	1.034
	(0.0887)	(0.1493)

Newey-West standard errors (3 lags) in parentheses. Session 3 included a dummy variable for period 20 to 30.

Table 5: Low Elasticity Treatments

1-Step Forecast	Coefficient on REE Forecasts whole sample	last 20 periods
Session 5 (T1)	0.333	0.717
	(0.0766)	(0.0927)
Session 6 (T1)	0.226	0.216
	(0.0633)	(0.1055)
Session 1 (T1)	-0.136	0.029
	(0.0468)	(0.1147)
Session 2 (T1)	-0.025	-0.059
	(0.0450)	(0.0923)
Session 3 (T1)	-0.168	-0.182
	(0.0455)	(0.0993)
Session 4 (T1)	-0.253	-0.249
	(0.0517)	(0.0993)
Session 3 (T2)	-0.220	0.000
	(0.0776)	(0.1295)
Session 4 (T2)	-0.0776	-0.002
	(0.0733)	(0.1765)

Newey-West standard errors (3 lags) in parentheses.

Table 6: Low Elasticity Treatments

2-Step Forecast	(Actual - REE Forecast)^2 whole sample	last 20 periods
Session 5 (T1)	0.116	0.007
	(0.0244)	(0.0369)
Session 6 (T1)	1.260	0.321
	(0.3941)	(0.0720)
Session 1 (T1)	0.787	0.9839
	(0.1044)	(0.1207)
Session 2 (T1)	0.1945	0.298
	(0.0332)	(0.0554)
Session 3 (T1)	9.123	10.119
	(1.6336)	(2.2004)
Session 4 (T1)	0.493	0.886
	(0.1106)	(0.1642)
Session 3 (T2)	3.307	1.102
	(0.7934)	(0.4665)
Session 4 (T2)	4.672	3.922
	(0.7331)	(0.8545)

Newey-West standard errors (3 lags) in parentheses.

Table 7: High Elasticity Treatments, Sessions 5 and 6

1-Step Forecast	Share of RPE-Forecasters (vs. REE) whole sample	last 20 periods
Session 5 (T2)	0.655	0.619
	(0.0391)	(0.0257)
Session 6 (T2)	0.761	0.503
	(0.0759)	(0.0696)

2-Step Forecast	Share of RPE-Forecasters (vs. REE) whole sample	last 20 periods
Session 5 (T2)	0.274	0.266
	(0.1024)	(0.1098)
Session 6 (T2)	0.231	-0.113
	(0.1164)	(0.1601)

Newey-West standard errors (3 lags) in parentheses.

Table 8: High Elasticity Treatments, Sessions 5 and 6

1-Step Forecast	Coefficient on REE Forecasts whole sample	last 20 periods
Session 5 (T2)	0.226	0.170
	(0.0549)	(0.1014)
Session 6 (T2)	0.180	0.527
	(0.1001)	(0.0654)

2-Step Forecast	$(\text{Actual - REE Forecast})^2$ whole sample	last 20 periods
Session 5 (T2)	0.078	0.062
	(0.0167)	(0.0164)
Session 6 (T2)	0.114	0.050
	(0.0272)	(0.0182)

Newey-West standard errors (3 lags) in parentheses.