

Bao, Te; Duffy, John; Hommes, Cars

Working Paper

Learning, Forecasting and Optimizing: An Experimental Study

Tinbergen Institute Discussion Paper, No. 12-015/1

Provided in Cooperation with:

Tinbergen Institute, Amsterdam and Rotterdam

Suggested Citation: Bao, Te; Duffy, John; Hommes, Cars (2012) : Learning, Forecasting and Optimizing: An Experimental Study, Tinbergen Institute Discussion Paper, No. 12-015/1, Tinbergen Institute, Amsterdam and Rotterdam

This Version is available at:

<https://hdl.handle.net/10419/87393>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

TI 2012-015/1
Tinbergen Institute Discussion Paper



Learning, Forecasting and Optimizing: An Experimental Study

Te Bao^a

John Duffy^b

Cars Hommes^a

^a *CeNDEF, Faculty of Economics and Business Administration, University of Amsterdam,
and Tinbergen Institute;*

^b *Department of Economics, University of Pittsburgh.*

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

More TI discussion papers can be downloaded at <http://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 1600

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Duisenberg school of finance is a collaboration of the Dutch financial sector and universities, with the ambition to support innovative research and offer top quality academic education in core areas of finance.

DSF research papers can be downloaded at: <http://www.dsf.nl/>

Duisenberg school of finance
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 8579

Learning, Forecasting and Optimizing: An Experimental Study*

Te Bao^a, John Duffy^b, and Cars Hommes^a

^aCeNDEF, University of Amsterdam and Tinbergen Institute

Valckenierstraat 65-67, 1018XE Amsterdam, NL

^bDepartment of Economics, University of Pittsburgh

4901 Posvar Hall, Pittsburgh, PA 15260 USA

February 16, 2012

*The authors thank Jasmina Arifovic, George Evans, Ernst Fehr, Andrew Schotter, Theo Offerman, Lise Vesterlund, Stephanie Wang, seminar participants at the University of Pittsburgh, CESS-CREED Amsterdam and the Symposium on Behavioral and Experimental Economics for helpful comments and discussions. We acknowledge financial support from the EU 7th framework collaborative project “Monetary, Fiscal and Structural Policies with Heterogeneous Agents (POLHIA)” Grant No. 225048 and from the Research Priority Area Behavioral Economics of the University of Amsterdam. Te Bao thanks the Willems foundation for funding his visit to the University of Pittsburgh to work on this paper.

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 converge to Rational Expectation Equilibrium (REE), but at very different speeds. 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 in treatment 4) where the forecast is made by the other member of their team, which suggests that “two heads are better than one” in finding REE.

JEL Classification: C91, C92, D83, D84

Keywords: Learning, Rational Expectations, Optimization, Experimental Economics, Bounded Rationality.

1 Introduction

Rational Expectations (RE) macroeconomic models have two crucial dimensions: 1) Rational agents correctly forecast future prices given all available information, that is, they do not make systematic forecast mistakes. 2) Given agents' rational expectations, these same agents solve optimization problems that determine their consumption and/or production decisions, which then, via market clearing, determine the realizations of prices and wages the agents were seeking to forecast; these data are then used to update forecasts. Thus, RE systems are *self-referential*; beliefs affect outcomes and outcomes affect beliefs.

Testing rational expectation models with field data is problematic as agents' expectations are not generally observable and economists may disagree as to what constitutes the "true" model in which agents' expectations are formed. An alternative approach is to test rational expectations models in the laboratory where it is possible to control the model that determines economic data and to elicit and use agents' expectations of future variables in the determination of that same data. However, the self-referential nature of RE models makes them difficult to test in the laboratory. As 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 macroeconomic 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 then determine the actual price realizations (the object of agent's forecasts), 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 a "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 a rational expectations equilibrium (REE) (see e.g., Hommes 2011 for a survey). In some instances, subjects learn a REE 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. One aim of this research is to assess whether the results obtained in the LtFE literature are robust when the optimization task is performed by an individual rather than by a computer

program. Moreover, our novel team specialization treatment 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. This type of team specialization set-up has not been previously explored in the laboratory.

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) development of the cobweb model. Furthermore, Muth (1961) proposed rational expectations in the context of this very 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 REE 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 using the LtFE design.

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 in play of a Beauty-Contest Game. 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 games 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 individuals. 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 all group members. By contrast, our team treatment involves *specialization* of tasks between the two group members, who share a common interest in maximizing their joint payoff.

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 adoption of different learning rules, and market evolution toward 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 the 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. They found none of their information treatments generate successful collusion, and information that encourage "imitating the best" learning leads to a 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 REE 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 a 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 subject's 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 (optimization) decisions from a 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 (same tables as in Treatments 1 and 2, respectively).
4. Treatment 4: the LtFE+LtOE Team treatment. In this treatment, there is a forecaster and a production manager in each two-agent 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 (same tables as in Treatments 1 and 2, respectively).

Prices in the experiment is restricted to be non-negative, so forecasters are also not allowed to submit negative price forecasts. We set 60 as the upper bound for the price prediction, because this is the maximum possible 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 subjects who participated in the 8 sessions of this experiment. No subject participated in more than one 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:

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

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, differing only 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:

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

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 (equation 1) and noting that the expected value of the noise term is zero, we have that:

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

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 78.70.

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 a table 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 shown 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, (though we did record data on the time it took to make certain decisions, as discussed later). The market price was not determined until all N subjects had submitted their price forecasts and/or quantity production decisions. Each round

