

UC Irvine

UC Irvine Previously Published Works

Title

Adaptive versus eductive learning: Theory and evidence

Permalink

<https://escholarship.org/uc/item/9d72r75v>

Authors

Bao, Te
Duffy, John

Publication Date

2016-04-01

DOI

10.1016/j.euroecorev.2015.12.007

Peer reviewed

Economic Review

Elsevier Editorial System(tm) for European

Manuscript Draft

Manuscript Number: EER-D-15-00679R1

Title: Adaptive versus Eductive Learning: Theory and Evidence

Article Type: FLA Regular Paper

Keywords: Rational Expectations, Adaptive Learning, Eductive Learning, Experimental Economics

Corresponding Author: Prof. John Duffy,

Corresponding Author's Institution: University of California, Irvine

First Author: Te Bao

Order of Authors: Te Bao; John Duffy

Manuscript Region of Origin: NETHERLANDS

Abstract: Adaptive and eductive learning are two widely used ways of modeling the process by which agents learn a rational expectation equilibrium (REE). In this paper we report on an experiment where we exploit differences in the conditions under which adaptive and eductive learning converge to REE so as to investigate which approach provides the better description of the learning behavior of human subjects. Our results suggest that the path by which the system converges appears to be a mixture of both adaptive and eductive learning model predictions.



Department of Economics
3151 Social Science Plaza
Irvine, CA 92697-5100
(949) 824-8341
(949) 824-2182 FAX
duffy@uci.edu

December 2, 2015

Jörg Oechssler, Editor
European Economic Review

RE: Ms. No. EER-D-15-00679

Dear Jörg,

Thank you for the opportunity to further revise the above referenced paper for *European Economic Review*.

Te Bao and I have made a number of changes that address all of the comments of the one 'old' referee –referee 1- and the one 'new' referee – referee 3 – that you recently sent to us. The changes we have made in response to these comments are detailed in the attached response to the referees. We hope that you like the new version of our paper and that you can now accept it for publication in *EER*.

We thank you once again for your time and consideration and we look forward to hearing of your decision.

Sincerely,

A handwritten signature in black ink, appearing to read 'John Duffy'.

John Duffy
Professor of Economics

Reply to Referee 1

We thank referee 1 for his constructive and helpful comments. Here we detail how we have addressed these comments in the revised draft of our paper.

Major points

[1] *In the current draft, the authors still consider T3 and T4 as two not eductively stable cases. This is not true, only T4 is not eductively stable. The distinction is even more important now that they found that with T4, but not with T3, convergence fails. The authors cannot reject eductive learning when convergence occurs in their experiment with T3. Why do they keep talking about T3 as an eductively unstable outcome? (pag. 14 "If Hypothesis 1 is rejected...") Do the authors still believe that they can test the original Guesnerie's criterion with a continuum of agents in an experiment with three agents? The condition in Gaballo (2013) is a generalization of the ones in Guesnerie (1992), and not a different case related to a finite number of agents (as they keep saying in the middle of pag. 5 and on top of pag. 11). N can be infinite! In the limit of $N \rightarrow \infty$ the condition is the same of Guesnerie (1992)! Therefore, to experimentally test the eductive learning criterion introduced by Guesnerie (1992), they do not have other choice that to refer to the conditions depending on N provided by Gaballo (2013)...unless they have an infinitely large lab!*

We agree there was some misunderstanding in the previous version. In the updated version, we discuss both Guesnerie (1992) and Gaballo (2013) when discussing the eductive stability conditions and clarify that the latter is a generalization of the former as it takes explicit account of the finite number, N, of agents in the economy. In the new draft, T3 is treated as an eductively stable treatment throughout the paper. We also state clearly that in the limit as $N \rightarrow \infty$ the condition in Gaballo (2013) is the same as the competitive market version in Guesnerie (1992).

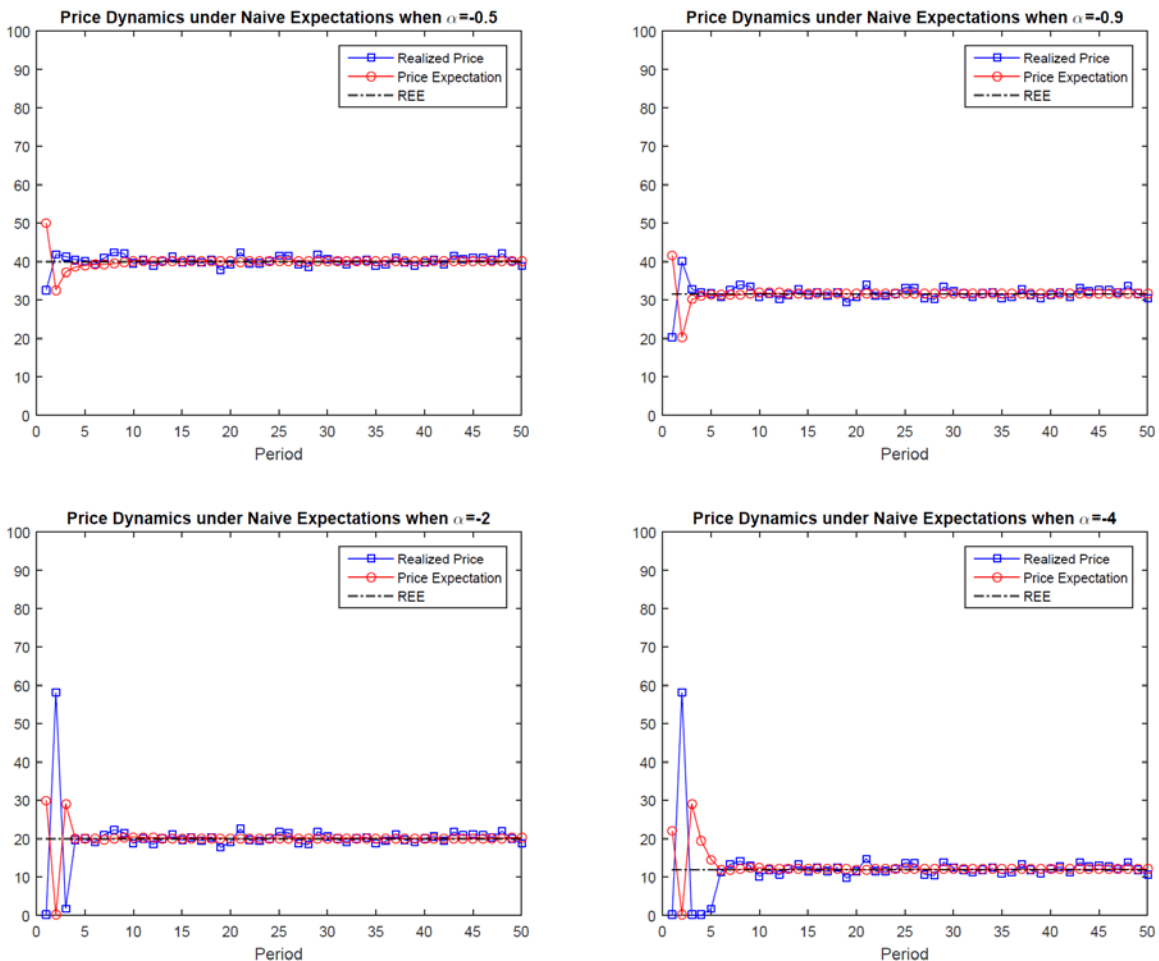
[2] *If the authors want to give relevance to the condition $\alpha = -2$ they can discuss a third possibility, that is, agents revise their expectation according to a naive rule. In this case, T1 and T2 only would be stable, whereas T3 and T4 unstable. This is different from adaptive learning for which all treatments are stable, and also different from eductive learning for which - with three agents - only T4 is unstable. In this context, the authors could put naive expectations in relation with Guesnerie's original conditions, but, again, they cannot say that convergence in T3 is in contrast with the eductive principle: this is definitely wrong!*

Thanks for this suggestion. We have added a brief discussion of naïve expectations in sections 4 and 5.6 and we now label T3 only as a treatment with strong negative feedback and instability under naive expectations. We also make sure that there is no place in the paper where T3 is referred to as eductively stable.

