Learning, forecasting and optimizing: an experimental study

Bao, T.; Duffy, J.; Hommes, C.

[Link to publication](#)

Learning, Forecasting and Optimizing: an Experimental Study[*]

Te Bao[a], John Duffy[b], and Cars Hommes[a]

[a]Department of Quantitative Economics, University of Amsterdam

Roetersstraat 11, 1018WB Amsterdam, NL

[b]Department of Economics, University of Pittsburgh

4901 Wesley W. Posvar Hall, PA 15260, Pittsburgh, USA

October 14, 2011

Abstract

Rational Expectations (RE) models have two crucial dimensions: 1) agents correctly forecast future prices given all available information, and 2) given expectations, agents solve optimization problems and these solutions in turn determine actual price realizations. Experimental testing of such models typically focuses on only one of these two dimensions. In this paper we consider both forecasting and optimization decisions in an experimental cobweb economy. We report results from four experimental treatments: 1) subjects form forecasts only, 2) subjects determine quantity only (solve an optimization problem), 3) they do both and 4) they are paired in teams and one member is assigned the forecasting role while the other is assigned the optimization task. All treatments converges to Rational Expectation Equilibrium (REE), but the at very different speed. We observe that performance is the best in treatment 1) and worst in the treatment 3). Most forecasters use an adaptive expectations rule. Subjects are less likely to make conditionally optimal production decision for given forecasts in treatment 3) where the forecast is made by themselves, than treatment 4) where the forecast is made by the other member of the team, which confirms "two heads are better than one" in finding REE.

1 Introduction

Rational Expectations (RE) macroeconomic models have two crucial dimensions: 1) Agents use the model to correctly forecast future prices given all available information, that is, they do not make systematic mistakes. 2) Given expectations, agents solve an optimization problem to determine their consumption or production decisions, which then, via market clearing, determines the realization of the price level they were seeking to forecast. Thus, RE systems are self-referential and assume rational, optimizing behavior on the part of agents. Testing such models in controlled laboratory experiments has proved difficult owing to the complicated nature of RE models. As Nobel Prize laureate Sargent (2008) observes:

> "Laboratory experiments using macroeconomics are rarer than those using microeconomics...I suspect that the main reason for fewer experiments in macro than in micro is that the choices confronting artificial agents within even one of the simpler recursive competitive equilibria used in macroeconomics are very complicated relative to the settings with which experimentalists usually confront subjects."

Experimentalists seeking to test RE macro models have dealt with the complicated nature of these models by reducing the dimensionality of the problem that subjects face. Two approaches have been taken.

In a "learning to forecast experiment, " – a design first proposed by Marimon and Sunder (1993) – subjects are asked to submit a forecast for a future economic variable (e.g., a price, inflation rate, foreign exchange rate, etc.), and they are rewarded solely on the basis of the accuracy of their forecast. Their forecast is then used as an input by a computer program to determine each individual's optimal quantities, as if the subjects themselves were capable of solving the optimization problem conditional on their forecast. The computer-determined quantities together with market clearing conditions determine actual price realizations (or the forecast variable of interest), and these realizations are then used to assess the subjects' forecast accuracy. Subjects, however, are not necessarily made aware of how their forecasts affect outcomes; for the subjects the determination of actual realizations of forecasted variables often amounts to a "black-box" process.

In a second, older experimental approach, known as the "learning to optimize experiment" (LtOE) design, subjects are asked to make economic decisions (to consume, invest, trade, produce, etc.) directly, without any elicitation of their forecasts of the relevant endogenous variables such as the market price. Of course, such forecasts can be determined implicitly based on subjects' decisions or are sometimes determined separately via some market mechanism (e.g., a double auction or a call market) that is often external to the theory being tested.

Studies using the LtFE approach find mixed evidence as to whether subjects are able to learn rational expectations (see e.g., Hommes 2011 for a survey). In some instances, subjects learn rational expectations via some adaptive learning process while in other instances subjects behave as trend extrapolators resulting in persistent deviations or cycles around the rational expectations equilibrium. Similarly, findings from LtOE studies have sometimes confirmed competitive equilibrium predictions and associated comparative statics predictions, but in other instances have generated outcomes that are at odds with RE model predictions, for instance, non-rational bubbles, excess volatility, etc.

In this paper we compare the LtFE and LtOE approaches in a common, economic decision-making task. Importantly, we also consider how behavior improves or deteriorates if we combine these two approaches. Our combined LtFE and LtOE design gets at the heart of the belief-outcome interaction that is the signature property of rational expectations models. We ask if convergence to the REE and efficiency are affected when subjects are asked to play both roles as forecaster and optimizer or if specialization of tasks by individuals alone (as in LtFE and LtOE designs) or within two-agent teams leads to a significant improvement in performance. A main aim of this research is to assess whether the results in the former LtFE literature are robust when the optimization task is performed by an individual rather than by a computer program. Moreover, the team specialization treatment that we add has a very natural, real-world interpretation: Organizational investors such as investment banks and pension funds usually employ both professional forecasters (researchers and economists) and production managers or traders.

The experimental environment we study is a simple, N-firm cobweb model economy – a negative expectation feedback system. This kind of feedback system arises naturally in commodity markets that were the inspiration for Ezekiel's (1938) devel-

opment of the cobweb model. Furthermore, Muth (1961) proposed rational expectations in the context of this same negative feedback cobweb model. Prior research indicates that under a LtFE design, market prices will converge very quickly to the RE equilibrium in this environment. In addition to LtFE, we consider three additional treatments where the subjects need to submit their production decision directly without a forecast (LtOE), or together with a forecast, or subjects are paired in teams and one submits a forecast which the other can use to determine a production decision.

We find some tendency for the market price to converge to the RE equilibrium price in all four treatments. Thus, the stabilizing effect of a negative feedback market is a robust feature. However, when the volatility and speed of convergence are compared, we find that the market price converges most quickly and reliably when subjects only make price forecasts as in the LtFE design. There is not much difference in performance between the treatments where subjects only make production decisions (LtOE) and where they form teams that specialize in one of the two tasks. The market price and quantity fluctuate the most and are the slowest to converge when subjects are required to do both tasks, forecasting and production decision-making. Our findings have important implications for both the design of experiments and for how to think about the representative agent firm: should it be viewed as an individual actor (e.g., the C.E.O.) or is it better to think of the representative firm as consisting of teams of individuals specialized in various tasks?

The rest of the paper is organized as follows: Section 2 discusses the related literature. Section 3 describes our experimental design. Section 4 presents the experimental results. Finally, section 5 concludes.

2   Related Literature

Our work is related to former LtFE and LtOE studies. Smith et al. (1988), Lim et al. (1994), Arifovic (1996), Lei et al. (2001) and Crockett and Duffy (2010) are some examples of LtOE studies. Adam (2007), Marimon et al. (1993), Marimon and Sunder (1993, 1994, 1995), Hommes et al. (2005, 2007) and Heemeijer et al (2009) are some representative works in the LtFE literature.

As we also have a treatment where subjects participate as members of teams, our experiment is related to the literature on the comparison of group and individual decisions. In the context of experimental macroeconomics and finance, Blinder and Morgan (2005) show that monetary policy decisions made by groups are not slower than those made by individuals, and are generally better; Kocher and Sutter (2005) find that groups learn faster, and can beat individuals when they are playing as opponents in Beauty-Contest Games. There is a parallel literature in experimental game theory on individual versus group decisions. The evidence is mixed on whether groups are more "rational" or self-interested than individuals. Bornstein and Yaniv (1998) find groups offer less and accept less in the ultimatum game relative to individuals. Cox (2002) shows that there is no significant difference between group and individual decisions in the trust game. Cason and Mui (1997) find that groups offer more in dictator games than indvididuals. In all of these group-versus-individual-studies, group members are asked to perform/participate in the same kind of the task, and the decision of the group is usually the average or majority choice of the group members. By contrast, our team treatment involves specialization of tasks between the two group members.

Our work is also related to the experiments on Cournot oligopoly. Offerman, Potters and Sonnemans (2002) demonstrate that giving subjects different information about other firms' behavior (information about the sum of the other firms' quantity only, about individual firm's quantity only or about individual firm's quantity and profit) can lead to different learning rules, and market evolution towards different equilibria (Walrasian, Collusive and Cournot-Nash). In our experiment, subjects have no information about other firm's quantity and profit at all. They also have no information about the relationship between the market price and total output. As the optimal quantity decision requires them to set price equal to marginal cost, the rational expectation equilibrium in this Cournot market is the same as the Walrasian outcome. Huck, Normann and Oechssler (1999) vary information available to subjects from full information about the market including others' decisions and profits and their own decision and profit to only their own decision and profit, while they still know the number of subjects in the market. They found none of the treatments generates successful collusion, and information that encourage "imitating the best" learning leads to Walrasian outcome, which confirms the prediction of Vega-Redondo (1997). Their NOIN treatment, where subjects have no information about others'

behavior is similar to the information we provide subjects except that their subjects know the number of firms in the market. Their NOIN treatment generates an outcome very close to the Walrasian outcome and that is why we chose this informational structure for our experiment. However, as they use constant marginal cost in their paper, the optimal quantity given a price prediction is piecewise linear, and generates no steady state. It is therefore not possible to test convergence to RE equilibrium using their experimental design.

3 Experimental Design

3.1 Treatments

Our experiment consists of four treatments that differ in the tasks assigned to participants and in the payoff scheme. Sample experimental instructions are provided in the Appendix. Subjects are playing the role of firms only, deciding on price forecasts or optimal production amounts or both.

1. Treatment 1: the LtFE treatment. In this treatment, subjects (firms) only make price forecasts. Each firm's production decision is calculated by the computer optimally, given the firm's price forecast. Each subject is paid according to the accuracy of his forecast alone. The forecasters know: the history of the market price they are attempting to forecast which is standard in the LtFE literature and the history of their own forecasts and payoffs. Each subject can read his payoff from the forecasting task for different prediction errors from the payoff table (See Appendix, "Payoff Table for the Forecaster").

2. Treatment 2: the LtOE treatment. In this treatment, subjects (firms) only make quantity (or production) decisions. Each subject knows the history of the market price, his own prior decisions and profits. Each subject makes a quantity decision only; there is no elicitation of a subjects price forecast. The market price is determined by the production decisions submitted by all firms in the market. Each subject is paid according to the profit his firm makes each period. He can read his payoff for different combinations of the market price

and his production (optmization) decisions from the payoff table (See Appendix, "Payoff Table for the Production Manager").

3. Treatment 3: the LtFE+LtOE Individual treatment. In this treatment, each subject plays the role of both forecaster and production manager. Each subject knows the history of the market price, his prior decisions and profits. Each subject makes both a price forecast and a quantity decision. The market price is determined by the quantity decisions of all firms in the market. Subjects are paid according to an equal weighted linear combination of the payoff functions used in the LtFE and LtOE treatments. Each subject can read his payoff for the forecasting task from the payoff table for forecasters, and his payoff from the production (optimization) task from the payoff for quantity decisions.

4. Treatment 4: the LtFE+LtOE Team treatment. In this treatment, there is a forecaster and a production manager in each team. The forecaster knows the history of market prices, and the production manager knows the history of his own production decisions and profits. The market price is determined by the production decisions of all firms in the market. Each subject is paid according to an equal weighted linear combination of the payoff functions used in the LtFE and LtOE treatments, exactly as in treatment 3. Subjects can read the payoff for the forecasting task from the payoff table for forecasters, and the payoff for the production task from the payoff for quantity decisions.

The price in the experiment is restricted to be non–negative, so forecasters are also not allowed to submit negative forecasts. We set 60 as the upper bound for the price prediction, because this is the maximum of the price (when all firms produce 0). The quantity decision should also be non– negative, and we set 20 as the upper bound for the quantity decision as the payoff for the production manager will be negative if he produces more than 20 units when the price is 0.