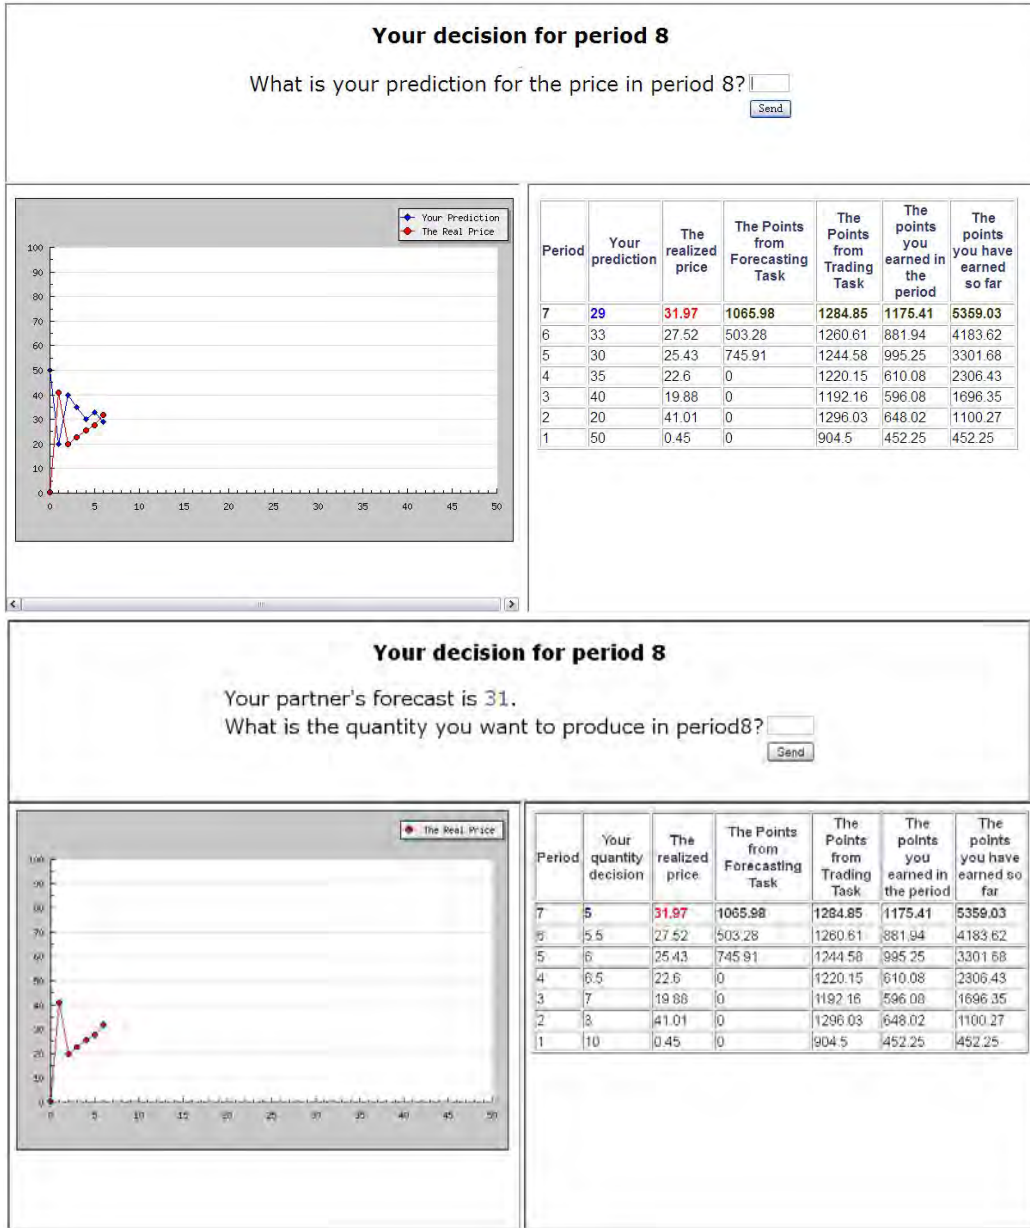


Figure 1: The computer interface for forecasters (top) and production managers (bottom).

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 (in points) is a decreasing function of their prediction error, and was given by:

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

Notice that subjects earn 0 if their price forecast error is greater than 7, and they a maximum of 1300 for a perfect forecast.

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

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

Notice that subjects get a baseline ‘salary’ of 1200 points 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, our set-up implies that the maximum possible 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 equals the REE price 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 payoff 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 and this conversion rate was known to subjects in advance.

4 Experimental Results

4.1 Aggregate Market Price

Figure 2 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. Thus, the general tendency for a negative feedback system to converge to REE is not greatly affected by the type of task that is 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.

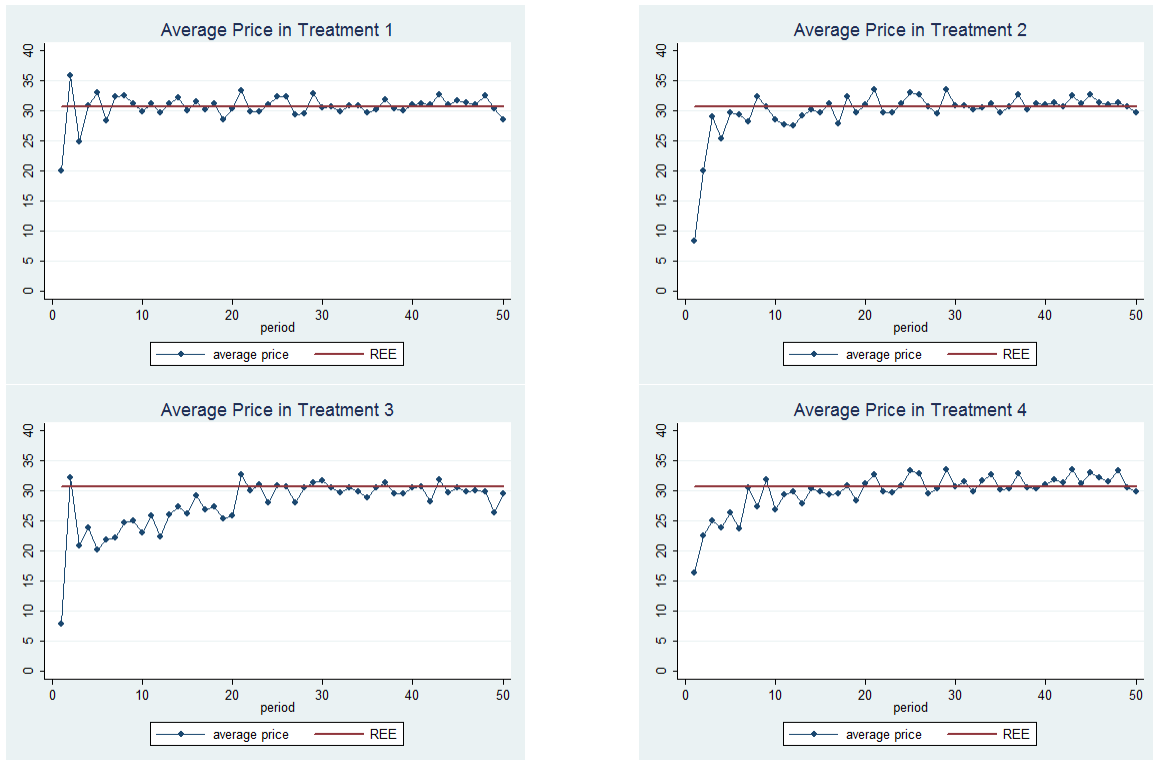


Figure 2: The average market price against the REE price in each of the four treatments.

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. Using this criterion, we count the number of periods required before convergence obtains across our 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. Comparing these time-to-convergence numbers, we observe that 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). We further observe that 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). Finally, treatment 4 converges slightly faster than treatment 2 on average, but that difference is not significant at 5% level according to a Wilcoxon Mann Whitney test.

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

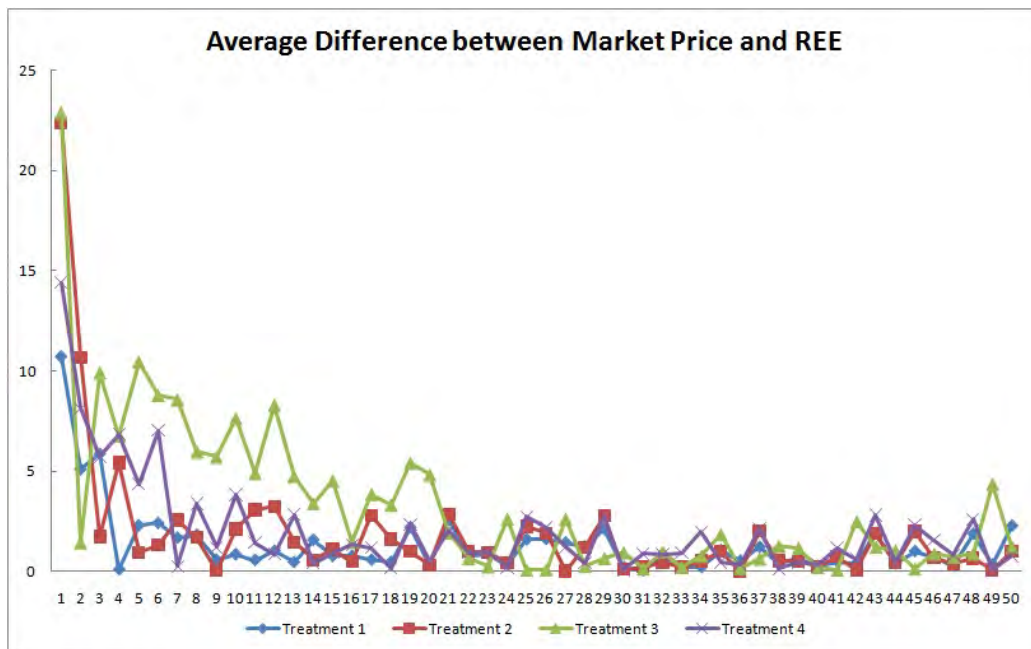


Figure 3: The distance between the fundamental price and the average of the market prices from all markets of each treatment.