[3] *The authors discuss about the speed of convergence without quoting any theoretical result about the speed of convergence in the adaptive learning literature. The authors find that as α decreases the speed of convergence increases: is this consistent with the theory?*

We checked the Evans and Honkapohja (2001) book and there is no direct discussion on the relationship between the speed of convergence and the size of α under adaptive learning. We checked the larger literature and found that Marcat and Sargent (1992) report that the convergence speed decreases when the absolute value of the coefficient for expectation feedback increases when agents use adaptive learning to form expectations on inflation rate. Their results are based simulations.

Using the cobweb model of our paper, we simulated the case of a single agent using adaptive learning who starts from $p_1^e = REE + 10$ for all four values of α . The simulation shows that it is indeed the case that the number of periods before convergence increases (i.e., convergence speed decreases) as $|\alpha|$ increases if agents are adaptive learners. The plots from this exercise are show below.



Mathematically, if we ignore the noise term and refer to mathematical Appendix C: when an agent starts from an arbitrary deviation from the REE, Δ , at each period t , his deviation from the REE is $\frac{\alpha(\alpha+1)(\alpha+2)\dots(\alpha+t-2)}{1 \times 2 \times 3 \times \dots \times (t-1)} \Delta$. This expression reveals that the smaller the absolute value of α , the smaller is the absolute value of the numerator, and the quicker it goes to 0. We now add a brief one sentence mention of this finding to the paper in section 5.2 on p. 27. We would be happy to elaborate on this analysis as shown above if you think it is warranted.

Minor points

- There should be a minus on the left hand term of (4).

Thank you. We corrected that.

- There are some (very few) typos (like pag. 12 "A noted" instead of "As noted"). The paper is generally well written, although some explanations and tables towards the end are less interesting (at least to me); some of them are based on quite rough classifications (for example the criterion about convergence that generated my concern about pictures 5 and 8 in the previous version).

We corrected the typo on page 12 and went through the paper again a number of times to check typos. We acknowledge that some classifications towards the end are a bit rough when we use self-reported data on strategies. But that is an issue that typically occurs for all experimental data of this type. We think it is worthwhile to include this data as it reveals information about subjects' self-retrospection, and our estimation exercise has greater validity if there is some accord with the self-reported data.

-The introduction is still 90% motivation and 10% results. Moreover, from the introduction, results appear to be more generic than that which they are. The authors need to be more informative and overall sharper. The introduction could improve by far.

Thank you for this perspective. We have rewritten parts of the introduction to make the questions we ask clearer and to give more weight and precision to our major experimental findings.

Reference

Marcet, Albert, and Thomas J. Sargent. "Speed of convergence of recursive least squares learning with ARMA perceptions." Economics Working Paper of University Pompeu Fabra, (1992).

Reply to Referee3

We thank referee 1 for his constructive and helpful comments. Here we detail how we have addressed these comments in the revised draft of our paper.

(1) The description of the literature in Section 2 is in general accurate, but two minor points are:

p4: "First, subjects in all of these prior studies do not precisely know the model of the economy ..." This is true for most learning-to-forecast market experiments, but e.g. Sonnemans and Tuinstra (2010) compare different treatments with limited as well as complete information.

p4: "Second, all prior experiments using the cobweb model employ a group design ..."

This is true for most cobweb experiments, but e.g. Hommes et al. (2000) discuss an individual learning-to-forecast experiment.

We replaced "all these works" by "most of these works" in this paragraph, and include the discussion on Sonnemans and Tuinstra (2010) and Hommes et al. (2000) in a qualifying footnote.

(2) p11, Eq. (4): it seems to me that a "-" sign is missing and the Eq. should read

Thank you. We corrected this typo.

(3) p18: "it seems that several subjects directly solved for the REE using ". On which evidence is this statement based?"

It was based on two facts. (1) there are a few, though not many subjects who choose option 2 "I make calculations based on the value of α " in the strategy question in the post experiment questionnaire. (2) in one session, a subject commented to the experimenter "this experiment is simple; I can just solve the linear equation and then give the same prediction every period" after the experiment was over. We include this anecdote in a footnote on the same page.

(4) p21, Table 3: For the four different treatments, the REE levels p^ are different with p^* lower for more unstable treatments. Might it be that lower price levels are more difficult to learn, is it more asymmetric w.r.t. the price interval? Why were these levels chosen differently? An alternative might have been to fix the REE level p^* across all treatments.*

In this dimension of the design, we face a trade-off between fixing the intercept μ of the expectation feedback function or the REE level p^* . Given that the intercept μ is included in the instructions to provide complete information of the economy, we decided to make intercept μ the same in all treatments so that the instructions are more comparable across treatments—they only vary in terms of α . In addition, if we want to have the same REE in all treatments, we would need a much larger μ , e.g. $\mu = 200$ for the treatment $\alpha = -4$. In that case, it would be difficult to set the axes of the different treatments on the same scale. Thus, while we acknowledge that the lower REE might be harder to learn, in particular in learning to forecast experiments with qualitative instead of quantitative information on the law of motion of the economy, we choose to fix μ in our experiment to address other concerns.

(5) The questionnaire to get some information about the use of adaptive versus educative learning seems important, in particular to get insight in any inductive reasoning of subjects. But it might have been better to ask open questions, in order not to impose or suggest any strategies to the subjects. Moreover, out of the four choices in the questionnaire, two (options 2 and 3) are classified as educative and only one (option 1) as adaptive learning. Might this difference bias the results?

This is a very good question. Actually, we asked the restricted form multiple choice question partly because each of us had prior experience asking open-ended questions. While the data from open-ended questions is richer, it is also more likely to have uninformative answers such as “I try to maximize profit” or “I used a very wise strategy” (among many other possibilities). For these reasons we chose to script the available answers and have subjects choose which answer best described their behavior.

In our view, it should be acceptable to assign two options to educative learning and one to adaptive learning. Options 2 and 3 are associated with individual and social versions of educative learning (or individual vs. common knowledge of rationality) respectively. It seemed better to combine them to capture all of the educative learners. To partly compensate for this, we also made the adaptive learning option the first option to give it more salience.

(6) The authors focus on a cobweb model with negative expectations feedback. I wonder what the results might be under positive feedback, with $0 < \alpha < 1$, particularly for a coefficient close to 1 (near unit root). Earlier laboratory experiments have shown that near unit root positive feedback markets do not easily converge. But both adaptive learning and educative learning predict convergence to RE for $0 < \alpha < 1$. The authors could perhaps briefly discuss this in the concluding section how adaptive versus educative learning would fare in a positive feedback environment.

We fully agree with the referee’s comment. It would be interesting to study a version of our model with positive feedback, where $0 < \alpha < 1$. Previous works, e.g., Hommes et al. (2005, 2008) find that positive feedback tends to generate bubbles and crashes in asset prices when the slope coefficient is close to 1. A useful future study would redo our analysis with positive feedback and $0 < \alpha < 1$ in which case both adaptive and educative learning theories predict convergence to REE. The focus of such studies would be on whether adaptive or educative learning was a better predictor of *individual* behavior in the positive feedback environment. We include this suggestion for future research in the last paragraph of the conclusion.