3.2 Number of Observations

We report results from 8 experimental sessions that were conducted using the CREED laboratory at the University of Amsterdam on April 27-29 and on May 3, 2011. There were a total of 180 experimental subjects who participated in the 8 sessions of this

experiment. No subject participated in more than one treatment or session. Each session involved multiple groups of $N = 6$ or $N = 12$ participants who interacted with one another for 50 periods in one of our four treatments, that is, we adopt a "between subjects" design. We refer to each independent observation, involving $N = 6$ or 12 subjects interacting together for 50 periods under the same treatment conditions as a "market." A summary of the number of markets (observations) and the number of participants per market for each of our four treatments is given in Table 1:

Treatment Number	Number of Firms Per Market	Number of Participants per Market	Total Number of Markets (Observations)	Total Number of Participants
1	6	6	4	24
2	6	6	7	42
3	6	6	7	42
4	6	12	6	72

Table 1: Characteristics of the Experimental Design

Notice that in treatments 1, 2 and 3 we always had 6 subjects (or firms) per market, while in our team treatment 4 we had 12 subjects per market so that each of the 6 "firms" consisted of a pair of players (a "team") who remained matched together for all 50 rounds of the market.

3.3 Theoretical Model

Let D be a nonnegative and monotonically decreasing demand function and let $S_{h,t}$ be the nonnegative supply function of firm h, derived from expected profit maximization. Let $p_{h,t}^e$ be the price forecast made by firm h at period t. The supply function may be rewritten as $S(p_{h,t}^e)$. We assume that all firms have the same supply function. Market demand is assumed to be exogenously given in our experiment. Subjects were exclusively in the role of firms.

The market price is determined by the market clearing condition for a cobweb economy, which is given by:

$$p_t = D^{-1}(\sum_h S_{h,t}) + \epsilon_t, \tag{1}$$

where $\epsilon_t \sim N(0,1)$ is the realization of an i.i.d. price shock in period t.

We assume there are H suppliers, only differing in the way they form expectations. We use a linear demand function $D(p_t) = a - bp_t$, where $a = 63, b = \frac{21}{20}$. We assume each firm has a cost function $c(q) = \frac{Hq^2}{2}$. The expected profit of a firm $\pi_{h,t}^e$ can be defined as:

$$\pi_{h,t}^e = p_{h,t}^e q_{h,t} - c(q_{h,t}) \tag{2}$$

Solving the profit maximization problem yields the optimal supply function for each firm: $S^*(p_{h,t}^e) = \frac{p_{h,t}^e}{H}$. If every firm makes supply decisions optimally, the total supply on the market will coincide with the mean price forecasts, $(\sum_h S^*(p_{h,t}^e) = \overline{p_t^e})$. Substituting this optimal market supply into the market clearing condition (3.3) and noting that the expected value of the noise term is zero, we have that:

$$p_t = \frac{20}{21}(63 - \overline{p_t^e}) \tag{3}$$

Imposing the RE assumption, we find the rational expectations equilibrium (REE) price, $p^* = 30.73$. The optimal supply in this REE is 5.12, and the profit for each firm is 78.7.

Subjects were not informed of the precise demand function as detailed in this section nor were they informed of the total quantity supplied (the quantity decisions of the other $N - 1$ subjects in their market). However, they were told that market demand was decreasing in the market price and that the market price was determined by market clearing, i.e. that supply equals demand -see the Instructions in the Appendix for specific details.

3.4 Computer Interface

Figure 1 provides an illustration of the computer interface that subjects saw in the experiment. The screen was divided into 3 mini pages. In the top mini page, subjects were prompted to submit their decisions, i.e., their price forecast or their quantity production choice. In the bottom left mini page they saw a graph plotting past market prices (the Real Price) and, if they were a forecaster, they also saw their past price forecast history (Your Prediction). Finally, in the bottom right mini page they saw

8

a table containing reporting the history of market prices, as well as their own prior decisions and their period and cumulative payoffs.

The top panel of Figure 1 shows the computer interface that forecasters saw in treatment 4. The computer interface the forecasters saw in treatment 1 is very similar to the one shown for forecasters in treatment 4, except that the history of past performance (points earned) was only for the forecasting task and not from the optimizing task as in treatment 4.

The bottom panel of Figure 1 shows the computer interface the production managers saw in treatment 4. At the start of each period these production managers were told "We wait for your partner to give a forecast." Once the forecaster/team partner has submitted his/her forecast, the production manager was informed of this forecast (as show in the bottom panel of Figure 1) and he or she then entered a quantity decision for the team. The computer interface that subjects see in treatment 2 is very similar to that shown in the bottom panel of Figure 1 except that there is no waiting phase, and the history of past performance is only for the optimization task instead of for both the forecasting and optimization tasks as in treatment 4. The computer interface in treatment 3 is also similar to the one shown in Figure 1, except that there is no waiting phase, and the same subject is asked to first submit a price forecast and then to submit a quantity decision. The history of past performance for treatment 3 is the same as for treatment 4 as the payoff functions are the same in these two treatments.

We note that there were no time constraints on decision-making in any of our treatments. The market price was not determined until all N subjects had submitted their price forecasts and/or quantity production decisions. Each round took no more than 3 minutes to complete (and was often much faster than that).

3.5 Payoffs

Subjects earned points during the experiment that were converted into Euros at the end of the experiment at a known and fixed rate. The payoff function for forecasters is a decreasing function of their prediction error, and was given by:

$$\text{Payoff for Forecasting Task} = \max\{1300 - \frac{1300}{49}(p_t - p_{h,t}^e)^2, 0\} \qquad (4)$$

Your decision for period 8

What is your prediction for the price in period 8? [ ]

Send

Period	Your prediction	The realized price	The Points from Forecasting Task	The Points from Trading Task	The points you earned in the period	The points you have earned so far
7	29	31.97	1065.98	1284.85	1175.41	5359.03
6	33	27.52	503.28	1260.61	881.94	4183.62
5	30	25.43	745.91	1244.58	995.25	3301.68
4	35	22.6	0	1220.15	610.08	2306.43
3	40	19.88	0	1192.16	596.08	1696.35
2	20	41.01	0	1296.03	648.02	1100.27
1	50	0.45	0	904.5	452.25	452.25

Your decision for period 8

Your partner's forecast is 31.
What is the quantity you want to produce in period8? [ ]

Send

Period	Your quantity decision	The realized price	The Points from Forecasting Task	The Points from Trading Task	The points you earned in the period	The points you have earned so far
7	5	31.97	1065.98	1284.85	1175.41	5359.03
6	5.5	27.52	503.28	1260.61	881.94	4183.62
5	6	25.43	745.91	1244.58	995.25	3301.68
4	6.5	22.6	0	1220.15	610.08	2306.43
3	7	19.88	0	1192.16	596.08	1696.35
2	3	41.01	0	1296.03	648.02	1100.27
1	10	0.45	0	904.5	452.25	452.25

Figure 1: The computer interface for forecasters (top) and production managers (bottom).

Notice that subjects earn 0 if their price forecast error is greater than 7.

The payoff function for the production (optimization) task was given by:

$$\text{Payoff from the Production Task} = p_t q_t - c(q_t) + 1200 \qquad (5)$$

Notice that subjects get a 1200 baseline salary, plus the actual profit earned by their firm, which depends on the market determined price, p_t and on the quantity, q_t, chosen by their firm. A firm's profit can be negative, so a subject's payoff can be smaller than 1200. However, we make sure that the maximum loss (the absolute value of negative profit) is 1200, so that each subject's total payoff can never be negative. As the profit for the firm when the market price is at the REE prediction is about 80, the maximum payoff earned by a subject as a forecaster or as a production manager is approximately the same, at around 1300 points.

Subjects in treatment 1 earn the payoff from the forecasting task only. Subjects in treatment 2 earn the payoff from the production task only. Subjects in treatments 3 and 4 each earn the equal weighted average of the payoffs from the forecasting and production tasks. These payoff functions were carefully explained to subjects in the written instructions and presented to subjects as Tables (see the Appendix). At the end of the experiment, subjects were paid 1 Euro for each 2600 points they earned in all 50 rounds of the experiment.

4   Experimental Results

4.1   Aggregate Market Price

Figure 1 plots the average market prices in each treatment against the REE price, $p^* = 30.73$. We see that the average price in all four treatments tracks the fundamental price very well, especially in the later periods of the experiment. So the general tendency for a negative feedback market to converge to REE is not affected too much by the type of task assigned to the market participants. However, the adjustment towards REE at the beginning of the experiment is fastest in treatment 1 and is slowest in treatment 3. The volatility of the market price is also smallest in treatment 1, and largest in treatment 3.

Figure 2: The average market price against the REE price in each of the four treatments.

As a first check on whether prices are converging to the RE prediction, we declare convergence to have occurred in the first period for which the difference between the market price and the REE price is less than 5 and stays below 5 forever after that period. We count the number of periods before convergence in the markets of the different treatments, as reported in Table 2. If there is no convergence according to our criterion, as is the case for 5 markets in treatment 3, then we count the number of periods to convergence as the full sample size of 50 periods. By comparing these numbers, we can see the market price converges faster in treatment 1 than in the other three treatments (the difference is significant at the 5% level according to a Wilcoxon Mann Whitney test using the independent market observations for each treatment). The convergence is faster in treatments 2 and 4 than in treatment 3 (the difference is significant at the 5% level according to a Wilcoxon Mann Whitney test). Treatment 4 converges slightly faster than treatment 2 on average, but that difference is not significant at 5% level according to Wilcoxon Mann Whitney test.

For a second view of convergence, Figure 3 plot the average difference between the market price and the REE price using data from all markets of each treatment.

Figure 3: The distance between the fundamental price and the average of the market prices from all markets of each treatment.

Figure 3 reveals that the difference decreases most rapidly toward zero in treat-

Treatment	Market	Number of Periods before Convergence
Treatment 1	Market 1	3
	Market 2	3
	Market 3	4
	Market 4	1
	Mean	2.75
	Median	3
Treatment 2	Market 1	17
	Market 2	33
	Market 3	13
	Market 4	12
	Market 5	11
	Market 6	4
	Market 7	28
	Mean	14.43
	Median	13
Treatment 3	Market 1	50
	Market 2	50
	Market 3	35
	Market 4	3
	Market 5	50
	Market 6	50
	Market 7	50
	Mean	42.29
	Median	50
Treatment 4	Market 1	36
	Market 2	10
	Market 3	13
	Market 4	25
	Market 5	6
	Market 6	10
	Mean	10.67
	Median	10

Table 2: The number of periods before convergence for each market.

ment 1 (circles), and most slowly in treatment 3 (triangles). Treatment 2 (squares) and treatment 4 are very similar to one another.

Finally we can test for convergence econometrically using a method suggested by Duffy (2008). The following linear equation is estimated:

$$p_{j,t} = \lambda_j p_{j,t-1} + \mu_j + \epsilon_{j,t} \tag{6}$$

This linear equation is stable if λ is smaller than 1, and has a long term equilibrium level $\frac{\hat{\mu}_j}{1-\hat{\lambda}_j}$. For a market j, we declare a *weak convergence* if reject $\hat{\lambda}_j \geq 1$ at 5% level, and a *strong convergence* if we can not reject $\frac{\hat{\mu}_j}{1-\hat{\lambda}_j} = 30.73$ at 5% level. The estimation results are shown in the appendix. We see from the results that:

1. All markets in all the treatments satisfy *weak convergence.*

2. All markets in treatment 1 and 2 satisfy *strong convergence.* All but one market in treatment 4 satisfy strong convergence. The equilibrium price in the one market of Treatment 4 that does not satisfy strong convergence is not very different from the REE ($\frac{\hat{\mu}_j}{1-\hat{\lambda}_j} = 32.08$). Only 2 out of 7 markets in treatment 3 satisfy strong convergence.