Figure 3 reveals that the difference decreases most rapidly toward zero in treatment 1 (diamonds), and most slowly in treatment 3 (triangles). Treatment 2 (squares) and treatment 4 (Xs) are very similar to one another.

Treatment	Market	Number of Periods to 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
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
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 to convergence for each market.

Finally we can test for convergence econometrically using a method suggested by Duffy (2008). For each market j , the following linear equation is estimated:

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

The results of this estimation exercise are reported in Appendix B. We note first that all of the estimated λ s and μ s are significantly different from 0 at the 5% level of significance. We also checked for evidence of serially correlated errors. For our estimation, the relevant upper bound of the Durbin-Watson Statistic, dU , ($n = 50, k' = 2$) is 1.445. We found that for each market, the estimated Durbin-Watson statistics were always greater than that upper bound, which implies that we cannot reject the null hypothesis of no first order serial correlation in the error terms. The estimated linear equation is stable if $|\lambda|$ is smaller than 1¹, and has a long-run equilibrium level $\frac{\hat{\mu}_j}{1-\hat{\lambda}_j}$. For each market j , we declare that *weak convergence* obtains if we can reject $\hat{\lambda}_j \geq 1$ at 5% level, and we say that *strong convergence* obtains if we cannot reject $\frac{\hat{\mu}_j}{1-\hat{\lambda}_j} = 30.73$ (the REE value) at the 5% level (using a Wald test). Summarizing the estimation results (as reported in Appendix B), we have:

1. All markets in all four treatments satisfy *weak convergence*.
2. All markets in treatments 1 and 2 satisfy *strong convergence*. All but one market in treatment 4 satisfies 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.

¹As all the estimated λ s are positive, we just need to check whether $\lambda \geq 1$ is rejected. This statement is equivalent to the claim that the price dynamics are stationary.

4.2 Individual–Level Decisions

4.2.1 Distribution of 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 price forecasts and optimization (quantity-choice) decisions is shown in Figure 4 using pooled data from all markets of the various treatments (treatments 1, 2 and 4 for the price forecasts and treatments 2, 3 and 4 for the quantity choices). Under rational expectations the CDF should be a step function switching from 0 to 100% at the RE price ($re=30.73$) or quantity ($qre=5.12$).

Figure 4 reveals that there is some heterogeneity in individual decisions across treatments with the largest departures from RE predictions occurring in treatment 3, 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 possible 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 the 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

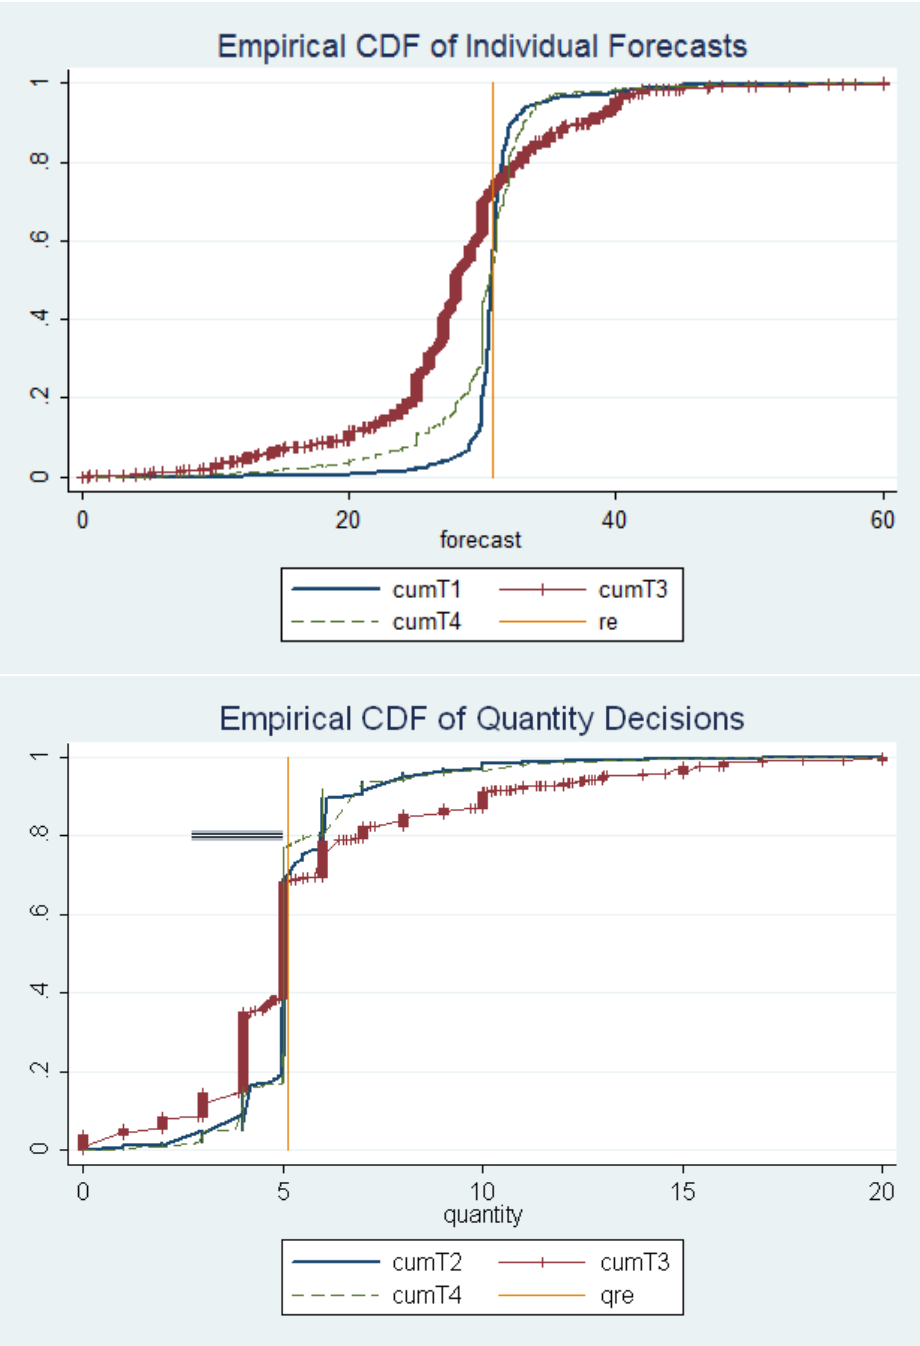


Figure 4: The empirical cdf of individual price forecasts and quantity decisions.

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 possible pairing of the three treatments. The tests 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 the 5% level). In particular, there is much greater heterogeneity in the quantity decisions of treatment 3 as compared with either treatments 2 or 4.

4.2.2 Time Taken to Make Decisions

We also collected data on the time it took for subjects to make their decision(s). Such data can be useful in understanding the cognitive difficulty of decision-making. In particular, Rubinstein (2007) provides evidence that choices requiring greater cognitive activity are associated with longer decision response times. While there was no decision time limit in our experiment (subjects could take as much time as they wished for each decision), the computer program that implements our experiment start counting (in seconds) when a subject first entered each new period, and stopped counting when he or she submitted his or her decision(s). Figure 5 plots the empirical CDF of the time taken by subjects in each period of treatments 1, 2 and 3². As subjects submit their forecasts and quantity decisions together in treatment 3, the decision-time data for treatment 3 is the total time taken for both the forecasting and optimizing tasks.

Figure 5 clearly reveals that subjects take less time to make their decisions in treatments 1 and 2 as compared with treatment 3. The average decision time per period is 19.83 seconds in treatment 1, 21.16 seconds in treatment 2 and 33.99 seconds in treatment 3. The difference between either treatment 1 or 2 and treatment 3 is significant at the 5% level according to a Wilcoxon Mann Whitney test. The difference between treatments 1 and 2 is not significant at the 5% level. These results confirm the notion that making two decisions, as in treatment 3, is indeed more cognitively challenging than making a single decision as in treatments 1 and 2.