Adaptive versus Eductive Learning: Theory and Evidence *

Te Bao[†] and John Duffy[‡]

December 2, 2015

Abstract

Adaptive and eductive learning are two widely used ways of modeling the process by which agents learn a rational expectation equilibrium (REE). In this paper we report on an experiment where we exploit differences in the conditions under which adaptive and eductive learning converge to REE so as to investigate which approach provides the better description of the learning behavior of human subjects. Our results suggest that the path by which the system converges appears to be a mixture of both adaptive and eductive learning model predictions.

JEL Classification: C91, C92, D83, D84

Keywords: Rational Expectations, Adaptive Learning, Eductive Learning, Experimental Economics.

*We thank the editor, Jörg Oechssler, three anonymous referees and participants in the Barcelona GSE Summer Forum on Theoretical and Experimental Macroeconomics, the Expectations in Dynamic Macroeconomic Models Workshop at the Bank of Finland, the 14th SAET Conference in Tokyo and the 2015 CEF conference in Taipei for helpful comments and suggestions on earlier drafts. We gratefully acknowledge financial support from the National Science Foundation of China under Grant No. 71301174.

[†]Division of Economics, Nanyang Technological University and Department of Economics, Econometrics and Finance, University of Groningen. P.O. Box 800, 9700 AV Groningen, The Netherlands. Email: t.bao@rug.nl.

[‡]Department of Economics, University of California, Irvine, CA, 92697 USA Email: duffy@uci.edu.

1 Introduction

How do agents learn a rational expectations equilibrium (REE) if they do not initially find themselves in such an equilibrium? This important, foundational question has generated a large literature in macroeconomics (see, e.g., surveys by Sargent (1993), Grandmont (1998), Evans and Honkapohja (2001)). In this paper we focus on two different but related approaches to addressing this question.

Perhaps the most widely used approach, beginning, e.g., with Bray (1982) is to suppose that agents are boundedly rational *adaptive* learners and to ask whether their use of a given real-time adaptive learning model that allows for a REE as a possible solution converges in the limit to that REE. An alternative, off-line approach, originally advocated by Guesnerie (1992, 2002), is to suppose that learning is a mental process involving collective introspection that takes place in some notional time and that leads agents to understand and instantly coordinate upon or “educate” the REE solution.¹ Both approaches to learning place restrictions on model parameters under which learning agents can learn the REE. The restrictions for learnability under adaptive learning may differ from the restrictions under educative learning and this difference serves as one means of identifying the learning process that agents are using. A second difference between the two learning approaches concerns the speed of convergence: if agents are adaptive learners who start out with beliefs different from the REE and if the REE is stable under adaptive learning, then it will generally take some time before agents’ forecasts converge to the REE value. By contrast, if agents are educative learners who understand the model and if the REE is educatively stable, then agents’ forecasts should instantaneously converge to the REE value. A third difference between the two learning approaches is that educative learning is an inherently social type of learning as it relies heavily on common knowledge of the rationality of other actors, whereas adaptive learning does not explicitly consider the behavior of other agents. The contribution of this paper is that we test the importance of all of these different features of the two learning approaches using controlled laboratory experiments with human subjects.

Evans (2001) highlights the differences between the adaptive and educative approaches to learning and invites empirical and experimental testing of the different

¹These two approaches are also considered as two broad classes for belief formation in a recent survey of expectations in macroeconomics by Woodford (2013).

theoretical predictions. Specifically he writes:

“Which is the appropriate way to model economic agents will ultimately be a matter for empirical and *experimental* research. It is likely that the answer *depends* on the circumstances, for example, in experiments, on the details of the setting and the types of *information* provided to the subjects. A plausible conjecture is that when a model is simple and transparent as well as *eductively stable*, agents will coordinate rapidly on the REE....If a model has no eductively stable REE, but has an REE that is adaptively stable, then a plausible conjecture is that there will still be convergence to the REE, at a rate governed by the accumulation of data....The eductive results provide a caution, however, that coordination in such cases may not be robust.” (Evans 2001, p. 581 emphasis added).

In this paper we follow up on Evans’s invitation to compare adaptive versus eductive learning approaches. Indeed, the manner in which agents might go about learning a rational expectations equilibrium is an important, fundamental yet unresolved issue; there are many ways to model this learning process and it would be useful to have a consensus on which approach (or combination of approaches) are more empirically valid than others.² Understanding the manner in which agents learn is also important for policy purposes. For instance, if agents can educe a REE prior to making decisions via the mental, collective introspective process described by eductive learning, then policy ineffectiveness propositions that arise under rational expectations may have full standing. However, if agents instead learn a REE adaptively over time, then policy interventions may be effective in the short-run in the determination of economic variables. Thus, the *manner* in which agents learn is an important empirical question.

Ideally, one would like to address the question of how agents form expectations using non-experimental field data, but unfortunately, properly incentivized field data on individual-level expectations are not generally available. Survey evidence, e.g., on inflationary expectations, consumer confidence, etc. *are* available, but these data are not properly incentivized in that constant rewards or, more typically, no reward at all for participation in such surveys, yield poor incentives to report truthful beliefs. Even setting such incentive problems aside, to use survey data on expectations one

²Here we focus on just two approaches, but there are several other approaches including Bayesian learning, evolutionary learning and near-rational (calculation-cost) learning.

would have to know the precise structure of the economic environment, i.e., the data generating process in which agents were forming their expectations, knowledge that is typically unavailable and/or subject to some dispute. For these reasons, a laboratory experiment offers the better means of collecting data on expectations as truthful revelation can be properly incentivized (using quadratic loss scoring rules) and the control of the laboratory allows for precise implementation of the model environment (data generating process) in which agents' expectations matter for the realizations of economic variables.³

We report results from an experiment where subjects form predictions about the single market price in a simple cobweb economy. The price forecasts determine actual prices according to an equilibrium reduced form equation of the model which is known to all subjects. Our main treatment variable is the slope parameter of this reduced form model, which is also known to subjects. We consider values for this parameter such that the REE solution is both adaptively and eductively stable and other parameterizations where the REE solution is adaptively stable but not eductively stable, taking into account the number of agents in the economy. As eductive stability relies upon common knowledge of rationality, we also vary whether several agents' forecasts are aggregated to determine market prices as in an oligopoly setting, or whether each subject acts as a monopolist so that their own price forecast uniquely determines the market clearing price.

Our main finding is that support can be found for both types of learning, adaptive and eductive, in our experimental data. More precisely, we find that prices converge reliably to the REE in experimental markets where both learning approaches predict convergence to the REE. Some of these markets converge in the first period as is consistent with eductive learning, while other markets take many more periods to converge, as is consistent with adaptive learning. In our treatment where the REE is learnable under adaptive learning but is not stable under eductive learning, we find that some markets still converge to the REE within 50 periods but that the majority of markets fail to converge to REE in this case. This non-convergence finding only obtains when agents form forecasts as part of a group of agents, as in our oligopoly setting, and not when they form forecasts independently of others as in our

³See Duffy (2016) for further arguments in support of using laboratory evidence to evaluate macroeconomic models and assumptions as well as a survey of the literature on experimental macroeconomics.

monopoly setting, and this distinction is consistent with the social learning aspect of the educative learning criterion. Finally, looking at individual price forecasts we find that we can characterize a majority of the subjects in our experiment as either adaptive or educative learners, with a roughly equal division between the two types.

The organization of the remainder of paper is as follows: section 2 discusses related literature, section 3 presents the theoretical model, section 4 discusses the experimental design and hypotheses, section 5 reports the experimental results, and section 6 concludes.