We see a large difference between treatment 3 and the other three treatments. The difference between treatments 3 and 4 in particular suggests that teamwork and specialization may help participants to make optimal decisions.

4.2 Individual–Level Decisions

We have seen that aggregate market price tracks the REE well in many markets. It is of interest to consider whether decisions at the individual level are also consistent with RE predictions. The empirical cumulative distribution function (CDF) of individual forecasts and optimization decisions is shown in Figure 4 using pooled data from all markets of the various treatments. Under rational expectation the CDF should be a step function switching from 0 to 1 at the RE price or quantity.

Figure 4 reveals that there is some heterogeneity in individual decisions across treatments with the largest departures from RE predictions occurring in treatment 3,
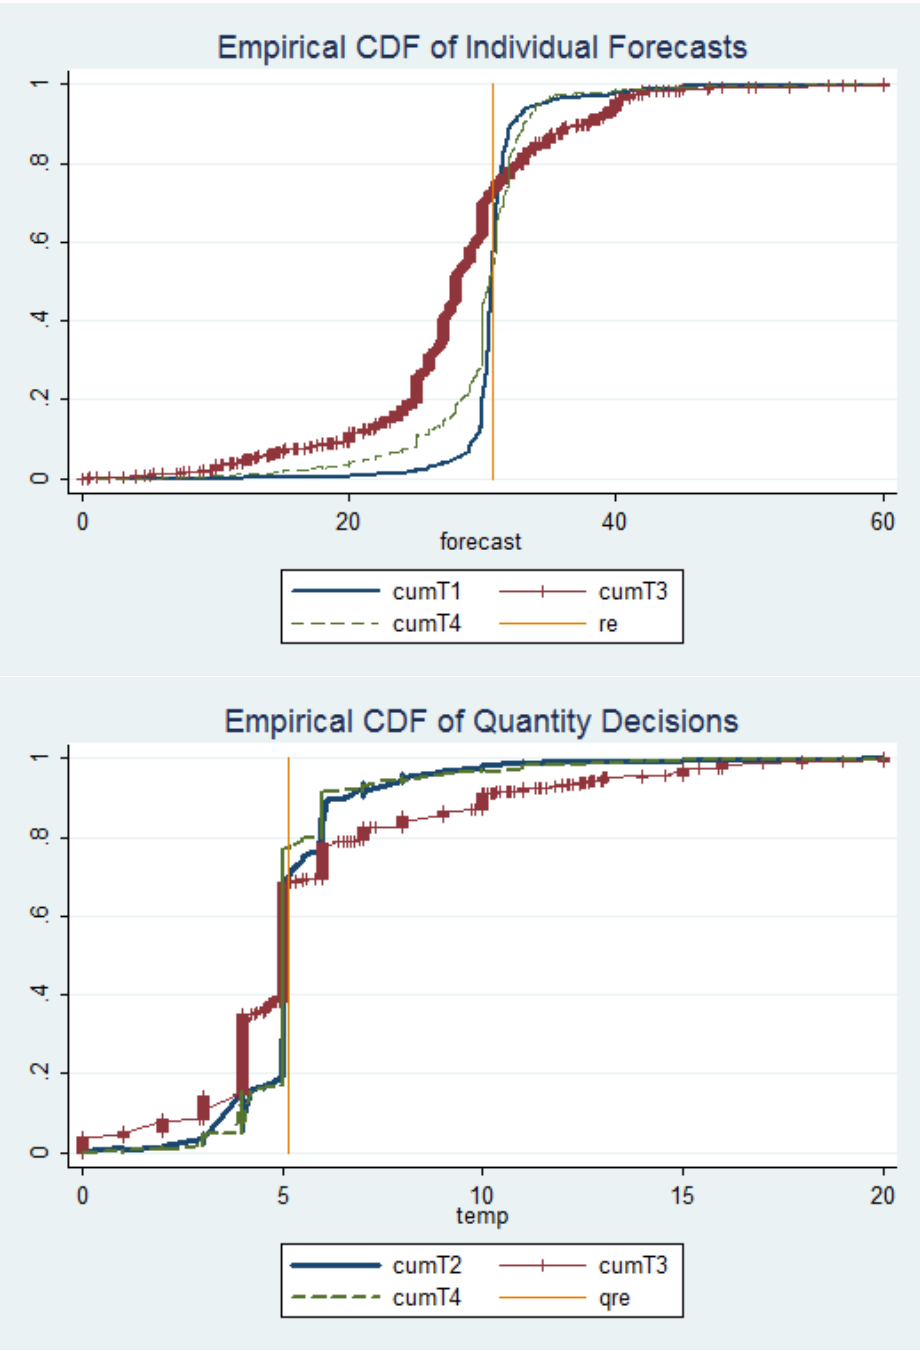
Figure 4: The empirical cdf of individual forecasting and quantity decisions.

a finding that is consistent with our findings using aggregate measures of prices and quantities.

Using the distribution of individual forecasts for the three treatments involving forecasting, we perform a one-sample Kolmogorov-Smirnov test of whether the distribution of individual forecasts is significantly different from the RE prediction, $p^* = 30.73$ (at the 5% level). We can reject the null hypothesis of no difference for all three treatments. The top panel of Figure 4 suggests that the distribution of individual forecasts is similar in treatments 1 and 4, while treatment 3 looks very different. For confirmation we perform a two-sample Kolmogorov-Smirnov test on whether the distribution of individual forecasts is the same between each of pairing of these three treatments, and we find that each treatment is significantly different from the others (at 5% level). Indeed, the ordering is such that Treatment 1 is closest to the RE price prediction, treatment 3 is furthest and treatment 4 is intermediate.

For the distribution of individual quantity decisions, we also perform a one-sample Kolmogorov-Smirnov test on whether the distribution of individual quantity decisions is significantly different from the RE prediction that all firms produce 5.12 units (at 5% level). We can again reject the null hypothesis of no difference for all three treatments involving quantity decisions. The lower panel of Figure 4 suggests that the distribution of individual quantity decisions is similar in treatments 2 and 4, while treatment 3 looks very different. We again perform a two–sample Kolmogorov-Smirnov test on whether the distribution of individual quantity decisions is the same between each pairing of the three treatments. The test indicate that there is no significant difference in the distribution of quantity decisions between treatments 2 and 4, and but there is a significant difference between treatment 3 and the other two treatments (at 5% level). In particular, there is much greater heterogeneity in the quantity decisions of treatment 3 as compared with either treatments 2 and 4.

4.3   Variance of the Market Price and M.S.D from REE

The variance of the market price and the mean squared deviation (M.S.D.) of prices from the REE in our experiment are shown in Table 3. We calculate these numbers for the whole experiment and the first and second 25 periods. Both measures follow the same order: $Treatment3 > Treatment2 > Treatment4 > Treatment1$, although

the difference between Treatments 1, 2 and 4 are very small in the second 25 periods, when the markets in these three treatments have converged to REE. This finding basically confirms our conjecture that "two heads are better than one" (in finding the REE). Treatment 3 generates the largest variance and deviance from REE probably because subjects are a little overloaded by the need to complete two tasks at the same time. Treatment 4 improves upon Treatment 3 because specialization promotes efficiency. Treatment 4 not only yields more frequent convergence to REE but it also takes no more time to complete compared with sessions of Treatment 3: a Treatment 4 session took between 1 hour and 20 minutes to 2 hours to complete while the two Treatment 3 sessions took 1 hour and 40 minutes and 2 hours, respectively, to complete.

Treatment	Market	Period 1-50		Period 1-25		Period 26-50	
		Variance	MSD from REE	Variance	MSD from REE	Variance	MSD from REE
Treatment 1	Market 1	8.4639	8.3246	15.9253	15.498	1.1862	1.1512
	Market 2	4.5009	4.4123	8.0549	7.7576	1.1042	1.0669
	Market 3	6.0093	5.8903	10.5023	10.2533	1.4662	1.5273
	Market 4	4.0495	3.9687	5.9651	5.7271	2.2995	2.2104
	Average	5.7559	5.649	10.1119	9.809	1.514	1.489
Treatment 2	Market 1	37.4148	42.1834	57.3784	80.1954	4.2428	4.1714
	Market 2	43.1768	47.091	75.7162	88.4803	5.7746	5.7017
	Market 3	6.2406	6.3842	9.6834	9.4855	3.0436	3.2829
	Market 4	30.6806	30.3493	49.8641	54.9473	3.3323	5.7514
	Market 5	24.5577	24.3453	44.5447	44.7759	3.9408	3.9148
	Market 6	21.3695	20.9732	40.6862	39.4943	1.4866	2.4521
	Market 7	11.9966	11.7587	19.3881	18.9612	4.2627	4.5562
	Average	28.9441	31.8862	47.5927	59.387	4.3537	4.3853
Treatment 3	Market 1	26.1905	27.9131	39.5377	38.5528	12.8353	17.2734
	Market 2	48.9827	65.194	54.7201	53.7902	26.2326	76.5979
	Market 3	76.5335	125.0443	117.0166	236.1772	4.9931	13.9114
	Market 4	26.9917	29.4857	51.0947	57.3338	1.3238	1.6376
	Market 5	20.3711	48.2724	24.2351	74.8411	10.1406	21.7038
	Market 6	6.9515	15.6452	12.1408	19.9058	2.0312	11.3847
	Market 7	60.2049	147.2105	63.8626	271.2218	4.9447	23.1991
	Average	44.6746	61.9093	65.5922	96.4635	11.3462	27.3551
Treatment 4	Market 1	14.3269	15.3855	23.0514	22.1382	3.8329	8.6329
	Market 2	17.2713	17.2323	30.9771	32.2878	2.0178	2.1768
	Market 3	18.4874	19.2729	25.9906	32.3755	6.0827	6.1703
	Market 4	36.5533	40.4508	61.4327	78.6819	2.2928	2.2197
	Market 5	9.0801	9.92666	13.8618	18.0667	1.8365	1.7866
	Market 6	28.9092	29.3816	45.1668	27.1176	3.1776	3.8889
	Average	20.7714	21.9416	33.4134	35.1113	3.2067	4.1459

Table 3: The MSD from REE and variance of price for each market.

4.4 Efficiency

We compare subjects' earnings in the experiment to the hypothetical case where all subjects play according to the REE predictions in all 50 periods. Subjects can earn 1300 points per period for the forecasting task when they play according to REE because they make no prediction errors, which means they earn 0.5 Euro each period,

and 25 Euros for all 50 periods. The profits they can earn for the production task is 1278.7 points per period when they play according to the REE, which means they earn 0.4918 Euro per period, and 24.59 Euros for 50 periods. We use the ratio of actual to hypothetical REE payoffs as a measure of efficiency. This measure can be greater than 100 percent in treatments with production decisions, because subjects can earn more by producing a little less than the REE prediction. These efficiency ratios reported in Table 4 are generally very high (more than 80%) in all four treatments. We see the ranking of efficiency level in all 50 periods is $Treatment2 > Treamtnet4 > Treament1 > Treatment3$, the ranking in the second 25 periods is $Treatment2 > Treatment1 > Treatment4 > Treamemt3$. Only the difference between efficiency in treatment 2 and other treatments is significant at 5% level according to Wilcoxon Mann Whitney test. The differences between the efficiency level in other treatments are not significant.