²There was a technical problem with decision-time capture in treatment 4, and as a consequence we cannot construct precise decision-time data for that treatment.

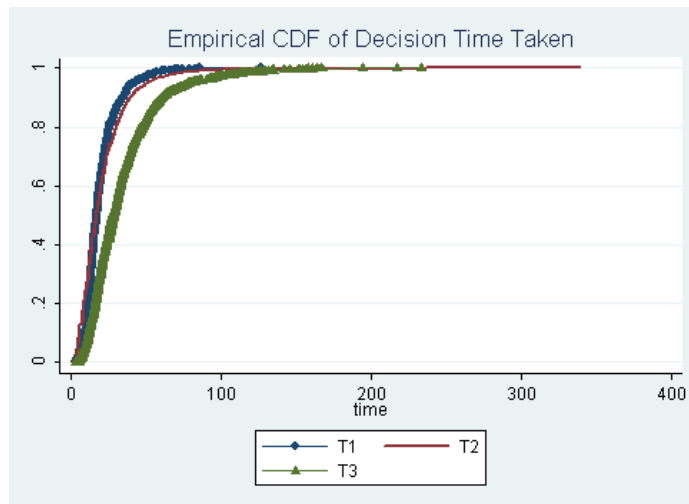


Figure 5: The empirical cdf of the time taken to complete decision tasks in treatments 1, 2 and 3. The unit of time is seconds, as measured on the horizontal axis.

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: MSD from REE and variance of prices 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, as reported in Table 4, are generally very high (exceed 80%) in all four treatments. The ranking of average efficiency over all 50 periods is $Treatment2 > Treamtnet4 > Treament1 > Treatment3$, while the ranking for the second 25 periods is $Treatment2 > Treatment1 > Treatment4 > Treament3$. Only the difference between efficiency in treatment 2 and the other treatments is significant at the 5% level according to Wilcoxon Mann Whitney test. The differences

between the efficiency levels in the other treatments are not significant.

Treatment	Market	Period 1-50		Period 1-25		Period 26-50	
		Average Earnings	Efficiency	Average Earnings	Efficiency	Average Earnings	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: Average earnings and 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 treatment 1, or the implicit efficiency of the forecasting task in treatment 2. For treatment 1, it is straightforward that the firm will produce as much as one sixth of the price prediction, and the profit of the firm can be calculated accordingly. For treatment 2, we can assume that subjects always make production decisions that are conditionally optimal for their implicit price 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 the 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 a 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 greater 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 treatments³, but there are no other significant differences in all pairwise comparisons between treatments. As we will see later in the paper, subjects in treatments 3 and 4 (especially treatment 3) do not make perfect production decisions given their forecasts. This result suggests that there is not much change in efficiency if subjects are boundedly rational in optimization tasks. However, if subjects are not fully rational with regard to the optimization task, this may result in inaccurate forecasts resulting in larger forecast efficiency losses.

This result also suggests that high efficiency levels in learning to optimize experiments should be treated with caution; even if efficiency metrics indicate that subjects are doing well on the optimization task, the implicit price forecasts may be far from rational. In this case, the team design with specialized roles provides a clearer view of the efficiency of the decision process for each task (and may improve efficiency levels for forecasting tasks).

³This 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).

Treatment	Periods	Avg. Payoff Forecasting	Avg. Payoff Optimization	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 optimization tasks.

4.5 Individual Forecasts

The upper panel of Figure 6 shows the average individual price forecasts in treatments 1, 3 and 4 against the REE. We observe that treatment 1 converges fastest, followed by treatment 4, and that treatment 3 is the slowest to converge. The lower panel of Figure 6 shows the average variance of individual forecasts in treatments 1, 3 and 4. We observe that heterogeneity of forecasts 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 that are often used in negative feedback markets (such as the one studied here) are adaptive expectations:

$$p_{i,t+1}^e = p_{i,t}^e + \lambda(p_t - p_{i,t}^e), \quad (7)$$

and trend extrapolation rules:

$$p_{i,t+1}^e = p_t + \gamma(p_t - p_{t-1}). \quad (8)$$

The estimated value for γ is typically negative in the negative-feedback market setting that we consider, so we will refer to the trend extrapolation rule as the “contrarian rule” to differentiate this rule from the trend-following version of the same rule where γ is positive. We estimate these two types of rules for each individual subject in our experiment. We call an estimation successful if it generates coefficient estimates that are statistically significant at the 5% level, and if there is no serial correlation in the errors. It turns out that more than 75% of subjects can be successfully characterized by *both* adaptive rules. In those cases we compare the R^2 value for each estimated