2 Related Literature

Our experiment employs a “learning-to-forecast” experimental design (as pioneered by Marimon and Sunder (1993)), that involves versions of the cobweb market model with negative feedback (or strategic substitutes). Hommes et al. (2000) provide the first experimental test of such a cobweb economy, and this study has been followed up by Sonnemans et al. (2004), Hommes et al. (2007), Heemeijer et al. (2009), Sonnemans and Tuinstra (2010), Bao et al. (2012, 2013) and Beshears et al. (2013). Hommes (2011) surveys the literature. The differences between the present study and those earlier papers are as follows. First, subjects in most of these prior studies do not precisely know the model of the economy (i.e., the data generating process). This lack of knowledge makes it impossible for subjects to apply educative learning as that type of learning (as demonstrated below) requires full knowledge of the model thereby enabling introspective reasoning about the proper price forecast. By contrast, subjects in our experiment *are* informed about the model economy and so they *can* in principle apply educative learning, or even directly solve for the REE using the perfect foresight condition. Second, most prior experiments using the cobweb model employ a group design where both learning and strategic uncertainty can influence the speed of the convergence to the REE. By contrast, we have both a group (“oligopoly market”) treatment *and* an *individual*-decision making (“monopoly” market) treatment. The monopoly treatment rules out strategic uncertainty as a factor and serves as an important baseline for assessing the extent of rational play among subjects. Third, all prior learning-to-forecast experiments involving linear cobweb models use a data generating process for the market price equation that has a coefficient on expected

prices, α , that is smaller than 1 in absolute value.⁴ Finally, we explicitly test restrictions on the stability of REE under the two different learning approaches. By contrast, most of the existing experimental literature on whether and how agents learn a REE in cobweb economies has been concerned with characterizing the type distribution of (adaptive) learning behaviors without regard to any stability under learning criteria, and certainly not a comparison of two different learning criteria, as we present in this paper.

Since subjects in our experiment are informed that prices are determined as a function of price forecasts, (i.e., they know the data generating process) our experiment is also related to an experimental literature on “guessing” or “beauty contest” games (see, e.g., Nagel (1995), Duffy and Nagel (1997), Ho et al. (1998) Grosskopf and Nagel (2008) among others). In these guessing games, subjects are asked to guess a number. The winning guess, (which is similar to a market price and which yields the winner a large prize), is a known function of the average guess (or average opinion which is similar to the mean price forecast). A main finding from this literature is that the winning number is initially very far from the rational expectations equilibrium though it gets closer to that prediction with experience.

In our experiment we consider forecasting by a group of three subjects (in our “oligopoly” setting) as well as an individual forecasting treatment (our “monopoly” setting) and we also examine whether our results for the monopoly treatment are closer to the REE relative to the oligopoly treatment. The winning number in beauty contest games is typically a linear function, $\rho \times$ the mean guess, where $\rho \in (0, 1)$ which is similar to a learning-to-forecast experiment with positive feedback (strategic complements). There are also some guessing game experiments where $\rho \in (-1, 0)$ such as Sutan and Willinger (2009). The difference between our work and their paper is that we provide a more detailed description of the model that generates the price that agents are seeking to forecast and we vary the value of ρ (equivalently, our α) so as to explore the implications of differing stability results under the adaptive and eductive approaches to learning. As in a typical macroeconomic model, we also add

⁴Sonnemans and Tuinstra (2010) compare different treatments with limited as well as complete information. Hommes et al. (2000) discuss an individual learning-to-forecast experiment. Hommes et al. 2007 report on a *nonlinear* cobweb model experiment where $|\alpha| > 1$ in the REE. However, none of these studies explicitly tests for differences between adaptive and eductive learning model predictions.

a shock term to the price determination equation, a setup that is not usually found in beauty contest/guessing games. Our framework could also be readily extended to a real intertemporal design where these shocks are autocorrelated.

Since we have both monopoly (individual decision-making) and oligopoly (group decision-making) treatments, our paper is related to experimental studies on oligopoly markets, for example, Bosch-Doménech and Vriend (2003), Huck et al. (1999), and Offerman et al. (2002). These oligopoly market experiments use learning-to-optimize designs where subjects submit a quantity choice directly and price forecasts are not elicited. By contrast, we ignore quantity choices and focus instead on price forecasts (expectations) using a learning-to-forecast design.⁵ Our monopoly versus oligopoly design is also helpful for investigating the important role played by common knowledge of rationality in educative learning. In this respect our paper is related to other experimental studies exploring the role of common knowledge of rationality in market settings, for example, the “money illusion” experiments of Fehr and Tyran (2005, 2007, 2008) and the asset market experiments of Akiyama et al. (2012, 2013).

Finally, we note that Gaballo (2013) has generalized Guesnerie’s (1992, 2002) perfectly competitive market version of the educative learning criterion to market settings with an arbitrary but finite number of firms, $N > 0$, so that the conditions under which an REE is educatively stable explicitly depends on N . We compare and contrast Gaballo’s educative stability conditions with Guesnerie’s educative stability conditions in our experimental design and analysis and find that Gaballo’s conditions matter for convergence to REE and the stability of price forecasts.

⁵In a learning-to-forecast design, subjects submit a price forecast and a computer program uses that forecast to optimally determine the subject’s quantity decision. By contrast, in a learning-to-optimize design, subjects submit a quantity choice directly; their price forecast is not elicited, though it is implicit in their quantity decision. See Bao et al. (2013) for a comparison of these two approaches.

3 Theoretical Model

3.1 Cobweb economy

We consider a simple version of a cobweb model as presented in Evans and Honkapohja (2001) that is based on Bray and Savin (1986). We chose the cobweb model as it was the model originally used by Muth (1961) to illustrate the notion of a REE. It is also simple enough to explain to subjects and has the critical feature that expectations matter for outcomes, here price realizations, while outcomes can in turn matter for beliefs as subjects interact under the same model environment repeatedly. The cobweb model is one of demand and supply for a single perishable good and consists of the two equations:

$$\begin{aligned}D_t &= a - bp_t, \\S_t &= cp_t^e + \eta_t.\end{aligned}$$

Here, D represents demand, S supply, a , b , and c are parameters, which are usually assumed to be positive, p_t is the period t price of the good, $p_t^e = E_{t-1}[p_t]$, and η_t is a mean zero supply shock.⁶

Assuming market clearing, the reduced form equation for prices is given by:

$$p_t = \mu + \alpha p_t^e + \nu_t, \tag{1}$$

where $\mu = \frac{a}{b}$, $\alpha = -\frac{c}{b}$, and $\nu_t = \frac{\eta_t}{b}$.

The system has a unique rational expectation equilibrium where

$$p_t^{e,*} = \frac{\mu}{1 - \alpha} \quad \text{and} \quad p_t^* = p_t^{e,*} + \nu_t. \tag{2}$$

3.2 Theoretical Predictions

As Evans (2001) shows, the unique REE of this model (2) is stable under adaptive learning (i.e., it is “learnable”) if $\alpha < 1$. However, under the educative learning

⁶Bray and Savin and Evans and Honkapohja use a somewhat richer model in which the supply equation, $S_t = cp_t^e + \delta w_{t-1} + \eta_t$, where w_{t-1} is an observable exogenous variable affecting supply, e.g., weather in period $t - 1$, that follows a known process (i.i.d. mean 0 or possibly AR(1)). For simplicity we study the case where $\delta = 0$, but we think it would also be interesting to study cases with such exogenous forcing variables as well.

approach, the REE is learnable only if $|\alpha| < 1$ (See, e.g., Evans (2001) or Evans and Honkapohja (2001, section 15.4).⁷ We exploit this difference in stability conditions as one means of identifying which learning process, adaptive or educative, characterizes the price forecasts of our human subjects.