		Period 1-50		Period 1-25		Period 26-50	
Treatment	Market	average earning	efficiency	average earning	efficiency	average earning	efficiency
Treatment 1	Market 1	20.44	89.27%	8.45	67.58%	11.99	95.94%
	Market 2	21.57	86.27%	9.47	75.79%	12.09	96.74%
	Market 3	21.50	86.00%	9.59	76.69%	11.91	95.31%
	Market 4	21.83	87.33%	10.50	84.03%	11.33	90.64%
	Average	21.80	87.22%	9.50	76.02%	12.30	98.41%
Treatment 2	Market 1	24.45	99.43%	11.64	94.70%	12.81	104.16%
	Market 2	23.98	97.53%	11.73	95.43%	12.25	99.64%
	Market 3	23.95	97.40%	12.19	99.18%	11.76	95.61%
	Market 4	24.47	99.50%	11.90	96.81%	12.56	102.19%
	Market 5	24.43	99.36%	12.03	97.85%	12.40	100.88%
	Market 6	24.33	98.96%	12.09	98.35%	12.24	99.56%
	Market 7	24.25	98.62%	12.14	98.71%	12.11	98.53%
	Average	24.27	98.69%	11.96	97.29%	12.30	100.08%
Treatment 3	Market 1	22.10	89.11%	9.68	78.07%	12.42	100.16%
	Market 2	18.57	74.87%	9.42	75.96%	9.15	73.78%
	Market 3	20.63	83.20%	7.08	57.07%	13.56	109.33%
	Market 4	21.18	85.42%	10.53	84.93%	10.65	85.91%
	Market 5	19.12	77.08%	9.06	73.04%	10.06	81.13%
	Market 6	22.78	91.87%	10.93	88.16%	11.85	95.58%
	Market 7	19.27	77.69%	8.39	67.67%	10.88	87.71%
	Average	20.52	82.75%	9.30	74.98%	11.22	90.51%
Treatment 4	Market 1	22.10	89.11%	10.07	81.22%	12.03	97.00%
	Market 2	21.80	87.90%	10.14	81.81%	11.66	93.99%
	Market 3	21.08	85.01%	9.36	75.48%	11.72	94.55%
	Market 4	20.60	83.06%	9.16	73.83%	11.44	92.30%
	Market 5	22.32	89.99%	10.09	81.38%	12.23	98.60%
	Market 6	22.13	89.25%	10.65	85.85%	11.49	92.64%
	Average	21.67	87.39%	9.91	79.93%	11.76	94.85%

Table 4: The efficiency for each market.

However, as the payoff functions for the forecasting and optimizing tasks were different, it is difficult to draw conclusions from the reported efficiency ratios across some of the treatments. One way to make the results more comparable is to examine implicit production decisions in treatment 1 and implicit price forecasts in treatment 2, and then calculate the implicit efficiency level of the production decisions in treat-

ment 1, or the implicit efficiency of forecasting task in treatment 2. For treatment 1, it is straightforward that the firm will produce as much as one sixth of the prediction, and the profit of the firm can be calculated accordingly. For treatment 2, we can assume that the subjects always make production decisions that are conditionally optimal for their implicit forecast, and therefore we calculate their implicit forecast as six times their quantity decision. Given these numbers we can calculate the efficiency level for both the forecasting and optimizing tasks for all four treatments in a consistent manner and we can define an efficiency index for all the treatments as the mean of the efficiency levels for the two tasks. This index, which allows for efficiency comparisons across the four treatments, is reported in Table 5.

Table 5 reveals that efficiency level for the implicit optimizing task in treatment 1 is as high as the comparable efficiency level of the optimizing task in treatment 2, and sometimes exceeds 100% in the second 25 rounds of the experiment. This suggests that the higher efficiency level reported for treatment 2 as compared with treatment 1 may be an artifact of the payoff function differences. Subjects performing the optimization task benefit from small, positive random shocks which result in higher market price. By contrast, both positive and negative shocks are equally penalizing for subjects performing the prediction task as both types of shocks lead to higher prediction errors.

Table 5 also reveals that the ranking of the overall efficiency index is $Treatment1 > Treatment4 > Treatment3 > Treatment2$. This ranking for the forecasting task is the same as the overall ranking, and the ranking for the optimizing task is $Treatment1 > Treatment2 > Treatment4 > Treatment3$. We conducted a Wilcoxon Mann Whitney test on market level efficiency for the two tasks and on the efficiency index for period 1-50. The result suggests that the efficiency level is significantly higher in treatment 1 in both tasks as well as for the efficiency index as compared with all other treatments. The efficiency for forecasting is significantly lower in treatment 2 as compared with the other treatments[1], but there are no other significant difference in pairwise comparisons between treatments. As we will see in later sections, subjects in treatments 3 and 4 (especially treatment 3) do not make perfect production decisions given their forecasts. This result suggests that it does not cause a lot of

---

[1]This result may be due to our assumption that the implicit forecast is 6 times the quantity , or the fact that the subjects do not act conditionally optimally to their implicit forecast (produce exactly one sixth of the implicit forecast).

differences in the efficiency if subjects are boundedly rational in optimization tasks. But if subjects are not fully rational with regard to the optimization task, this may lead to inaccurate forecasts as well resulting in lower forecast efficiency.

This result also suggests that caution needs to be taken in explaining a high efficiency level in learning to optimize experiments, because even if the efficiency says there is no problem in the optimization task, the implicit forecasts made by the subjects may be far from rational. In this case, the team design with specialized roles provides a clearer view of the decision process in each task, and restores the efficiency level in forecasting.

Treatment	Periods	Avg. Payoff Forecasting	Avg. Payoff Optimizing	Efficiency Forecasting	Efficiency Optimization	Efficiency Index
Treatment 1	Period 1-50	21.80	24.55	87.22%	99.85%	93.54%
	Period 1-25	9.50	12.26	76.02%	99.68%	87.85%
	Period 26-50	11.99	12.30	98.41%	100.03%	99.22%
Treatment 2	Period 1-50	14.45	24.27	57.79%	98.69%	78.24%
	Period 1-25	5.78	11.96	46.24%	97.29%	71.76%
	Period 26-50	8.67	12.30	69.36%	100.08%	88.68%
Treatment 3	Period 1-50	17.63	23.39	70.53%	95.14%	82.84%
	Period 1-25	7.19	11.41	57.48%	92.81%	75.15%
	Period 26-50	10.45	11.98	83.57%	97.47%	90.52%
Treatment 4	Period 1-50	19.08	24.27	76.31%	98.68%	87.50%
	Period 1-25	7.87	11.96	62.93%	97.24%	80.09%
	Period 26-50	11.21	12.31	89.69%	100.12%	94.91%

Table 5: The breakdown of efficiency into forecasting and optimizing tasks.

4.5   Individual Forecast

The left panel of Figure 5 shows the average individual price forecasts in treatments 1, 3 and 4 against the REE. We can see that Treatment 1 converges fastest, followed by treatment 4, and that treatment 3 is the slowest to converge. The right panel of Figure 5 shows the average variance of individual forecasts in treatments 1, 3 and 4. We observe that heterogeneity of supply decisions is greatest in treatment 3, and there is not much difference between treatments 1 and 4.

Prior experimental work (Heemeijer et al, 2009) suggests that subjects tend to use simple heuristics in learning to forecast experiments. Two natural candidates they often use in negative feedback markets are adaptive expectations:

$$p_{i,t+1}^e = p_{i,t}^e + \lambda(p_t - p_{i,t}^e), \tag{7}$$
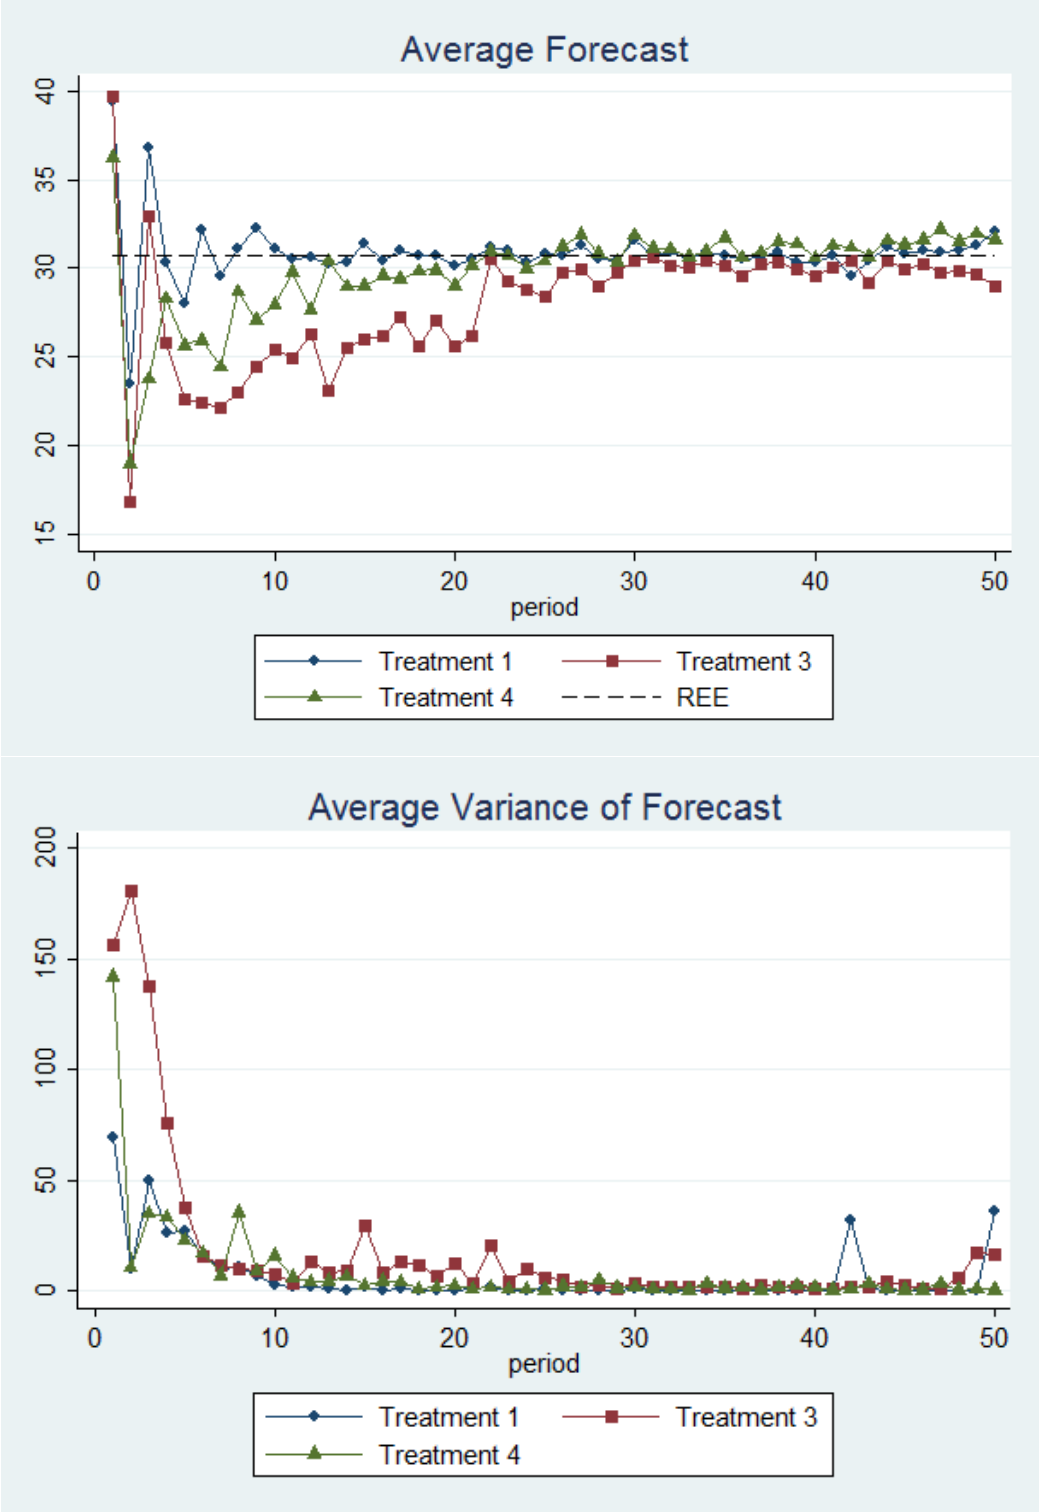
Figure 5: The upper panel shows the average individual forecasts in Treatment 1, 3 and 4. The bottom panel shows the average of the group variance of individual forecasts in Treatments 1, 3 and 4.

and trend extrapolation rules:

$$p^e_{i,t+1} = p_t + \gamma(p_t - p_{t-1}). \tag{8}$$

The estimated γ is usually negative in the market setting we consider, so we use the term "contrarian rule" to differentiate this rule from the trend–following rule where γ is positive. We estimate these two types of rules. We examine the estimation results after performing the estimation. We call an estimation successful if it generates coefficient estimates that are significant at the 5% level, and there is no serial correlation. If both rules are successful for the same individual, we compare the R^2 of each estimated model and characterize the individual as following the rule with larger R^2. It turns out that more than 75% of subjects can be categorized by either rule in all treatments. The distribution of individual subjects over the types of forecasting rules is shown in Table 6 and Figure 6, while the Tables in Appendix B show the estimation results for the subjects who can be successfully identified with one rule:

Treatment	Adaptive	Contrarian	Neither
Treatment 1	66.67%	12.50%	20.83%
Treatment 3	52.38%	23.81%	23.81%
Treatment 4	50.00%	27.78%	22.22%

Table 6: The fraction of subjects who are characterized by one type of forecasting rule or neither.