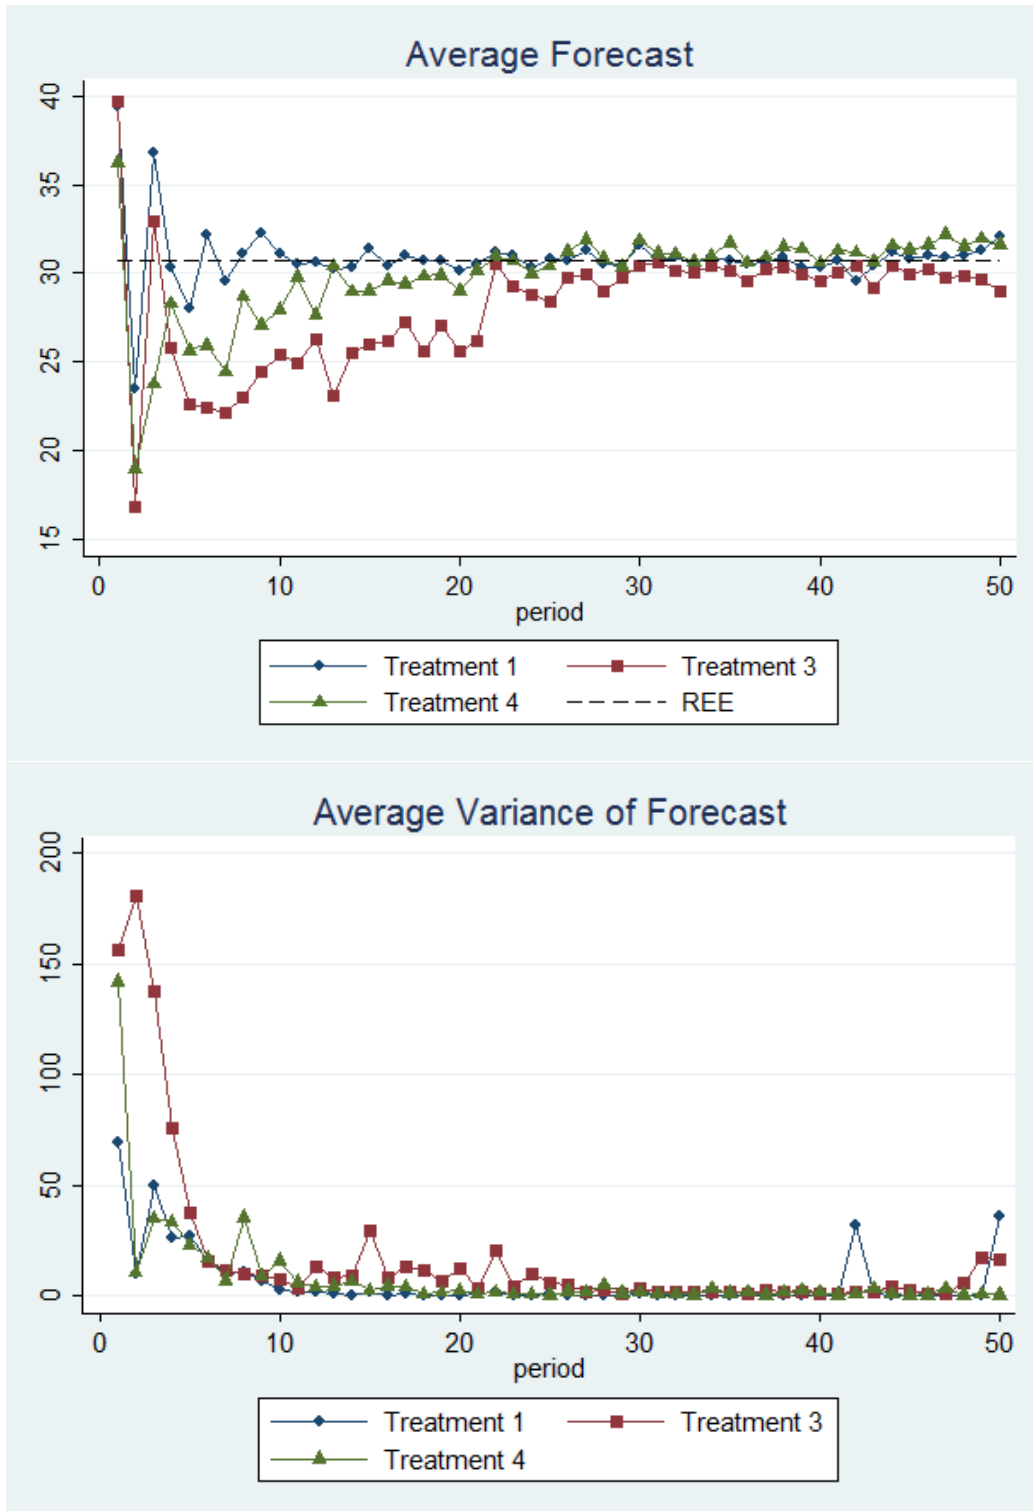


Figure 6: The upper panel shows the average individual forecasts in Treatment 1, 3 and 4. The bottom panel shows the average of the group variance of individual forecasts in Treatments 1, 3 and 4.

model and characterize the individual as following the rule with larger R^2 . The distribution of individual subjects over the types of forecasting rules is shown in Table 6 and Figure 7, while the Tables in Appendix B show the estimation results for the subjects who can be successfully identified using a single 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 of the two forecasting rules (Adaptive or Contrarian) or Neither of the two rules.

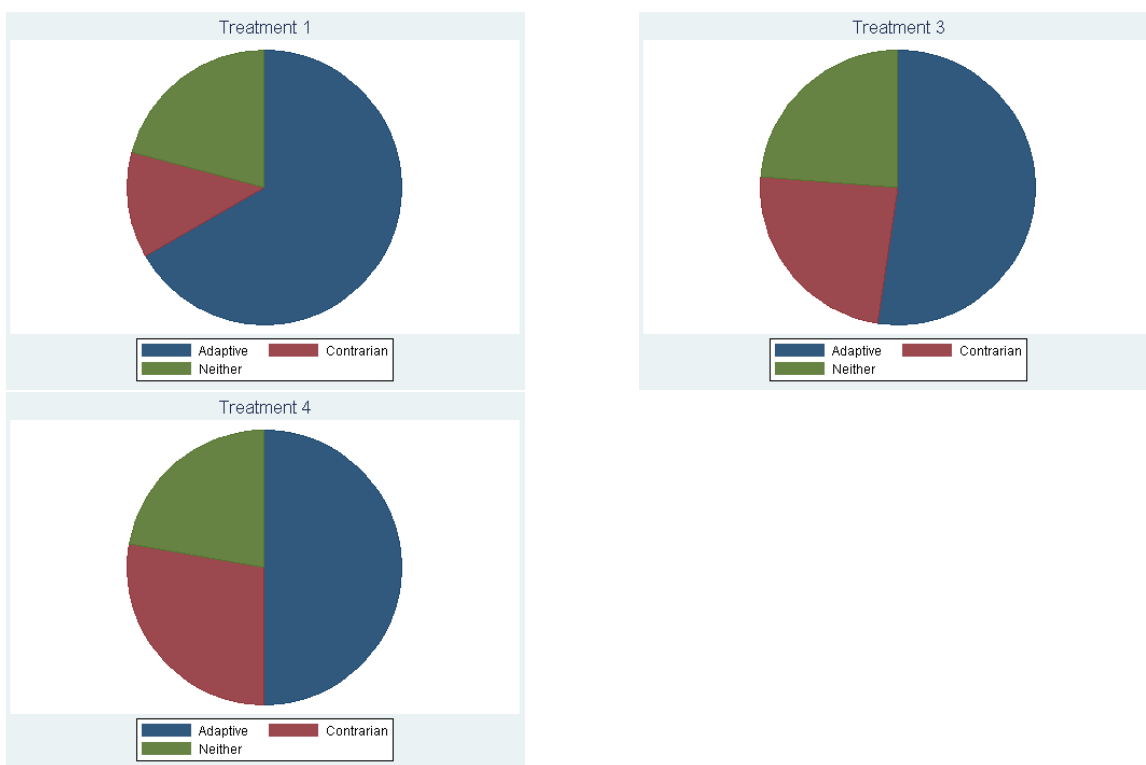


Figure 7: The fraction of subjects who are characterized by one type of forecasting rule or neither rule in treatment 1 (top left), 3 (top right) and 4 (bottom).

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 that the market price is most stable when there are overwhelmingly more subjects using the adaptive rule.

4.6 Individual Supply Decision

4.6.1 Descriptive Statistics

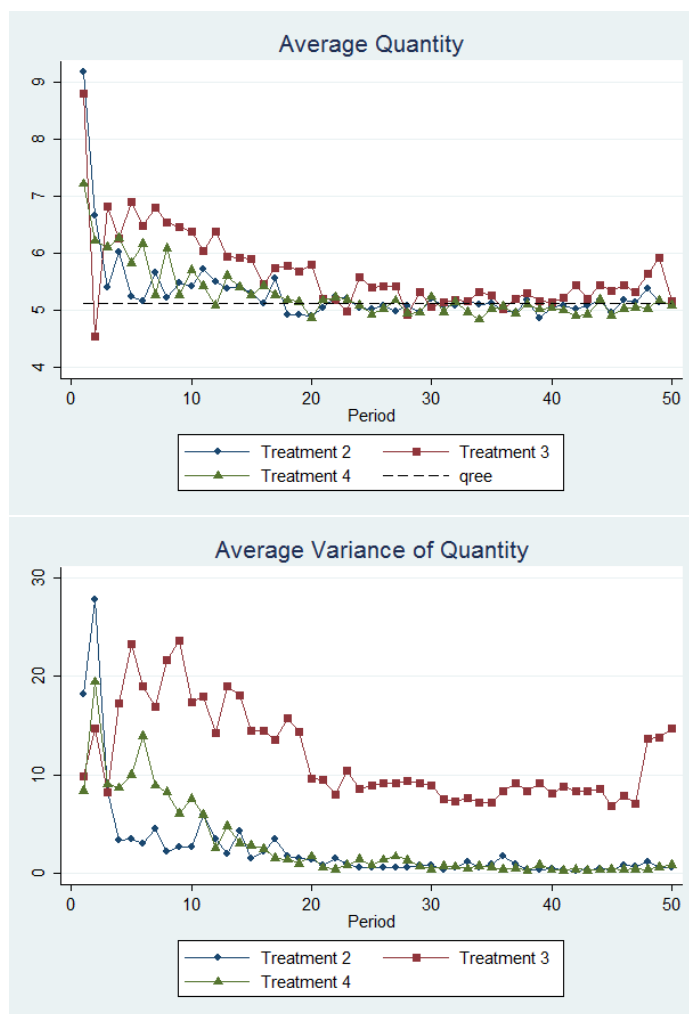


Figure 8: Upper panel: the average individual supply in Treatments 1, 3 and 4. Bottom panel: the average variance of individual supply in Treatments 1, 3 and 4.