To be more precise about these two learning processes, adaptive learning consists of a general class of backward looking learning rules that make use of past information and the specific type of adaptive learning rule that we consider in this paper is “least squares learning” which is widely used. In supposing that agents learn in this particular adaptive fashion, we assume that they do not know or they ignore any information about the price determination equations of the economy. Instead, they start out by choosing a random prediction for the price in period 1, p_1^e . The adaptive agents’ “perceived law of motion” for the price at time t is that it is equal to some constant, a , plus noise, ϵ_t , i.e., $p_t^e = a + \epsilon_t$, which has the same functional form as the REE solution. Given this perceived law of motion and the assumption that the adaptive agents are least squares learners, it follows that, in each period $t > 1$, an agent’s price forecast is equal to the sample average of all past prices given the available history:

$$p_t^e = \frac{1}{t-1} \sum_{s=1}^{t-1} p_s. \quad (3)$$

Evans and Honkapohja (2001, section 2.3 and 2.4 for a simple linear case with an additional exogenous variable) provide a general proof, based on matrix operations, as to why the REE in this simple cobweb system is learnable via adaptive, least squares learning provided that $\alpha < 1$.⁸ For readers without prior knowledge about adaptive learning to capture the idea of this modeling approach, we provide an alternative proof for the non-stochastic version of this model (namely, ignoring the noise term ν_t since it has mean zero and small variance) based on mathematical induction in the

⁷We recognize that other learning approaches may impose different restrictions on the parameters of the cobweb economy to ensure converge to the REE. For example, Hommes and Wagener (2010) find that when agents use the evolutionary learning model of Brock and Hommes (1997), the market price may converge to a locally stable two cycle when $\alpha \in [\frac{1}{2}, 1]$.

⁸See also Evans and Honkapohja (2001), p. 149 for a more sophisticated non-stochastic non-linear negative feedback model with decreasing gain in the multivariate case based on Evans and Honkapohja, (2000).

Appendix.

For the experiment we parameterized the cobweb model as follows. We set $\mu = 60$ and $\nu_t \sim N(0, 1)$. We drew the sequence of random shocks, $\{\nu_t\}$ just once and then used the *same* shock sequence in all of our experimental sessions. Thus, μ and $\{\nu_t\}$ are the same in all of our experimental sessions. We consider four different values for the parameter α , which serves as the primary treatment variable in our experiment. Specifically, our four treatment values for α are $T1 : \alpha = -0.5$, $T2 : \alpha = -0.9$, $T3 : \alpha = -2$ and $T4 : \alpha = -4$. The REE price predictions associated with these four different choices for α are $T1 : p^{e,*} = 40$, $T2 : p^{e,*} = 31.58$, $T3 : p^{e,*} = 20$ and $T4 : p^{e,*} = 12$, respectively. The rationale for these different values for α has to do with differing predictions for the stability of the REE under the two different learning approaches.

Eductive learning has two versions, the basic, single-dimensional version found in Guesnerie (1992) and a more general, multi-dimensional version in Guesnerie (2002). Eductive learning is based on iterated elimination of strategies that are never best responses, in our case, the elimination of unlikely price forecasts.⁹ This iteration occurs in a competitive market environment where each individual producer has no market power (therefore, the rational expectations equilibrium corresponds to the *competitive equilibrium*). Each producer has perfect *individual rationality*, namely, each producer can perfectly solve for the rational expectations equilibrium of the system, and common knowledge of such rationality by other producers is also assumed.¹⁰ The eductive learning model describes the learning process by which agents iteratively eliminate non-rationalizable strategies (price forecasts) from their strategy space. If this process leads to elimination of all other strategies aside from predicting the rational expectations equilibrium, then the rational expectations equilibrium is said to be *eductively stable*. We would like to emphasize that in this sense, eductive learning is a *social learning* process (Vriend, 2000) as the agents' learning behavior is also conditioned upon others' decisions, while adaptive learning is essentially an *individual learning* process where agents learn from the history of the realized market

⁹Eductive learning is the counterpart of rationalizability (Bernheim 1984, Pearce 1984) in games.


¹⁰We acknowledge that the ability of agents to solve the REE from equation (1) is not explicitly included as a part of individual rationality as defined by Guesnerie (1992) on page 1257. But immediately after that, on page 1258, Guesnerie makes the comment that the rational expectations equilibrium is also the unique Nash Equilibrium of the game. Therefore, the subjects should be able to find the REE/Nash Equilibrium by solving the game as if they are perfectly rational.

price and interact with other agents only indirectly.

In this paper, we focus on the single-dimensional version of eductive learning as in Guesnerie (1992) since our cobweb model is a simple, one product market. The eductive learning process works in the following way: in notional period 0, each agent knows that it is rational to forecast $p_t = \frac{\mu}{1-\alpha}$. Further, since all agents know the data generating process, $p_t = \mu + \alpha p_t^e$, that prices should be non-negative and that $\alpha = -\frac{c}{b} < 0$, it follows that agents can logically rule out the possibility that any other agent would forecast prices greater than μ ¹¹, and so it can be regarded as common knowledge that no one is going to forecast $p_t > \mu$. In notional period 1, knowing that no one is going to make a price forecast that is larger than μ , and substituting this constraint into the price equation, $p_t = \mu + \alpha p_t^e$, agents should all infer that no one will forecast prices lower than $\mu + \alpha\mu = (1 + \alpha)\mu$. In notional period 2, using the same reasoning, agents can rule out price forecasts greater than $\mu + \alpha(\mu + \alpha\mu) = (1 + \alpha + \alpha^2)\mu$, etc. More generally, in notional period t , the new forecast boundary created by this iterative process will be $(1 + \alpha + \alpha^2 + \dots + \alpha^t)\mu$. If $|\alpha| < 1$, this process will tighten the interval range of possible price forecasts to a single point, the REE. When $|\alpha| < 1$, in the limit, the two boundaries become a single point, $\lim_{t \rightarrow \infty} \sum_{s=1}^t \alpha^s \mu = \frac{\mu}{1-\alpha}$. This iterative, notional time eductive learning process is illustrated in Figure 1. By contrast, when $\alpha < -1$, agents cannot rule out any price forecasts starting from notional period 1, because $\mu + \alpha\mu < 0$. Hence, in the case where $\alpha < -1$, the REE is not eductively stable, though as noted earlier, it is stable under the adaptive learning dynamics. This difference in the criteria for convergence to REE is the main hypothesis that our experiment addresses.

Thus far, our theoretical conditions for stability under eductive learning are derived under the perfectly competitive market model of Guesnerie (1992), which presumes that all firms are atomistic and have no market power. However, since it is not possible to host infinitely many producers in the lab, we should also consider how the finite number of firms in our oligopoly treatment matters for the stability of REE under eductive learning. Recently, Gaballo (2013) has derived a generalized version of the eductive stability condition that can be used for the oligopoly cobweb market

¹¹Since the literature on eductive learning typically assumes that $\alpha < 0$ as the starting point, when we prove that the REE is not eductively stable when $|\alpha| > 1$, we only focus on $\alpha < -1$, case because the $\alpha > 1$ case is already ruled out by the assumption that $\alpha < 0$.