Generally speaking, the distribution of subjects over the different rules is not very different across the three treatments. In all three treatments 50% or more subjects can be categorized by the adaptive rule. There are relatively more subjects using the contrarian rule in Treatments 3 and 4 as compared with Treatment 1. If we relate the result here to the stability of the markets, it seems the market price is most stable when there are overwhelmingly more people using the adaptive rule.

4.6 Individual Supply Decision

4.6.1 Descriptive Statistics

The average supplies in treatment 2, 3 and 4 are plotted against the REE supply in the top panel of Figure 7. As with prices, we see that quantity in treatment 3
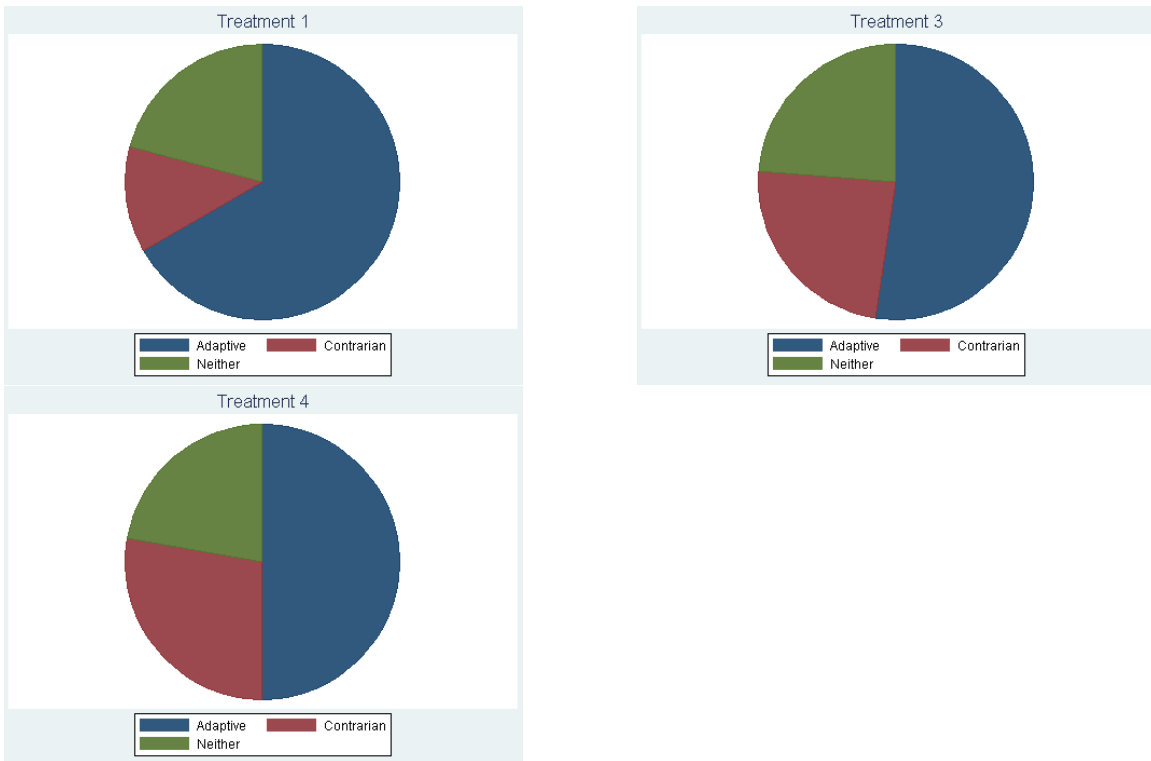
Figure 6: The fraction of subjects who are characterized by one type of forecasting rule or neither in treatment 1 (top left), 3 (top right) and 4 (bottom).
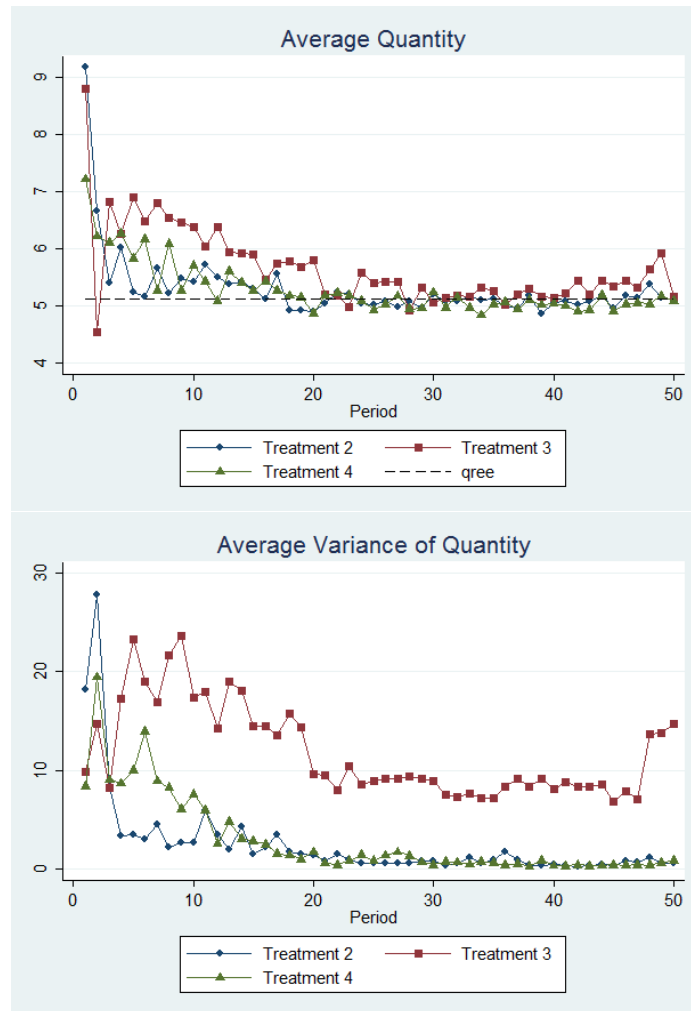
Figure 7: Upper panel: the average individual supply in Treatments 1, 3 and 4. Bottom panel: the average variance of individual supply in Treatments 1, 3 and 4.

25

converges towards the REE level in a rather sluggish manner, and there is not much difference in the average quantity supplied over time between treatments 2 and 4. The bottom panel of Figure 7 shows the average variance of supply in each treatment. We again observe that the heterogeneity of supply decisions is greatest in treatment 3, and there is not much difference between treatments 2 and 4.

4.6.2   Conditional Optimality of Production Decision

Figure 8: The average distance between actual supply and the conditionally optimal supply in Treatments 3 and 4.

If the production manager acts optimally with respect to the forecaster's forecast, he should decide to supply 1/6 of the firm's price prediction. Do production managers make decisions in this manner? Figure 8 shows the average difference between the supply chosen by the production manager and the optimal supply given his own or his paired forecaster's forecast in treatments 3 and 4, respectively. If the production managers make decisions optimally, this difference should be zero.

Figure 8 reveals that the production managers in treatment 4 on average make

supply decisions that are closer to the conditionally optimal quantity choie given their partners' price forecast. This also indicates that the production managers generally trust their partners. Although trust should not be an issue in treatment 3, where the forecast and supply decisions are made by the *same* person, we observe that subjects in treatment 3 generally fail to make production decisions that are optimal given their own price forecasts. We suspect that the reason for this difference in treatment 3 as compared with treatment 4 is that doing both tasks (as is required in treatment 3) is indeed very difficult for a single individual, that is, there is a greater cognitive load in treatment 3 as compared with treatment 4.

4.6.3   Estimation of Supply Strategies

We are interested in the possible cause of the deviation of managers' supply decisions from the conditionally optimal decision given price predictions in treatments 3 and 4. To address this issue further, we estimate a simple production strategy specification:

$$q_t = c_0 + c_1 p_t^e. \tag{9}$$

If the production manager is a conditional optimizer, the regression result should yield that $c_0 = 0$, $c_1 = 1/6$ for each individual firm. There are certainly many other independent variables that could also be included in the specification of the production decision. As the production managers in Treatment 4 do not see information such as the price forecast history, and the forecaster and production managers in Treatment 3 should have incorporated all other information into the predictions they made for themselves, this equation is most suitable for comparing the two treatments. We discard the estimations with serial correlation, leaving 13 (out of 42) successful estimations for Treatment 3, and 18 (out of 36) successful estimations for Treatment 4. The results can be found in the Appendix C.

We can classify subjects in their role as production managers according to three types:

1. Unconditional supply, if c_0 is significantly different from 0 at 5%, and c_1 is not. This means the subject probably just choose to supply a constant number of goods.

2. Conditional optimal supply, if c_1 is significantly different from 0 at 5%, c_0 is not significant, and the null hypothesis $c_1 = 1/6$ can not be rejected at 5% level. This means that the subject choose to supply the conditionally optimal quantity for the given price forecast.

3. Hybrid strategy, if both c_0 and c_1 are significant. This means that the subject probably choose a constant as a psychological anchor, and adjusted it a little for different expected price levels.

It turns out all the successful estimations can be classified in this way. The graph below shows the shares of the three different types of production strategies. We use C to denote the use of the constant supply strategy, O to denote use of the conditionally optimal supply strategy and H to denote use of the hybrid strategy. There are 4 subjects using the constant supply strategy, 2 using the conditionally optimal supply strategy and 7 using a hybrid strategy in Treatment 3. There is 1 subject using the constant supply strategy, 9 using the conditionally optimal supply strategy and 8 using the hybrid strategy in Treatment 4. Thus, about half of all subjects (for whom we could identify a supply strategy) use a hybrid strategy in both treatments. For the remaining population, a majority uses the constant supply strategy in Treatment 3 while in Treatment 4, the majority uses the conditionally optimal strategy. This result suggests that subjects do behave in a systematically different manner between treatments 3 and 4. In treatment 3, many subjects choose to use the constant supply strategy which requires minimal cognitive cost, but which destabilizes the market when they choose the wrong (usually too high) quantity. In treatment 4, subjects in the production manager role trust their partners's forecasts to a reasonable degree, which facilitates their greater use of the conditionally optimal strategy.