The average supplies in treatment 2, 3 and 4 are plotted against the REE supply in the top panel of Figure 8. As with prices, we see that quantity in treatment 3 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 8 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

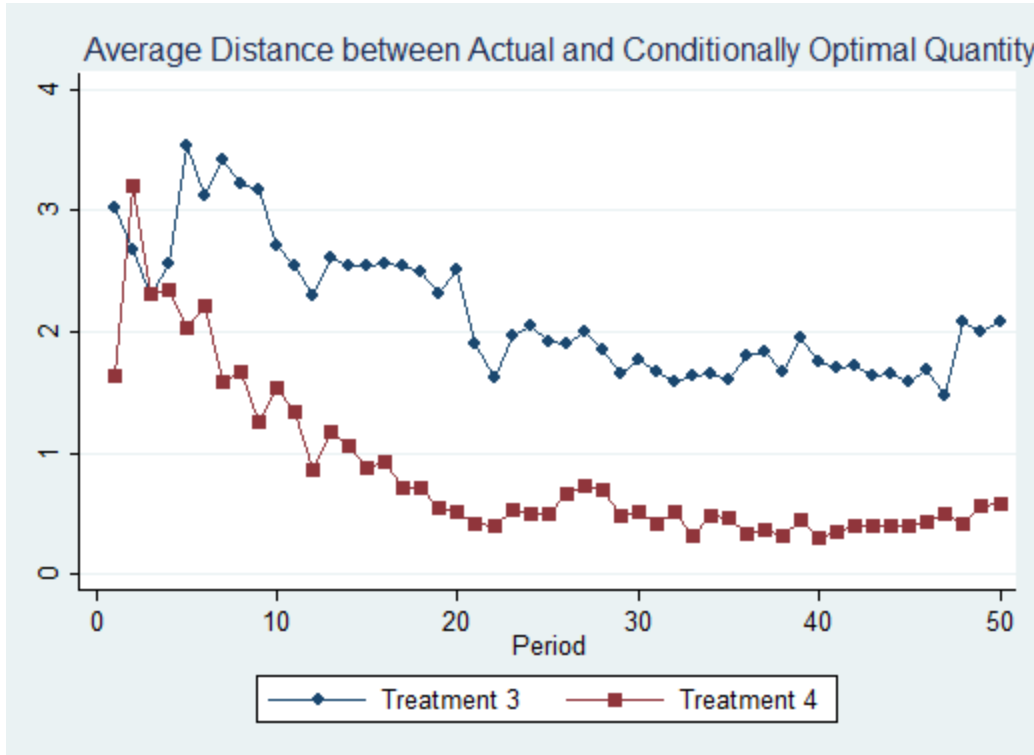


Figure 9: The average distance between actual supply and the conditionally optimal supply in Treatments 3 and 4.

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 9 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 production managers make decisions optimally, this difference should be zero.

Figure 9 reveals that the production managers in treatment 4 on average make supply decisions that are closer to the conditionally optimal quantity choice 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 does 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 and that is why we work with it. We discard any estimations for production decisions with serial correlation in the error term leaving 13 (out of 42) successful estimations for Treatment 3, and 18 (out of 36) successful estimations for Treatment 4. The results are 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 the 5% significance level while c_1 is not significantly different from zero. In this case, the subject is essentially supplying constant amount to the market each period.

2. Conditional optimal supply, if c_1 is significantly different from 0 at the 5% significance level, c_0 is not significantly different from zero and the null hypothesis $c_1 = 1/6$ cannot be rejected at the 5% level. In this case the subject chose to supply the conditionally optimal quantity for the given price forecast.
3. Hybrid strategy, if both c_0 and c_1 are significantly different from zero. In this case, the subject probably chooses a constant as a psychological anchor, and adjusts it a little for different expected price levels.

All the successful estimations can be classified according to one of these three types. Figure 10 shows the shares of the three different types of production strategies in treatments 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. 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 using all available information (i.e., they do not make systematic mistakes) and 2) Given these expectations, agents solve optimization problems and their solutions then determine actual price realizations, that is, there is



Figure 10: Distribution of estimated production strategies in treatments 3 and 4. Here C denotes use of the constant supply strategy, O denotes use of the conditionally optimal supply strategy and H denotes use of the hybrid strategy.

belief-outcome interaction. These two dimensions have been previously addressed *separately* in learning to forecast experiments (LtFE) and in learning to optimize experiments (LtOE). In this paper we design comparable LtFE and LtOE treatments for the same model, and we 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 a cobweb economy with negative feedback.

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 combined LtFE+LtOE design of treatment 3 is the least reliable and slowest to converge to REE. However, the latter treatment would seem to correspond most closely to what is expected of agents in rational expectations models, i.e., that individual agents are good at both forecasting *and* optimizing.

The estimation of individual forecast rules suggests that there is not much difference in the price prediction strategies subjects use across the different treatments of our experiment. 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 study 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 (close to RE predictions) than behavior 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 do not converge, at least in the learning to forecast experiments that have been used to date in such positive feedback environments.

Reference

1. Adam, K., "Experimental Evidence on the Persistence of Output and Inflation," *Economic Journal* 117, (2007), 603-636.
2. Arifovic, J., "The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies," *Journal of Political Economy* 104, (1996), 510-541.
3. Bao, T., C. Hommes, J.H. Sonnemans and J. Tuinstra, "Individual Expectations, Limited Rationality and Aggregate Outcomes," *Journal of Economic Dynamics and Control*, forthcoming.
4. Blinder, A.S., and J. Morgan, "Are Two Heads Better Than One? Monetary Policy by Committee," *Journal of Money, Credit and Banking* 37, (2005), 789-811.
5. Bornstein, G., and I. Yaniv, "Individual and Group Behavior in the Ultimatum Game: Are Groups more 'Rational' Players?," *Experimental Economics* 1, (1998), 101-108.
6. Cason, T., and V. Mui, "A Laboratory Study of Group Polarization in the Team Dictator Game," *Economic Journal* 107, (1997), 1465-1483.
7. 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.
8. Crockett, S., and J. Duffy, "A Dynamic General Equilibrium Approach to Asset Pricing Experiments," Working Paper, 2010.
9. Duffy, J., "Macroeconomics: A Survey of Laboratory Research", working paper, University of Pittsburgh, Working Paper, 2008.
10. Ezekiel, M., "The Cobweb Theorem," *Quarterly Journal of Economics* 52, (1938), 255-280.
11. 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.

12. 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, (1993), 119-137.
13. Heemeijer, P., C. Hommes, J.H. Sonnemans 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.
14. Hommes, C.H., J.H. Sonnemans, J. Tuinstra and H. van de Velden, "Coordination of Expectations in Asset Pricing Experiments," *Review of Financial Studies*, 18 (2005), 955-980.
15. Hommes, C.H., J.H. Sonnemans, J. Tuinstra and H. van de Velden, "Learning in Cobweb Experiments." *Macroeconomic Dynamics* 11, (2007), 8-33.
16. Huck, S., H.T. Normann, and J. Oechssler, "Learning in Cournot Oligopoly—An Experiment," *Economic Journal* 109, (1999) C80-C95.
17. Kocher, M., and M. Sutter, "When the 'Decision-maker' Matters: Individual versus Team Behavior in Experimental 'Beauty Contest' Games," *Economic Journal* 115, (2005), 200-223.
18. Lei, V., C.N. Noussair and C.R. Plott, "Nonspeculative Bubbles in Experimental Asset Markets: Lack of Common Knowledge of Rationality vs. Actual Irrationality," *Econometrica* 69, (2001), 831-859.
19. 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.
20. Marimon, R. and S. Sunder, "Indeterminacy of Equilibria in a Hyperinflationary World: Experimental Evidence," *Econometrica* 61, (1993), 1073-1107.
21. Marimon, R. and S. Sunder, "Expectations and Learning under Alternative Monetary Regimes: An Experimental Approach," *Economic Theory* 4, (1994), 131-62.
22. Marimon, R., S.E. Spear and S. Sunder, "Expectationally Driven Market Volatility: An Experimental Study," *Journal of Economic Theory* 61, (1993), 74-103.