The Iterative Process of Eductive Learning in Notional Periods

Figure 1: An illustration of the iterative process in notional time under eductive learning. The process creates a boundary, B_t , in notional time period t , and excludes numbers that are larger/smaller than this boundary in even/odd notional periods. When $|\alpha| < 1$ the boundaries move closer to each other with each iteration so that the interval eventually tightens to a single point, i.e., $\lim_{t \rightarrow \infty} \sum_{s=1}^t \alpha^s \mu = \frac{\mu}{1-\alpha}$.

model with a finite number $N > 0$ of firms. This condition is:

$$-\frac{N}{N-2} < \alpha_N < 1. \quad (4)$$

Here α_N refers to the expectation feedback coefficient in the price generation function when there are N firms (by contrast with α for the perfectly competitive version of the same model). Notice that in the limit, as $N \rightarrow \infty$, this condition becomes exactly the same as the competitive market version of Guesnerie (1992). According to Gaballo's criteria, our oligopoly $T3$ treatment, where $N = 3$ and $\alpha = -2$ is eductively stable since $-2 \in [-3, 1]$. To address whether the finiteness of the firm population size matters for our experimental findings, we added treatment $T4$ where $\alpha = -4$ as a robustness check.¹² With $N = 3$ and $\alpha = -4$, Gaballo's criterion for eductive stability is *not* satisfied. If oligopoly markets (with $N = 3$) converge to the REE in treatment $T4$, that finding would serve as strong evidence in favor of adaptive over eductive learning. In this paper, we will use the condition defined in (4) as the condition for eductive stability.

We note that there is another type of stability condition, namely stability under naïve expectations, where $p_t^e = p_{t-1}$. This naïve view of expectation formation predicts that the REE will be stable or unstable depending on whether $|\alpha| < 1$ or $|\alpha| \geq 1$.¹³

¹²We thank a referee who recommended us to run this additional treatment.

¹³For example, when $\alpha = -2$, naïve expectations results in convergence to an oscillatory two

While naïve expectations is not the focus of our study, it also predicts differences in the stability of the REE between our treatments T1-T2 and T3-T4. Later, in section 5.6, we check whether subjects are behaving in this naïve manner and find little evidence for naïve expectations, but we mention this possibility here as a further way of distinguishing between the predictions of our various experimental treatments.

As noted above, our primary treatment variable is α , which takes on four different values $\{-.5, -.9, -2, -4\}$ in our experiment. As a second treatment variable, we also differentiate between monopoly and oligopoly markets. In the oligopoly treatment there are $N = 3$ firms forming price forecasts in each market while in the monopoly treatment there is a single firm forming price forecasts in each market. Since educative learning is a social learning process, only the oligopoly market design provides an environment where *both* adaptive learning and educative learning can be properly implemented. In that case, both learning theories predict that subjects will learn the REE in treatments $T1$ ($\alpha = -.5$) and $T2$ ($\alpha = -.9$) and $T3$ ($\alpha = -2$), but under treatment $T4$ ($\alpha = -4$) the REE is “learnable” only if agents are adaptive learners. This is our main hypothesis to be tested. In addition, as a robustness check on the oligopoly market behavior, we also explore a monopoly market treatment involving individual decision-making. By contrast with the oligopoly design, in the monopoly design, the REE is learnable under educative learning in all four treatments, since educative learning assumes that agents have no difficulty solving for the REE by themselves. In other words, *the predictions of the two learning models diverge only in the oligopoly $T4$ treatment*. Finally, we also consider differences in the speed of convergence to the REE; when an REE is stable under educative learning, convergence should, in principle, be instantaneous while under adaptive learning, it can take several periods for the economy to converge to the REE depending on initial price forecasts.

4 Experimental Design and Hypotheses

We employ a 4×2 experimental design where the treatment variables are (1) the four different values of the slope coefficient, α , and (2) the number of subjects in one experimental market: either just one subject—the “monopoly” market case or

period cycle where prices alternate between 0 and 60.

three subjects—the “oligopoly” market case. The monopoly versus oligopoly design is helpful for investigating the role of strategic uncertainty in decision-making. Eductive learning assumes first that agents are perfectly rational, i.e., that they can perfectly solve for the REE when they have complete information about the model (as in our design) and second, that there is common knowledge of rationality, namely, that each player knows the other players are rational, each knows that the others know that they are rational and so on. In monopoly markets, common knowledge of rationality is not an issue since the single agent faces no uncertainty about his own level of rationality. Deviations from REE in the monopoly market are violations of the assumption of individual rationality. However, since individual rationality is a precondition for common knowledge of rationality, the extent to which deviations from individual rationality arise in the monopoly markets is useful for understanding behavior in the oligopoly markets where common knowledge of rationality plays a critical role. If there is some doubt as to whether other market participants can form rational expectation forecasts, as evidenced by monopoly market forecasts that are not immediately equal to REE values, then forecasting the REE price in the oligopoly setting may no longer be a best response. If it takes some time for subjects to learn the REE in the monopoly setting, then it should take at least as much time or longer for subjects to learn the REE in the oligopoly setting, as group members would first have to establish that there was common knowledge of rationality.

As noted earlier, our four treatment values for α are as follows:

Treatment 1 ($T1$): weak negative feedback treatment, $\alpha = -0.5$. The REE of this treatment is stable under both adaptive and eductive learning, as well as when agents use naïve expectations.

Treatment 2 ($T2$): medium negative feedback treatment, $\alpha = -0.9$. The REE of this treatment is stable under both adaptive and eductive learning, as well as when agents use naïve expectations.

Treatment 3 ($T3$): strong negative feedback treatment, $\alpha = -2$. The REE of this treatment is stable under both adaptive and eductive learning, but not stable when agents use naïve expectations.

Treatment 4 ($T4$): very strong negative feedback treatment, $\alpha = -4$. The REE of this treatment is stable only under adaptive learning, but not stable under

eductive learning or when agents use naïve expectations.

Our experiment makes use of a learning to forecast (“LtFE”) experimental design. Subjects play the role of an advisor who makes price forecasts only. In our monopoly treatment, each market consists of a single forecaster. The time t price forecast of subject i , $p_{i,t}^e$, determines the price forecast for that market, i.e., $p_t^e = p_{i,t}^e$ which is then used to determine the actual market price, p_t , for that monopoly market according to equation (1). By contrast, in the oligopoly treatment, each market involves three forecasters. We use the mean of the three subjects’ individual price forecasts for period t as the market price forecast for period t , i.e., $p_t^e = \frac{1}{3} \sum_{i=1}^3 p_{i,t}^e$, which is then used to determine the actual market price, p_t , for each oligopoly market, again according to equation (1). In both treatments, subjects are paid according to the accuracy of their *own* price forecast and are thus incentivized to provide good price forecasts. Since subjects are paid according to their forecast accuracy and *not* according to the profit from their production decision, they have *no incentive* to take their *market power* into consideration. Under our incentive system, predicting the REE (competitive equilibrium) is the only Nash Equilibrium of the prediction game, where the forecasting error is minimized, and the payoff is maximized for every subject in the same market. Thus, the atomistic assumption underlying Guesnerie’s eductive stability condition may not be unreasonable in our setting. If subjects were instead paid according to the profit their firm earned, they might have an incentive to play the Cournot-Nash equilibrium, or the collusive equilibrium, which are different from the competitive market settings that are typically used in both the adaptive and eductive learning literatures.

An important issue is how to allow for eductive learning, which is an off-line, notional time concept. It is not clear how to capture or measure this type of learning in real time. Here we focus on the stability differences as pointed out by Evans (2001) as our main test of whether agents are eductive or adaptive learners. Nevertheless, we wanted to be sure that subjects understood the model and had sufficient time for the introspective reasoning required under eductive learning. Under adaptive learning, there is no assumption that agents know the model while under eductive learning it is assumed that agents do know the model. What we have chosen to do is to fully inform subjects about the model, in particular about the price determination equation, (1) – see the written experimental instructions in the Appendix for the details on how this information was presented to subjects. Thus the agents in our model have more

information than is typically assumed under adaptive learning specifications, but at the same time, they have all the information they need to be educative learners. We felt that, in order to put the two learning approaches on an equal footing for comparison purposes we would have to eliminate any informational differences between the two learning approaches, which could serve as a confounding factor in our analysis. Thus we provide subjects with complete and common information about the model across all of our six treatments. Further, we did not impose any time limits on subjects' decision-making so as not to limit the type of introspective reasoning associated with the educative approach. Indeed, we captured subjects' decision time as a variable in order to better understand whether there were any differences in decision time across treatments $T1-T4$, or between individuals and groups in our monopoly and oligopoly treatments.