5   Conclusion

Rational Expectations (RE) macro models have two crucial dimensions: 1) Agents correctly forecast future prices (no systematic mistakes). 2) Given these expectations, agents solve optimization problems and their optimal production (or consumption or trading) strategies then determine actual price realizations, that is, there is

Figure 9: This graph shows the share of different types of estimated production strategies treatment 3 and 4. We use C to denote the use of the constant supply strategy, O to denote use of the conditionally optimal supply strategy and H to denote use of the hybrid strategy.

belief-outcome interaction. These two dimensions have been previously addressed separately in learning to forecast experiments (LtFE) and learning to optimize experiments (LtOE). In this paper we design comparable LtFE and LtOE treatments for the same model, and add two additional treatments where subjects perform both tasks either independently or as members of a team. Our paper shows that all the approaches give the same qualitative, long-run result, namely convergence to the REE in the context of cobweb economy.

Among all the treatments, the LtFE treatment converges more quickly and reliably than the other three treatments. We suspect this is because the forecasting task is considerably easier than the optimizing task and therefore behavior in LtFE studies should be regarded as an upper bound on the rationality that can be achieved in a laboratory experimental evaluation of RE models. The estimation of individual forecast rules suggests that there is not much difference in the price prediction strategies subjects use across the different treatments. However, estimation of the supply strategies suggests that there are differences in strategies used between treatments 3 and 4. The current macroeconomic literature usually only takes bounded rationality in forecasting into the theoretical models, and the implication of our result for future theoretical work is that it may be worthwhile to also take bounded rationality in optimization into account.

We also find evidence in support of the notion that "two heads are better than one" in the sense that behavior in treatment 4 is more rational than that in treatment 3, even in the aspect of consistency (how close the production decision is to the conditionally optimal decision for the given price forecast). This finding also goes along with the real life observation that large financial institutes usually have separate forecasting and trading departments, and rarely let one department perform the task of the other.

In future research it would be desirable to consider experiments with comparable LtFE and LtOE treatments in different market contexts from the one considered here. In particular it would be of interest to apply our same approach to a market with positive expectation feedback, where prices usually doe not converge, at least in the learning to forecast experiments that have been used to date in such environments.

Reference

1. Adam, K., "Experimental Evidence on the Persistence of Output and Inflation," Economic Journal, 117, 2007, 603-635.

2. Arifovic, J., "The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies", Journal of Political Economy, 104 (3), 510-541.

3. Bao, T., Hommes, C., Sonnemans, J., Tuinstra, J., "Individual Expectations, Limited Rationality and Aggregate Outcomes," Working Paper University of Amsterdam, 2010.

4. Bornstein, G., and I. Yaniv, "Individual and Group Behavior in the Ultimatum Game: Are Groups more 'Rational' Players?," Experimental Economics, 1, 101-108.

5. Cason, T., and V. Mui, "A Laboratory Study of Group Polarization in the Team Dictator Game," Economic Journal, 107, 1465-1483.

6. Cox, J. C., "Trust, Reciprocity, and Other-Regarding Preferences: Groups vs. Individuals and Males vs. Females," in (R. Zwick, and A. Rapoport, eds.), Advances in Experimental Business Research, pp. 331-50, Dordrecht: Kluwer Academic Publishers.

7. Crockett, S., and J. Duffy, "A Dynamic General Equilibrium Approach to Asset Pricing Experiments," Working Paper.

8. Duffy, J. "Macroeconomics: A Survey of Laboratory Research", working paper, University of Pittsburgh, 2008.

9. Fehr, E. and J.-R. Tyran, "Limited Rationality and Strategic Interaction: the Impact of the Strategic Environment on Nominal Inertia", Econometrica 76 (2008), 353-394.

10. Gode, D.K., and S. Sunder, "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality," Journal of Political Economy, 101(1), 119-137.

11. Heemeijer, P., Hommes, C., Sonnemans J., and J. Tuinstra, "Price Stability and Volatility in Markets with Positive and Negative Expectations Feedback:

An Experimental Investigation," Journal of Economic Dynamics and Control 33 (2009), 1052-1072

12. Hommes, C.H., Sonnemans, J.H., Tuinstra, J. and van de Velden, H., "Coordination of Expectations in Asset Pricing Experiments," Review of Financial Studies, 2005, 18 (3), 955-980.

13. Hommes, C.H., Sonnemans, J.H., Tuinstra, J.,and van de Velden, H., "Learning in Cobweb Experiments", Macroeconomic Dynamics, 11, 2007, 8-33.

14. Huck, S., Norman, H.T., and J. Oechssler, "Learning in Cournot Oligopoly–An Experiment," Economic Journal, 109 (454), 80-95.

15. Kocher, M., and M. Sutter, "When the 'Decision-maker' Matters: Individual versus Team Behavior in Experimental 'Beauty Contest' Games," Economic Journal, 115, 200-223.

16. Lim S.S., E.C. Prescott and S. Sunder, "Stationary Solution to the Overlapping Generations Model of Fiat Money: Experimental Evidence," Empirical Economics, 19, 1994, 255-77.

17. Marimon, R. and S. Sunder, "Indeterminacy of Equilibria in a Hyperinflationary World: Experimental Evidence," Econometrica, 61(5), 1993, 1073-1107.

18. Marimon, R. and S. Sunder, "Expectations and Learning under Alternative Monetary Regimes: An Experimental Approach," Economic Theory, 4, 1994, 131-62.

19. Marimon, R., Spear, S. E., and S. Sunder, "Expectationally Driven Market Volatility: An Experimental Study," Journal of Economic Theory, 61(1), 1993, 74-103.

20. Muth, J.F., "Rational Expectations and the Theory of Price Movements," Econometrica, 29(3), 1961, 315-335.

21. Offerman, T., Potters, J., and J. Sonnemans, "Imitation and Belief Learning in an Oligopoly Experiment," Review of Economic Studies 69, 973-97.

22. Smith, V. L., Suchanek. G. L. and A. W. Williams, "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets," Econometrica, 56(5), 1988, 1119-1151.

23. Vega-Redondo, F., "The Evolution of Walrasian Behavior," Econometrica, 65(2), 375-384.

A   Appendix

A.1   Experimental Instructions

A.1.1   Instruction for the Forecaster

(Normal for T1, in () for T4)

General information

In this experiment you participate in a market. Your role in the market is a professional Forecaster for a large firm, and the firm is a major Producer of one product sold in the market. In each period the firm asks you to make a prediction of the market price of the product. The price should be predicted one period ahead, since producing the good takes some time. You are going to advise the firm for 50 successive time periods. (At the beginning of the experiment, you and another participant a Production Manager as your partner are assigned into a team and will keep cooperating throughout the experiment.) In each period you have to make a prediction for the price in the next period, and your firm (partner) makes a decision about the quantity of the good the firm should produce. Your forecast is the only information the firm (production manager) has on the future market price. The more accurate your prediction is, the better the quality of your firm's (partner's) decision will be, and the more profit your firm can earn. In each period, (each of) you (and your partner) will get a payoff based on the accuracy of your prediction (and the quality of production decisions).

The information you can refer to consists of a plot of the past prices and your predictions, a table containing the history of your past forecasts, [production decisions] and payoff (of your team) in forecasting [(and production)] tasks. (You partner sees a plot of the past price, a table containing the history of his/her supply decisions and the payoff of your team in forecasting and production tasks.)

About the price determination

The price is determined by the market clearing condition, meaning that it will be such that the supply equals demand.

The supply on the market is determined by the production decision of the producers. There are several large producers on this market and each of them is advised by a forecaster like you. Usually, higher price predictions make a firm to produce a larger quantity, which increases the supply and vise versa. Total supply is largely determined by the sum of the individual supplies of these producers, although there may be small random fluctuations caused by transportation delay or other reasons.

The size of the demand depends upon the price. When the price goes up, the demand will go down.

About your job

Your only task in this experiment is to predict the market price in each time period as accurately as possible. At the beginning of the experiment you are asked to give a prediction for the price period 1. When all forecasters have submitted their predictions for the first period, the firms (production managers) will determine the quantity to supply, and the market price for period 1 will be determined and made public to all forecasters. Based the accuracy of your prediction in period 1, your earnings will be calculated.

Subsequently, you are asked to enter your prediction for period 2. When all participants have submitted their prediction [(and production decisions)] for the second period, the market price for that period will be made public and your earnings will be calculated, and so on, for all 50 consecutive periods. The information you can refer to consists of all previous prices, your predictions and earnings.

About your payoff

You payoff depends on the (both) performance of your forecasting task (and your partner's production decision task. Each of you and your partner will get one half of the payoff for the forecasting task and one half of the payoff for the quantity production task). The payoff for the forecasting task depends on the accuracy of your predictions. The earnings shown on the computer screen will be in terms of points. The maximum possible points you can make for the forecasting task is 1300 for each period, and the larger your prediction error is, the fewer points you can make. You will earn 0 points if your prediction error is larger than 7. There is a Payoff Table on your table, which shows the points you can earn for different prediction errors. Your action will have no impact on the payoff from the production task.

We will pay you in cash at the end of the experiment based on the points you earned. You earn 1 euro for each 2600 points you make.

A.1.2   Instruction for the Production Manager

(Normal for T2, in () for T4, and [] for T3)

General information

In this experiment you participate in a market. Your role in the market is a Production Manager of a large firm, and the firm is a major Producer of one product sold in the market. In each period the firm asks you to make a decision on the quantity your firm will supply to the market. You are going to play this role for 50 successive time periods.

(At the beginning of the experiment, you and another participant, a Forecaster as your partner are assigned into a team and will keep cooperating throughout the experiment. In each period you will receive a prediction for the price in this period from your partner, and make a decision about how much goods your firm should produce.) The better the quality of your decision is, the more profit your firm can earn.

The information you can refer to consists of a plot of the past prices, a table containing the history of your past decisions and the payoff (of your team) in (forecasting and) production tasks. (You partner sees a plot of the past price and his/her own forecasts, a table containing the history of his/her past forecasts and the payoff of your team in forecasting and production tasks.)

About the price determination

The price is determined by the market clearing condition, meaning that it will be such that the supply equals demand.

The supply on the market is determined by the production decision of the producers. Usually, higher price predictions make a firm to produce a larger quantity, which increases the supply and vise versa. Total supply is largely determined by the sum of the individual supplies of these producers, although there may be small random fluctuations caused by transportation delay or other reasons.

The size of the demand depends upon the price. When the price goes up, the demand will go down.

About your job

Your task in this experiment is to [make a prediction on the market price and] decide the quantity the firm will supply. At the beginning of the experiment (you receive the forecaster's prediction for the price period 1. When all forecasters have submitted their predictions for the first period, the decision makers including) [you make a prediction of the market price and] you determine the quantity to supply for period 1, and when all the participants submitted their [forecasts and] decisions, the market price for period 1 will be determined and made public to all forecasters. Based on [the accuracy of your prediction and] the profit of your firm in period 1, your earnings in the first period will be calculated.

Subsequently, (you receive the forecaster's prediction for period 2, and) you make [the prediction and] the production decisions for the second period. When all participants have submitted [their prediction and] production decisions for the second period, the market price for that period will be calculated and made public and your earnings will be calculated, and so on, for all 50 consecutive periods.

About your payoff

You payoff depends on the ([both]) performance of your production task ([and your] partner's [forecasting task.) Each of [you] and your partner [will get one half of the payoff for the forecasting task and one half of the payoff for the production task]).The payoff for the production task is the same as the profit of the firm. The earnings shown on the computer screen will be in terms of points. You do not need to calculate your payoff yourself. There is a Payoff Table for Production Task on your table, which shows the points you can earn for a given market price in the row (, [for which you could use your] partner's [forecast as a proxy) and your production decision in the column. You payoff from the forecasting task is decreasing in your prediction error, and you can also refer to the other payoff table to see how much you can earn for a given prediction error.] If you really want to know how the numbers in the payoff table is calculated you can read the last part of the instruction, which you can skip

otherwise.

We will pay you in cash at the end of the experiment based on the points you earned. You earn 1 euro for each 2600 points you make.