23. Muth, J.F., "Rational Expectations and the Theory of Price Movements," *Econometrica* 29, (1961), 315-335.
24. Rubinstein, A., "Instinctive and Cognitive Reasoning: A Study of Response Times," *Economic Journal* 117, (2007), 1243-1259.
25. Offerman, T., J. Potters and J.H. Sonnemans, "Imitation and Belief Learning in an Oligopoly Experiment," *Review of Economic Studies* 69, (2002), 973-97.
26. Sargent, T.J., "Evolution and Intelligent Design," *American Economic Review* 98, (2008), 537.
27. Smith, V.L., G.L. Suchanek and A.W. Williams, "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets," *Econometrica*, 56, (1988), 1119-1151.
28. Vega-Redondo, F., "The Evolution of Walrasian Behavior," *Econometrica* 65, (1997), 375-384.

A Appendix

A.1 Experimental Instructions

A.1.1 Instruction for the Forecaster

(Treatment 1; modifications for Treatment 4 shown in ().)

General information

In this experiment you participate in a market. Your role in the market is as 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 for 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 who is your partner, are assigned to form a team and you will keep cooperating together 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 will 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, 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 all your 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. (Your 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 lead a firm to produce a larger quantity, which increases the supply and vice 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 for the product 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 period 1. When all forecasters have submitted their price predictions for the first period, the firms (production managers) will then determine the quantity to supply, and the market price for period 1 will be determined and made public to all participants. Based on your price prediction (and your partner's production decisions), your earnings for period 1 will be calculated.

Subsequently, you are asked to enter your prediction for period 2. When all participants have submitted their predictions (and production decisions) for the second period, the market price for that period will be determined and 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

(Your payoff depends on both the performance of your forecasting task and your partner's production decision task. You and your partner will each get one half of the payoff from the forecasting task and one half of the payoff from 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 earn. You will earn 0 points if your prediction error is larger than 7. You have a Payoff Table which shows the points you can earn for different prediction errors. (Your forecast accuracy will not affect your payoff from the production task, but more accurate forecasts may

help your partner to make better production decisions).

We will pay you in cash at the end of the experiment based on the points you earned in all 50 periods. You earn 1 Euro for each 2600 points you earn.

A.1.2 Instruction for the Production Manager

(Treatment 2; modifications for treatment 4 shown in (), modifications for treatment 3 shown in [].)

General information

In this experiment you participate in a market. Your role in the market is as 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 who is your partner, are assigned to form a team and you will keep cooperating together 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 the quantity of 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 past prices, a table containing the history of your past decisions and the payoff (of your team) in (forecasting and) production tasks. (Your 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 decisions of the producers. Usually, higher price predictions lead a firm to produce a larger quantity, which

increases the supply and vice 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 for the product 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 have 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

Your payoff depends on the ([both the]) 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. You are given a Payoff Table for the Production Task on 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. Your payoff from the forecasting task is decreasing in your prediction error, and you can also refer to the other payoff table for the Prediction Task 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 are calculated then you can read the last part of these instructions, which you can skip otherwise.

We will pay you in cash at the end of the experiment based on the points you earned in all 50 periods. You earn 1 Euro for each 2600 points you earn.

The equation that determines your payoff from the production task

The payoff for the 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 a baseline "salary". But 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

market	λ	p-value $ \lambda \geq 1$	μ	R^2	MSE	Equilibrium	p-value Wald Test	Durbin-Watson
p11	0.1863	0.0000	24.9741	0.1113	7.6785	30.6921	0.9378	2.9626
p12	0.1698	0.0000	25.5817	0.1495	3.9079	30.8131	0.8073	3.1043
p13	0.2108	0.0000	24.3511	0.1838	5.0071	30.8546	0.7586	2.7229
p14	0.1827	0.0000	25.2209	0.185	3.3689	30.8573	0.6910	2.5946
p21	0.6059	0.0000	11.5545	0.5238	18.1876	29.3174	0.3667	2.6236
p22	0.4733	0.0000	15.3094	0.3085	30.4796	29.0684	0.2673	2.6857
p23	0.0056	0.0000	30.0469	0.0001	6.3699	30.2151	0.1540	1.9987
p24	0.5191	0.0000	14.8545	0.4287	17.8924	30.8868	0.9015	2.7677
p25	0.2666	0.0000	22.3075	0.1239	21.9634	30.4165	0.7306	2.5310
p26	0.5411	0.0000	14.4979	0.5543	9.7218	31.5946	0.3778	2.7611
p27	0.2891	0.0000	22.0542	0.2156	9.6062	31.0231	0.6383	3.0010
p31	0.4189	0.0000	19.0151	0.3137	18.3482	32.7227	0.0597	2.4415
p32	0.4488	0.0000	19.4622	0.2966	35.1703	35.3101	0.0028	2.4295
p33	0.4197	0.0000	14.0236	0.195	62.8946	24.1653	0.0008	2.1841
p34	0.2126	0.0000	22.9448	0.0734	25.5324	29.1408	0.0818	2.7188
p35	0.2351	0.0000	19.5698	0.0896	18.9326	25.5849	0.0000	2.2688
p36	0.1740	0.0000	23.023	0.0974	6.4048	27.8739	0.0000	2.6301
p37	0.7604	0.0012	5.5085	0.6605	20.8671	22.9902	0.0055	2.7725
p41	0.2182	0.0000	25.077	0.1149	12.9444	32.0757	0.0408	2.6890
p42	0.0674	0.0000	28.1808	0.0093	17.4663	30.2188	0.4225	2.3721
p43	0.3334	0.0000	19.9712	0.2167	14.7826	29.9588	0.3500	2.5853
p44	0.4994	0.0000	14.6045	0.3604	23.8659	29.1717	0.2652	2.5729
p45	0.3080	0.0000	20.7341	0.2789	6.6841	29.9616	0.1507	2.3922
p46	0.3530	0.0000	19.2175	0.199	23.6375	29.7035	0.3385	2.6715

Table 7: Results from estimation of the equation $p_{j,t} = \lambda_j p_{j,t-1} + \mu_j + \epsilon_{j,t}$ for each market j of the different treatments. The results start with p_{11} , which is the first market of treatment 1, and continue through p_{46} , which is the sixth market of treatment 4. Also reported are Wald tests of strong convergence and Durbin-Watson test statistics for first-order serial correlation.

C Identified Forecasting Rules

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: Estimation results for subjects in Treatment 1 who could be successfully categorized by one of the two forecasting rules. In the “Type” column, “A” means adaptive rule while “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: Estimation results for subjects in Treatment 3 who could be successfully categorized by one of the two forecasting rules. In the “Type” column, “A” means adaptive rule while “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: Estimation results for subjects in Treatment 4 who could be successfully categorized by one of the two forecasting rules. In the “Type” column, “A” means adaptive rule while “C” means contrarian/trend extrapolation rule.

D Estimated Supply Strategies

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: Estimated coefficients of the supply strategy used by subjects in Treatment 3. In the “Type” column, “C” means use of the constant supply strategy, “O” means use of the conditionally optimal supply strategy and “H” means use of the hybrid strategy.

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: Estimated coefficients of the supply strategy used by the subjects in Treatment 4. In the “Type” column, “C” means use of the constant supply strategy, “O” means use of the conditionally optimal supply strategy and “H” means use of the hybrid strategy.

E Payoff Tables

E.1 Payoff Table for Forecasters

Below is the payoff table for the forecasting task in treatments 1,3 and 4.

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=?	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 11: The payoff table for forecasters.

E.2 Payoff Table for Production Managers

Below is the payoff table for the forecasting task in treatments 2,3 and 4.