4.1 Hypotheses

Based on the theoretical analysis of the section 3, we formulate the following testable hypotheses. The underlying prior is that agents are adaptive learners, and so the results favor educative learning if the hypotheses are rejected.

Hypothesis 1. *The market price and price expectations in all treatments converge to the unique rational expectation equilibrium given in (2).*

As in section 2, both adaptive and educative learning theories predict that market price forecasts and market prices will converge to the REE in treatments $T1-T3$. In treatments $T4$, the REE is learnable under adaptive learning. It is learnable under educative learning in the monopoly design but not in the oligopoly design for the competitive version of educative learning. If Hypothesis 1 is rejected, and forecasts and prices do not converge to the REE in treatment $T4$, the experimental results favor educative learning over adaptive learning.

Hypothesis 2. *Given that the market price and price expectations converge to the REE, convergence never takes place in the first period of the experiment.*

Since convergence under adaptive learning takes place more gradually and in real time while educative learning happens in notional time, convergence should take place

in the first real period that is incentivized for monetary payment if agents are educative learners, or after a few periods if agents use adaptive learning and depending on initial conditions. If Hypothesis 2 is rejected, the experimental results favor educative learning over adaptive learning.

Hypothesis 3. *Agents spend no more time in making their decisions in each period of treatment T4 as compared with each period of treatments T1 – T3 in the oligopoly design.*

Since educative learning can involve considerable introspective reasoning in notional time, which we take to be the period prior to the first incentivized market forecasting period, it may require more time for agents to reach a decision. In particular, the REE is predicted to be more difficult to learn under educative learning in treatment T4 as compared with treatments T1 – T3 in the oligopoly design. Since decision time is a typical measure of the cognitive cost to agents of decision-making, if Hypothesis 3 is rejected, it suggests that making a decision in treatment T4 is indeed more difficult than in treatments T1 – T3, and the results would thus favor educative learning over adaptive learning.

4.2 Number of Observations

The experimental data was collected in a number of sessions run at the CREED Lab of the University of Amsterdam. Subjects had no prior experience with our experimental design and were not allowed to participate in more than a single session of our experiment. Each session consisted of 50 periods over which the treatment parameters for that session were held constant (i.e., we used a “between subjects” experimental design). Table 1 provides a summary of the number of subjects or markets (independent observations) for each of our six treatments. Note that in the monopoly treatment, each subject acted alone in a single market, so the number of subjects equals the number of independent observations (markets) in that setting. By contrast, in the oligopoly treatment, each market consisted of three firms (subjects), so while we have more subjects in our oligopoly treatments, we nevertheless have fewer 3-firm markets (independent observations) for the oligopoly treatments. Each session averaged about 1 hour and 10 minutes in duration. The average payoff was 21.9 euros across all four monopoly treatments and 18.8 euros across all four oligopoly

Treatment	Monopoly	Oligopoly	Total No.
Conditions	No. Markets /Subjects	No. Markets / Subjects	Subjects
T1	14 / 14	10 / 30	44
T2	12 / 12	10 / 30	42
T3	13 / 13	11 / 33	46
T4	14 / 14	10 / 30	44
Totals	53 / 53	41 / 123	176

Table 1: Number of Markets (Independent Observations) and Subjects in the eight Treatments of the Experiment

treatments.

4.3 Computer Screen

Figure 2 shows the computer screen we developed for the experiment in the treatment where $\alpha = -0.5$. Subjects were asked to enter a forecast number in the box and then to click “send” to submit their forecast in each period. Since the price and price expectation were restricted to be non-negative, the range of possible prices should be $[0, 60]$ according to equation (1).¹⁴ However, restricting the price forecast range to $[0, 60]$ would be equivalent to directly imposing the first step in the educative learning process. Therefore, we restricted the price forecast range to $[0, 100]$ in the experiment, which is less suggestive. The subjects are told in the instructions that neither the price nor their price predictions can be negative. The upper forecast bound of 100 was not indicated in the written instructions, but subjects would see a pop-up window indicating that a forecast larger than 100 is not allowed if they attempted to submit a price forecast that was greater than 100. Notice that the computer decision screen presented subjects with information and graphs of past prices, their own prior price forecasts as well as realizations of shocks. The screen was refreshed with updated information once all subjects had submitted forecasts and the market price was determined. Notice further that at the top of the decision screen, the price determination equation (1) with the treatment specific value of α was always available for subjects to view and it appeared next to the input box where

¹⁴If $p_t^e > 0$ and given that $\alpha < 0$ it follows that $p_t = 60 + \alpha p_t^e < 60$.

they were asked to submit their price prediction in each period.

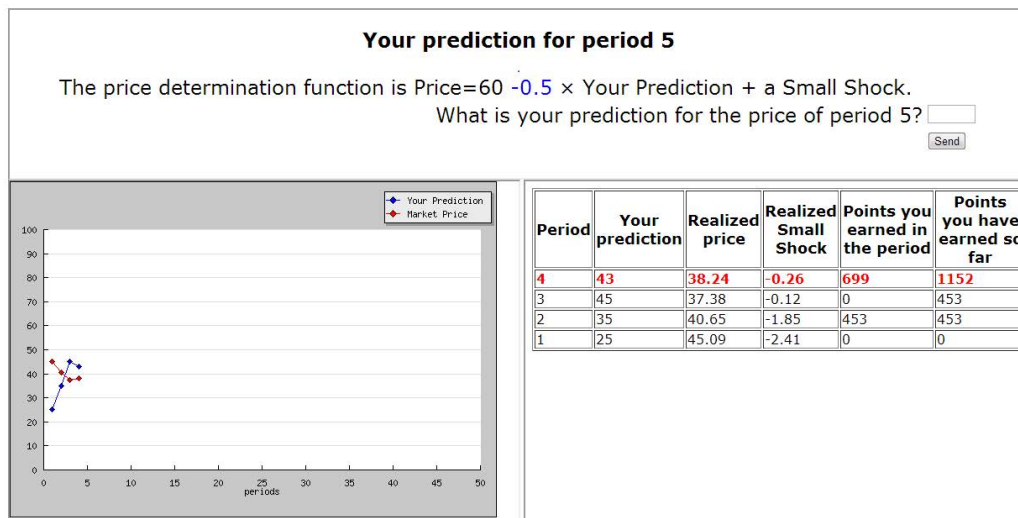


Figure 2: The computer decision screen used in the experiment for the treatment where $\alpha = -0.5$ and the subject is a monopolist in the market. Note: the price and price expectations shown in this figure are random inputs by the authors for illustration purposes and are not taken from any experimental data.

4.4 Payoff Function

The period payoff function for subjects (in points) is a decreasing quadratic function of their market price prediction error, and was given by:

$$\text{Payoff for subject } i\text{'s forecast in period } t = \max \left[1300 - \frac{1300}{49} (p_t - p_{i,t}^e)^2, 0 \right]. \quad (5)$$

This payoff function was carefully explained to subjects in the written instructions. Notice that subjects earn 0 points if their own, individual price forecast error is greater than 7 and they earn a maximum of 1300 points for a perfect forecast. Subjects' point totals from all 50 periods were converted into to euros at the end of each session at a known and fixed rate of 1 euro for every 2600 points. Thus, over 50 periods, each subject's maximum earnings were $(1300 \times 50) / 2600 = 25$ euros.