The equation that calculates the payoff for the production task

The payoff for production task can be written as the following equation:

$$\text{Payoff from the Production Task} = p_t q_t - c(q_t) + 1200$$

Where p_t is the market price of this good, and you can use your partner's prediction as a proxy. q_t is the amount of product you decide to let the firm produce. $c(q_t) = 3q_t^2$, which is the cost function. Therefore $p_t q_t - c(q_t)$ is the net profit of the firm, which coincides in numbers with your bonus. The higher the profit of the firm, the higher your bonus will be. You get 1200 points as the basis salary. The profit of the firm can be negative, so the payoff from the production task can be smaller than 1200.

B Testing Convergence using Linear Estimation

C Identified Forecasting Rules

D Estimated Supply Strategies

E Payoff Tables

E.1 Payoff Table for Forecasters

E.2 Payoff Table for Production Managers

market	λ	μ	R^2	MSE	Equilibrium	P-value Wald Test
p11	0.1863	24.9741	0.1113	7.6785	30.6921	0.9378
p12	0.1698	25.5817	0.1495	3.9079	30.8131	0.8073
p13	0.2108	24.3511	0.1838	5.0071	30.8546	0.7586
p14	0.1827	25.2209	0.1850	3.3689	30.8573	0.691
p21	0.6059	11.5545	0.5238	18.1876	29.3174	0.3667
p22	0.4733	15.3094	0.3085	30.4796	29.0684	0.2673
p23	0.0056	30.0469	0.0001	6.3699	30.2151	0.1540
p24	0.5191	14.8545	0.4287	17.8924	30.8868	0.9015
p25	0.2666	22.3075	0.1239	21.9634	30.4165	0.7306
p26	0.5411	14.4979	0.5543	9.7218	31.5946	0.3778
p27	0.2891	22.0542	0.2156	9.6062	31.0231	0.6383
p31	0.4189	19.0151	0.3137	18.3482	32.7227	0.0597
p32	0.4488	19.4622	0.2966	35.1703	35.3101	0.0028
p33	0.4197	14.0236	0.1950	62.8946	24.1653	0.0008
p34	0.2126	22.9448	0.0734	25.5324	29.1408	0.0818
p35	0.2351	19.5698	0.0896	18.9326	25.5849	0.0000
p36	0.1740	23.0230	0.0974	6.4048	27.8739	0.0000
p37	0.7604	5.5085	0.6605	20.8671	22.9902	0.0055
p41	0.2182	25.0770	0.1149	12.9444	32.0757	0.0408
p42	0.0674	28.1808	0.0093	17.4663	30.2188	0.4225
p43	0.3334	19.9712	0.2167	14.7826	29.9588	0.3500
p44	0.4994	14.6045	0.3604	23.8659	29.1717	0.2652
p45	0.3080	20.7341	0.2789	6.6841	29.9616	0.1507
p46	0.3530	19.2175	0.1990	23.6375	29.7035	0.3385

Table 7: This table shows the estimation results of $p_{j,t} = \lambda_j p_{j,t-1} + \mu_j + \epsilon_{j,t}$ on the markets from different treatments. The results start with p_{11}, which is the first market in treatment 1, and through p_{46}, which is the sixth market in treatment 4.

Participant	Type	Coefficient	$p-value$	R^2	MSE
exp12	A	0.7694	0.0000	0.3903	2.3471
exp13	A	0.7377	0.0000	0.3061	6.4884
exp14	A	0.8699	0.0000	0.4108	5.7149
exp16	C	-0.3954	0.0000	-0.4849	3.0398
exp21	A	0.4213	0.0000	0.3754	0.8548
exp22	A	0.8927	0.0000	0.3079	3.3428
exp23	A	0.5972	0.0000	-0.0826	2.8951
exp24	A	0.7315	0.0000	0.5123	2.1833
exp26	A	0.869	0.0000	0.5607	1.6952
exp32	A	0.8157	0.0000	0.0349	18.6036
exp33	A	0.7843	0.0000	0.3526	5.3878
exp34	C	-0.8417	0.0000	0.4292	2.3739
exp35	A	0.8046	0.0000	0.1779	10.4585
exp36	A	0.5127	0.0000	0.1428	1.0798
exp41	A	0.9088	0.0000	0.0314	19.9845
exp42	A	0.4992	0.0000	0.6754	0.5255
exp43	C	0.0179	0.0000	0.5062	1.2445
exp44	A	0.7407	0.0000	0.6352	0.8499
exp45	A	0.8464	0.0000	0.5273	0.9977

Table 8: This table shows the estimation results for the subjects who can successfully categorized by one forecasting rule in Treatment 1. In the "Type" column, "A" means adaptive rule, "C" means contrarian/trend extrapolation rule.

Participant	Type	Coefficient	$p-value$	R^2	MSE
exp11	A	0.9247	0.0000	0.568	9.1737
exp12	C	-0.4489	0.0000	0.4915	5.4309
exp13	A	0.8778	0.0000	0.4717	14.7741
exp14	A	0.9245	0.0000	0.4849	10.3821
exp16	A	0.673	0.0000	0.1749	12.4204
exp21	A	0.8436	0.0000	0.6769	12.1518
exp22	C	-0.6319	0.0000	0.6336	10.343
exp23	C	-0.4639	0.0000	0.4633	10.0276
exp24	C	-0.4922	0.0000	0.6665	8.8984
exp25	A	0.7225	0.0000	0.6602	11.5174
exp26	C	-0.7042	0.0000	0.3808	22.5964
exp31	A	0.7621	0.0000	0.1423	76.3659
exp32	A	0.5417	0.0000	0.5974	24.5421
exp33	A	0.6442	0.0000	0.724	10.3266
exp34	A	0.6899	0.0000	0.468	26.425
exp36	C	0.6408	0.0000	0.7821	10.767
exp51	A	0.4989	0.0000	0.513	5.803
exp52	C	-0.153	0.0003	0.7033	2.943
exp53	A	0.7542	0.0000	0.4008	9.3458
exp54	A	0.4362	0.0000	0.4895	6.6524
exp55	A	0.7914	0.0000	0.3238	44.0411
exp56	A	0.9086	0.0000	0.5106	11.6252
exp61	A	0.584	0.0000	-0.3159	2.7391
exp62	A	0.8386	0.0000	0.758	1.3188
exp63	A	0.8655	0.0000	0.9578	0.1774
exp65	A	0.7223	0.0000	0.8278	0.7259
exp66	C	0.0356	0.0000	0.5704	3.2036
exp72	A	0.416	0.0000	0.8604	7.1865
exp73	C	-0.1777	0.0004	0.9276	2.8821
exp74	A	0.2722	0.0001	0.6247	24.5324
exp75	C	-0.4258	0.0000	0.8927	6.3127
exp76	A	0.4953	0.0000	0.8135	6.8963

Table 9: This table shows the estimation results for the subjects who can successfully categorized by one forecasting rule in Treatment 3. In the "Type" column, "A" means adaptive rule, "C" means contrarian/trend extrapolation rule.

Participant	Type	Coefficient	$p-value$	R^2	MSE
exp11	A	0.6495	0.0000	0.8129	2.0813
exp12	C	-0.4201	0.0000	0.0841	5.7926
exp14	C	-0.0851	0.0000	0.3344	8.1675
exp16	C	-0.3519	0.0000	0.0871	4.6497
exp21	C	-0.2769	0.0000	0.3691	7.351
exp22	C	-0.6555	0.0000	0.7364	3.8006
exp25	C	-0.352	0.0011	0.2997	19.7205
exp26	A	0.8179	0.0000	0.8584	1.7657
exp31	A	0.8627	0.0000	0.7207	4.3887
exp32	A	0.507	0.0000	0.6858	2.4103
exp33	A	0.4594	0.0000	0.4918	5.3313
exp34	A	0.777	0.0000	0.866	1.4416
exp35	A	0.6202	0.0000	0.5198	6.4221
exp36	C	-0.3001	0.0169	-0.1436	18.7005
exp42	C	-0.6367	0.0000	0.7612	7.1209
exp43	C	-0.5521	0.0000	0.7237	10.4271
exp44	A	0.7716	0.0000	0.3301	15.5369
exp46	C	-0.6195	0.0000	0.8119	4.7725
exp51	A	0.8902	0.0000	0.2657	6.8258
exp52	A	0.5709	0.0000	0.3391	5.1909
exp53	A	0.7164	0.0000	0.3517	4.646
exp54	A	0.6875	0.0000	0.8253	0.9544
exp55	A	0.7167	0.0000	0.6162	1.3439
exp56	A	0.865	0.0000	0.4349	4.2283
exp62	A	0.8027	0.0000	0.4998	12.045
exp63	A	0.7674	0.0000	0.4927	15.8082
exp64	A	0.6532	0.0000	0.9321	1.3295
exp66	A	0.9196	0.0000	0.7403	7.6132

Table 10: This table shows the estimation results for the subjects who can successfully categorized by one forecasting rule in Treatment 4. In the "Type" column, "A" means adaptive rule, "C" means contrarian/trend extrapolation rule.

Payoff Table							
Payoff from Forecasting Task $= \max[1300 - \frac{1300}{49}(\text{Your Prediction Error})^2, 0]$							
1300 points equal 0.5 euro							
error	points	error	points	error	points	error	points
0	1300	1.85	1209	3.7	937	5.55	483
0.05	1300	1.9	1204	3.75	927	5.6	468
0.1	1300	1.95	1199	3.8	917	5.65	453
0.15	1299	2	1194	3.85	907	5.7	438
0.2	1299	2.05	1189	3.9	896	5.75	423
0.25	1298	2.1	1183	3.95	886	5.8	408
0.3	1298	2.15	1177	4	876	5.85	392
0.35	1297	2.2	1172	4.05	865	5.9	376
0.4	1296	2.25	1166	4.1	854	5.95	361
0.45	1295	2.3	1160	4.15	843	6	345
0.5	1293	2.35	1153	4.2	832	6.05	329
0.55	1292	2.4	1147	4.25	821	6.1	313
0.6	1290	2.45	1141	4.3	809	6.15	297
0.65	1289	2.5	1134	4.35	798	6.2	280
0.7	1287	2.55	1127	4.4	786	6.25	264
0.75	1285	2.6	1121	4.45	775	6.3	247
0.8	1283	2.65	1114	4.5	763	6.35	230
0.85	1281	2.7	1107	4.55	751	6.4	213
0.9	1279	2.75	1099	4.6	739	6.45	196
0.95	1276	2.8	1092	4.65	726	6.5	179
1	1273	2.85	1085	4.7	714	6.55	162
1.05	1271	2.9	1077	4.75	701	6.6	144
1.1	1268	2.95	1069	4.8	689	6.65	127
1.15	1265	3	1061	4.85	676	6.7	109
1.2	1262	3.05	1053	4.9	663	6.75	91
1.25	1259	3.1	1045	4.95	650	6.8	73
1.3	1255	3.15	1037	5	637	6.85	55
1.35	1252	3.2	1028	5.05	623	6.9	37
1.4	1248	3.25	1020	5.1	610	6.95	19
1.45	1244	3.3	1011	5.15	596	error≥7	0
1.5	1240	3.35	1002	5.2	583		
1.55	1236	3.4	993	5.25	569		
1.6	1232	3.45	984	5.3	555		
1.65	1228	3.5	975	5.35	541		
1.7	1223	3.55	966	5.4	526		
1.75	1219	3.6	956	5.45	512		
1.8	1214	3.65	947	5.5	497		

Figure 10: The payoff table for forecasters.