Payoff Table for Production Task (2500 points=1 hour)

predecessor	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	526	432	335	228	117	0
1	1200	1198	1190	1176	1156	1130	1098	1060	1016	966	910	848	780	706	628	540	448	350	246	136	20
2	1200	1190	1192	1179	1160	1135	1104	1067	1024	975	920	859	792	719	640	556	464	367	264	156	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	1186	1168	1145	1116	1081	1040	993	940	881	816	745	668	586	498	401	300	199	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	526	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	1108	1072	1029	980	925	864	797	724	645	560	468	372	269	160
9	1200	1206	1206	1200	1188	1170	1148	1116	1080	1038	990	936	876	810	738	660	576	484	390	298	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	1180	1150	1121	1087	1048	1003	952	894	831	762	687	606	521	428	331	230
14	1200	1211	1216	1215	1208	1185	1155	1126	1092	1053	1008	957	900	837	768	693	612	526	432	335	228
15	1200	1212	1218	1218	1212	1200	1182	1158	1128	1089	1043	991	934	871	803	728	653	571	480	384	280
16	1200	1213	1220	1221	1216	1205	1188	1166	1136	1097	1050	1003	946	883	815	740	665	583	492	396	290
17	1200	1214	1224	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	1178	1152	1119	1080	1035	984	927	864	795	720	639	552	459	360
19	1200	1216	1228	1230	1228	1220	1206	1184	1158	1126	1086	1040	988	940	878	810	736	655	570	478	380
20	1200	1217	1228	1233	1232	1225	1212	1190	1164	1132	1092	1046	994	946	884	816	742	661	576	484	390
21	1200	1218	1230	1236	1236	1230	1218	1200	1176	1146	1106	1060	1008	958	900	834	760	679	594	504	410
22	1200	1219	1234	1242	1242	1236	1224	1206	1182	1154	1114	1068	1016	966	906	840	766	685	600	510	416
23	1200	1220	1238	1248	1248	1242	1230	1212	1188	1158	1118	1072	1020	970	910	844	770	689	604	514	420
24	1200	1221	1240	1250	1250	1244	1232	1214	1190	1160	1120	1074	1022	972	912	846	772	691	606	516	422
25	1200	1222	1242	1254	1254	1248	1236	1218	1194	1164	1124	1078	1026	976	916	850	776	695	610	520	426
26	1200	1223	1244	1256	1256	1250	1238	1220	1196	1166	1126	1080	1028	978	918	852	778	697	612	522	428
27	1200	1224	1246	1258	1258	1252	1240	1222	1198	1168	1128	1082	1030	980	920	854	780	699	614	524	430
28	1200	1225	1248	1260	1260	1254	1242	1224	1200	1170	1130	1084	1032	982	922	856	782	701	616	526	432
29	1200	1226	1250	1262	1262	1256	1244	1226	1202	1172	1132	1086	1034	984	924	858	784	703	618	528	434

[This table continues at the back of this page.]

Figure 12: The payoff table for production managers, page 1.

production	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	1157	1128	1083	1023	975	912	843	768	697	630
31	1500	1526	1550	1569	1579	1580	1572	1556	1534	1507	1475	1438	1395	1346	1291	1230	1163	1090	1011	926	835
32	1800	1826	1851	1869	1879	1880	1872	1856	1834	1807	1775	1738	1695	1646	1591	1530	1463	1390	1311	1226	1135
33	1800	1830	1854	1872	1884	1890	1891	1884	1872	1856	1834	1807	1775	1738	1695	1646	1591	1530	1463	1390	1311
34	1800	1831	1856	1872	1886	1895	1896	1891	1880	1865	1846	1825	1804	1783	1762	1741	1720	1700	1680	1660	1640
35	1800	1832	1858	1878	1892	1900	1902	1906	1912	1918	1924	1929	1934	1939	1944	1949	1954	1959	1964	1969	1974
36	1800	1833	1860	1881	1896	1905	1906	1906	1906	1906	1906	1906	1906	1906	1906	1906	1906	1906	1906	1906	1906
37	1800	1834	1862	1884	1900	1910	1914	1917	1920	1923	1926	1929	1932	1935	1938	1941	1944	1947	1950	1953	1956
38	1800	1835	1864	1887	1904	1915	1920	1923	1926	1929	1932	1935	1938	1941	1944	1947	1950	1953	1956	1959	1962
39	1800	1836	1866	1890	1908	1920	1926	1930	1933	1936	1939	1942	1945	1948	1951	1954	1957	1960	1963	1966	1969
40	1800	1837	1868	1893	1912	1925	1932	1936	1939	1942	1945	1948	1951	1954	1957	1960	1963	1966	1969	1972	1975
41	1800	1838	1870	1896	1916	1930	1938	1942	1945	1948	1951	1954	1957	1960	1963	1966	1969	1972	1975	1978	1981
42	1800	1839	1872	1900	1921	1935	1944	1948	1951	1954	1957	1960	1963	1966	1969	1972	1975	1978	1981	1984	1987
43	1800	1840	1874	1903	1925	1940	1949	1954	1957	1960	1963	1966	1969	1972	1975	1978	1981	1984	1987	1990	1993
44	1800	1841	1876	1906	1929	1945	1954	1959	1962	1965	1968	1971	1974	1977	1980	1983	1986	1989	1992	1995	1998
45	1800	1842	1878	1909	1933	1950	1959	1964	1967	1970	1973	1976	1979	1982	1985	1988	1991	1994	1997	2000	2003
46	1800	1843	1880	1911	1936	1954	1963	1968	1971	1974	1977	1980	1983	1986	1989	1992	1995	1998	2001	2004	2007
47	1800	1844	1882	1914	1940	1959	1968	1973	1976	1979	1982	1985	1988	1991	1994	1997	2000	2003	2006	2009	2012
48	1800	1845	1884	1917	1944	1964	1974	1979	1982	1985	1988	1991	1994	1997	2000	2003	2006	2009	2012	2015	2018
49	1800	1846	1886	1920	1948	1969	1979	1984	1987	1990	1993	1996	1999	2002	2005	2008	2011	2014	2017	2020	2023
50	1800	1847	1888	1923	1952	1974	1984	1989	1992	1995	1998	2001	2004	2007	2010	2013	2016	2019	2022	2025	2028
51	1800	1848	1890	1926	1956	1979	1989	1994	1997	2000	2003	2006	2009	2012	2015	2018	2021	2024	2027	2030	2033
52	1800	1849	1892	1929	1960	1984	1994	1999	2002	2005	2008	2011	2014	2017	2020	2023	2026	2029	2032	2035	2038
53	1800	1850	1894	1932	1964	1989	1999	2004	2007	2010	2013	2016	2019	2022	2025	2028	2031	2034	2037	2040	2043
54	1800	1851	1896	1935	1968	1994	2004	2009	2012	2015	2018	2021	2024	2027	2030	2033	2036	2039	2042	2045	2048
55	1800	1852	1898	1938	1972	2000	2010	2015	2018	2021	2024	2027	2030	2033	2036	2039	2042	2045	2048	2051	2054
56	1800	1853	1900	1941	1976	2005	2015	2020	2023	2026	2029	2032	2035	2038	2041	2044	2047	2050	2053	2056	2059
57	1800	1854	1902	1944	1980	2010	2020	2025	2028	2031	2034	2037	2040	2043	2046	2049	2052	2055	2058	2061	2064
58	1800	1855	1904	1947	1984	2015	2025	2030	2033	2036	2039	2042	2045	2048	2051	2054	2057	2060	2063	2066	2069
59	1800	1856	1906	1950	1988	2020	2030	2035	2038	2041	2044	2047	2050	2053	2056	2059	2062	2065	2068	2071	2074
60	1800	1857	1908	1953	1992	2025	2035	2040	2043	2046	2049	2052	2055	2058	2061	2064	2067	2070	2073	2076	2079

Figure 13: The payoff table for production managers, page 2.