Payoff Table for Production Task (2600 points=1 euro)

price\quantity	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
0	1200	1197	1188	1173	1152	1125	1092	1053	1008	957	900	837	768	693	612	525	432	333	228	117	0
1	1200	1198	1190	1176	1156	1130	1098	1060	1016	966	910	848	780	706	626	540	448	350	246	136	20
2	1200	1199	1192	1179	1160	1135	1104	1067	1024	975	920	859	792	719	640	555	464	367	264	155	40
3	1200	1200	1194	1182	1164	1140	1110	1074	1032	984	930	870	804	732	654	570	480	384	282	174	60
4	1200	1201	1196	1185	1168	1145	1116	1081	1040	993	940	881	816	745	668	585	496	401	300	193	80
5	1200	1202	1198	1188	1172	1150	1122	1088	1048	1002	950	892	828	758	682	600	512	418	318	212	100
6	1200	1203	1200	1191	1176	1155	1128	1095	1056	1011	960	903	840	771	696	615	528	435	336	231	120
7	1200	1204	1202	1194	1180	1160	1134	1102	1064	1020	970	914	852	784	710	630	544	452	354	250	140
8	1200	1205	1204	1197	1184	1165	1140	1109	1072	1029	980	925	864	797	724	645	560	469	372	269	160
9	1200	1206	1206	1200	1188	1170	1146	1116	1080	1038	990	936	876	810	738	660	576	486	390	288	180
10	1200	1207	1208	1203	1192	1175	1152	1123	1088	1047	1000	947	888	823	752	675	592	503	408	307	200
11	1200	1208	1210	1206	1196	1180	1158	1130	1096	1056	1010	958	900	836	766	690	608	520	426	326	220
12	1200	1209	1212	1209	1200	1185	1164	1137	1104	1065	1020	969	912	849	780	705	624	537	444	345	240
13	1200	1210	1214	1212	1204	1190	1170	1144	1112	1074	1030	980	924	862	794	720	640	554	462	364	260
14	1200	1211	1216	1215	1208	1195	1176	1151	1120	1083	1040	991	936	875	808	735	656	571	480	383	280
15	1200	1212	1218	1218	1212	1200	1182	1158	1128	1092	1050	1002	948	888	822	750	672	588	498	402	300
16	1200	1213	1220	1221	1216	1205	1188	1165	1136	1101	1060	1013	960	901	836	765	688	605	516	421	320
17	1200	1214	1222	1224	1220	1210	1194	1172	1144	1110	1070	1024	972	914	850	780	704	622	534	440	340
18	1200	1215	1224	1227	1224	1215	1200	1179	1152	1119	1080	1035	984	927	864	795	720	639	552	459	360
19	1200	1216	1226	1230	1228	1220	1206	1186	1160	1128	1090	1046	996	940	878	810	736	656	570	478	380
20	1200	1217	1228	1233	1232	1225	1212	1193	1168	1137	1100	1057	1008	953	892	825	752	673	588	497	400
21	1200	1218	1230	1236	1236	1230	1218	1200	1176	1146	1110	1068	1020	966	906	840	768	690	606	516	420
22	1200	1219	1232	1239	1240	1235	1224	1207	1184	1155	1120	1079	1032	979	920	855	784	707	624	535	440
23	1200	1220	1234	1242	1244	1240	1230	1214	1192	1164	1130	1090	1044	992	934	870	800	724	642	554	460
24	1200	1221	1236	1245	1248	1245	1236	1221	1200	1173	1140	1101	1056	1005	948	885	816	741	660	573	480
25	1200	1222	1238	1248	1252	1250	1242	1228	1208	1182	1150	1112	1068	1018	962	900	832	758	678	592	500
26	1200	1223	1240	1251	1256	1255	1248	1235	1216	1191	1160	1123	1080	1031	976	915	848	775	696	611	520
27	1200	1224	1242	1254	1260	1260	1254	1242	1224	1200	1170	1134	1092	1044	990	930	864	792	714	630	540
28	1200	1225	1244	1257	1264	1265	1260	1249	1232	1209	1180	1145	1104	1057	1004	945	880	809	732	649	560
29	1200	1226	1246	1260	1268	1270	1266	1256	1240	1218	1190	1156	1116	1070	1018	960	896	826	750	668	580

(This table continues at the back of this page.)

Figure 11: The payoff table for production managers, page 1.

price\quantity	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
30	1200	1227	1248	1263	1272	1275	1272	1263	1248	1227	1200	1167	1128	1083	1032	975	912	843	768	687	600
31	1200	1228	1250	1266	1276	1280	1278	1270	1256	1236	1210	1178	1140	1096	1046	990	928	860	786	706	620
32	1200	1229	1252	1269	1280	1285	1284	1277	1264	1245	1220	1189	1152	1109	1060	1005	944	877	804	725	640
33	1200	1230	1254	1272	1284	1290	1290	1284	1272	1254	1230	1200	1164	1122	1074	1020	960	894	822	744	660
34	1200	1231	1256	1275	1288	1295	1296	1291	1280	1263	1240	1211	1176	1135	1088	1035	976	911	840	763	680
35	1200	1232	1258	1278	1292	1300	1302	1298	1288	1272	1250	1222	1188	1148	1102	1050	992	928	858	782	700
36	1200	1233	1260	1281	1296	1305	1308	1305	1296	1281	1260	1233	1200	1161	1116	1065	1008	945	876	801	720
37	1200	1234	1262	1284	1300	1310	1314	1312	1304	1290	1270	1244	1212	1174	1130	1080	1024	962	894	820	740
38	1200	1235	1264	1287	1304	1315	1320	1319	1312	1299	1280	1255	1224	1187	1144	1095	1040	979	912	839	760
39	1200	1236	1266	1290	1308	1320	1326	1326	1320	1308	1290	1266	1236	1200	1158	1110	1056	996	930	858	780
40	1200	1237	1268	1293	1312	1325	1332	1333	1328	1317	1300	1277	1248	1213	1172	1125	1072	1013	948	877	800
41	1200	1238	1270	1296	1316	1330	1338	1340	1336	1326	1310	1288	1260	1226	1186	1140	1088	1030	966	896	820
42	1200	1239	1272	1299	1320	1335	1344	1347	1344	1335	1320	1299	1272	1239	1200	1155	1104	1047	984	915	840
43	1200	1240	1274	1302	1324	1340	1350	1354	1352	1344	1330	1310	1284	1252	1214	1170	1120	1064	1002	934	860
44	1200	1241	1276	1305	1328	1345	1356	1361	1360	1353	1340	1321	1296	1265	1228	1185	1136	1081	1020	953	880
45	1200	1242	1278	1308	1332	1350	1362	1368	1368	1362	1350	1332	1308	1278	1242	1200	1152	1098	1038	972	900
46	1200	1243	1280	1311	1336	1355	1368	1375	1376	1371	1360	1343	1320	1291	1256	1215	1168	1115	1056	991	920
47	1200	1244	1282	1314	1340	1360	1374	1382	1384	1380	1370	1354	1332	1304	1270	1230	1184	1132	1074	1010	940
48	1200	1245	1284	1317	1344	1365	1380	1389	1392	1389	1380	1365	1344	1317	1284	1245	1200	1149	1092	1029	960
49	1200	1246	1286	1320	1348	1370	1386	1396	1400	1398	1390	1376	1356	1330	1298	1260	1216	1166	1110	1048	980
50	1200	1247	1288	1323	1352	1375	1392	1403	1408	1407	1400	1387	1368	1343	1312	1275	1232	1183	1128	1067	1000
51	1200	1248	1290	1326	1356	1380	1398	1410	1416	1416	1410	1398	1380	1356	1326	1290	1248	1200	1146	1086	1020
52	1200	1249	1292	1329	1360	1385	1404	1417	1424	1425	1420	1409	1392	1369	1340	1305	1264	1217	1164	1105	1040
53	1200	1250	1294	1332	1364	1390	1410	1424	1432	1434	1430	1420	1404	1382	1354	1320	1280	1234	1182	1124	1060
54	1200	1251	1296	1335	1368	1395	1416	1431	1440	1443	1440	1431	1416	1395	1368	1335	1296	1251	1200	1143	1080
55	1200	1252	1298	1338	1372	1400	1422	1438	1448	1452	1450	1442	1428	1408	1382	1350	1312	1268	1218	1162	1100
56	1200	1253	1300	1341	1376	1405	1428	1445	1456	1461	1460	1453	1440	1421	1396	1365	1328	1285	1236	1181	1120
57	1200	1254	1302	1344	1380	1410	1434	1452	1464	1470	1470	1464	1452	1434	1410	1380	1344	1302	1254	1200	1140
58	1200	1255	1304	1347	1384	1415	1440	1459	1472	1479	1480	1475	1464	1447	1424	1395	1360	1319	1272	1219	1160
59	1200	1256	1306	1350	1388	1420	1446	1466	1480	1488	1490	1486	1476	1460	1438	1410	1376	1336	1290	1238	1180
60	1200	1257	1308	1353	1392	1425	1452	1473	1488	1497	1500	1497	1488	1473	1452	1425	1392	1353	1308	1257	1200

Figure 12: The payoff table for production managers, page 2.

Participant	c_0	p-value	c_1	p-value	R^2	MSE	Type
exp13	3.6256	0.0006	0.0270	0.4007	0.0145	1.6985	C
exp14	6.6367	0.0250	-0.0437	0.6285	0.0049	8.3384	C
exp15	2.2833	0.0328	0.0986	0.0030	0.1552	0.6952	H
exp21	-0.0537	0.9133	0.1656	0.0000	0.7515	0.4390	O
exp24	15.5911	0.0000	-0.2960	0.0001	0.2346	12.4429	H
exp25	9.7377	0.0000	-0.0963	0.0409	0.0801	3.8086	H
exp26	17.9097	0.0000	-0.4436	0.0000	0.6938	3.9053	H
exp42	9.9127	0.0000	-0.1609	0.0003	0.2166	2.4122	H
exp46	8.0166	0.0000	-0.0318	0.3819	0.0157	2.0399	C
exp54	3.8616	0.0001	0.0241	0.5690	0.0067	1.1388	C
exp55	2.1090	0.0000	0.0726	0.0002	0.2287	1.1636	H
exp63	1.7970	0.0243	0.1129	0.0001	0.2470	0.2817	H
exp64	0.6905	0.1161	0.1613	0.0000	0.6936	0.2877	O

Table 11: This table shows the estimated coefficients of the supply strategy used by the subjects in Treatment 3.

Participant	c_0	p-value	c_1	p-value	R^2	MSE	Type
q12	0.3517	0.5014	0.1531	0.0000	0.6296	0.0969	O
q13	2.8890	0.0000	0.0670	0.0000	0.3376	0.0800	H
q14	0.0648	0.9744	0.1748	0.0060	0.1357	4.4876	O
q16	3.9698	0.0000	0.0318	0.0041	0.1467	0.0519	H
q21	0.1838	0.6911	0.1642	0.0000	0.7061	0.1834	O
q22	0.7196	0.0125	0.1391	0.0000	0.8170	0.0691	H
q23	2.7288	0.0348	0.0788	0.0769	0.0612	1.1603	C
q26	0.3254	0.8720	0.1643	0.0109	0.1191	3.7786	O
q31	6.4322	0.0000	-0.0349	0.3841	0.0155	1.4020	H
q32	0.4993	0.0965	0.1535	0.0000	0.8184	0.0906	O
q33	1.4861	0.0373	0.1191	0.0000	0.3333	0.2983	H
q34	0.9924	0.4204	0.1424	0.0004	0.2055	1.4840	O
q35	3.4949	0.0007	0.0515	0.1297	0.0457	0.8554	H
q43	13.4421	0.0000	-0.2624	0.0000	0.6169	1.8068	H
q45	0.3209	0.2657	0.1558	0.0000	0.8434	0.0975	O
q54	2.1929	0.0000	0.0936	0.0000	0.3941	0.0739	H
q56	0.2222	0.4602	0.1563	0.0000	0.8465	0.0679	O
q62	-0.2469	0.2875	0.1733	0.0000	0.9171	0.0828	O

Table 12: This table shows the estimated coefficients of the supply strategy used by the subjects in Treatment 4.We use C to denote the use of the constant supply strategy, O to denote use of the conditionally optimal supply strategy and H to denote use of the hybrid strategy.