5 Experimental Results

5.1 Market Dynamics

5.1.1 Monopoly Markets

Figure 3 plots the *average* price expectation (forecast) against the respective REE price expectation using data from *all* markets (individual observations) of each of the four monopoly treatments. According to the theoretical analysis in Section 3, convergence in terms of the market price or in terms of price expectations are equivalent. We choose to plot the time path of average *price expectations* instead of the *market prices* to limit the influence of the noise term ν_t .¹⁵ We observe that the mean price expectation in all four treatments appears to converge to the REE, although at different speeds (we will quantify this speed of convergence later in section 5.2). The adjustment towards the REE is observed to be fastest in *T1* and slowest in *T4*.

Figure 4 plots the disaggregated price expectation paths in each *individual* market for each of the four monopoly treatments against the respective REE. As this figure reveals, it may take 25 periods or more for some markets to converge, e.g., treatments *T3 – T4*, and there are a lot of extreme outcomes, e.g., price forecasts such as 0 and 60. From these results we preliminarily conclude that adaptive learning is correct in predicting the convergence outcome across all four treatments including treatments *T3-T4*, however the time path of convergence for some markets often resembles a real-time demonstration of the educative, introspective learning process, in particular, the dampened cycling of price expectations over time in some markets. Further, if we look at self-reported strategies from a questionnaire solicited from subjects following the end of the experiment (as we do later in section 5.6), it seems that several subjects directly solved for the REE using $p^{e,*} = \frac{\mu}{1-\alpha}$, which indicates that those subjects were applying educative reasoning.¹⁶

Table 2 reports the mean and variance of price expectations across all markets in

¹⁵The results for market prices are thus very similar to the results for market price forecasts, but due to the noise term, market prices have a higher variance.

¹⁶At the end of one experimental session during the payment phase, a subject commented: “this experiment is simple; I just solve the linear equation and then give the same prediction in every period.”




Figure 3: The average expectation against the REE in each of the four treatments of the monopoly design.

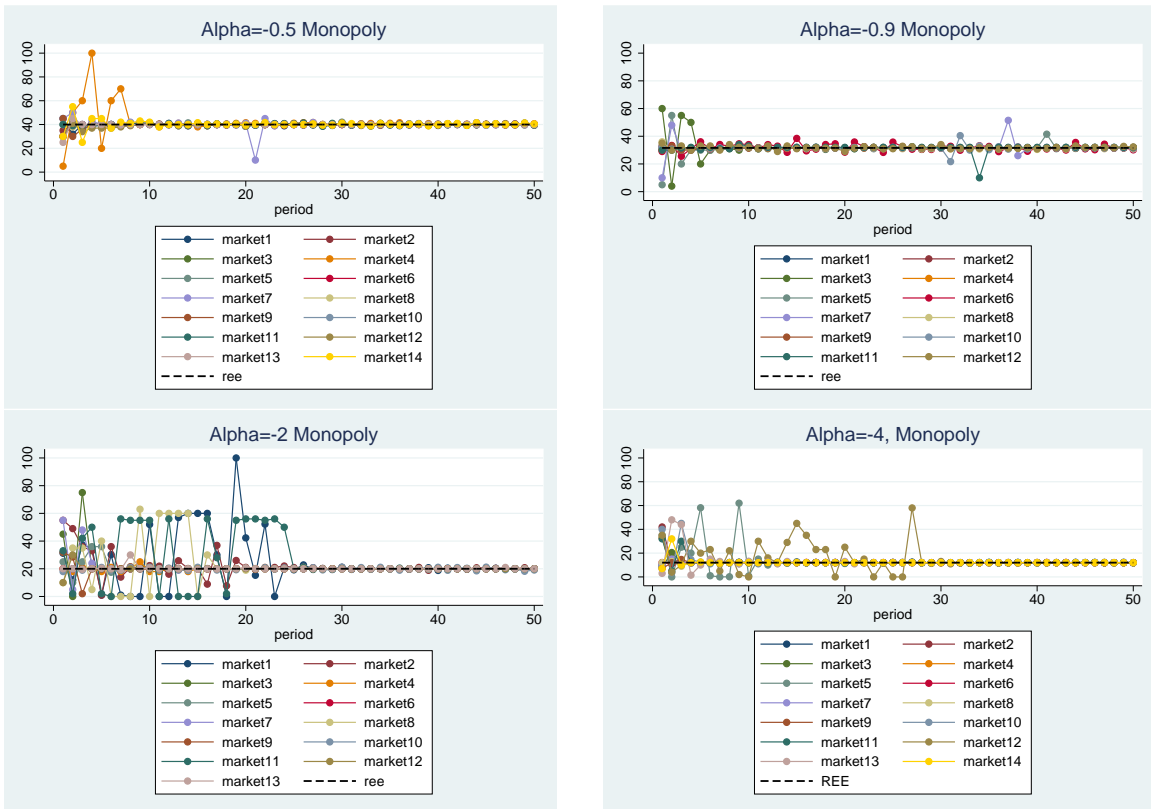


Figure 4: Disaggregated price expectations against the REE in each individual market of the four treatments of the monopoly design.

each of the four monopoly treatments for the entire sample of 50 periods as well as for the first 25 and the last 25 periods of the sample. Confirming the impression given in Figures 3-4, we observe that, on average, price forecasts converged to the REE in each monopoly treatment and that the variance in these forecasts over all 50 periods was greatest in treatment $T3$. When attention is restricted to the last 25 rounds, the variance is greatest for treatment $T4$.

Treatment	REE	Periods 1-50		Periods 1-25		Periods 26-50	
		Mean	Variance	Mean	Variance	Mean	Variance
$\alpha = -0.5$	$p^* = 40$	40.02	14.18	39.97	28.30	40.07	0.23
$\alpha = -0.9$	$p^* = 31.58$	31.62	10.79	31.65	17.51	31.60	4.51
$\alpha = -2.0$	$p^* = 20$	21.46	90.76	22.88	172.70	20.03	0.13
$\alpha = -4.0$	$p^* = 12$	12.63	28.40	13.17	49.82	12.09	6.48

Table 2: Mean and variance of price expectations in each treatment ($\alpha = -0.5, -0.9, -2, -4$) of the monopoly setting.

5.1.2 Oligopoly Markets

Figure 5 plots the average price expectations (forecasts) against the respective REE price expectation using data from *all* markets of each of the four oligopoly treatments. We see that the average price expectation in all four treatments either converges to the REE or to a neighborhood around the REE although, again, such convergence happens at different speeds. The adjustment towards the REE is again observed to be fastest in $T1$ and slowest in $T4$.

Figure 6 plots the disaggregated average price expectations for each of the three-firm markets (independent observations) against the respective REE price for all four oligopoly treatments. Compared with the monopoly treatment, the convergence to REE in the oligopoly setting seems to be faster and more reliable in the eductively stable treatments, $T1 - T3$. By contrast, in the eductively unstable oligopoly market treatment $T4$, the volatility of price expectations appears to be greater and more persistent than in the oligopoly $T1-T3$ treatments or by comparison with the monopoly $T4$ treatment. Indeed, 6 of the 10 oligopoly markets in the $T4$ treatment failed to converge to the REE within the 50 periods allowed according to our convergence criterion as described in the next section. This finding suggests that the oligopoly market

setting may facilitate learning when the REE is eductively stable as this environment aggregates the already near-rational expectations of other agents and may thus speed up the achievement of common knowledge of rationality. However, when the REE is not eductively stable so that the REE is not rationalizable, we frequently observe non-convergence to the REE or greater volatility in market price expectations that do converge to the REE.

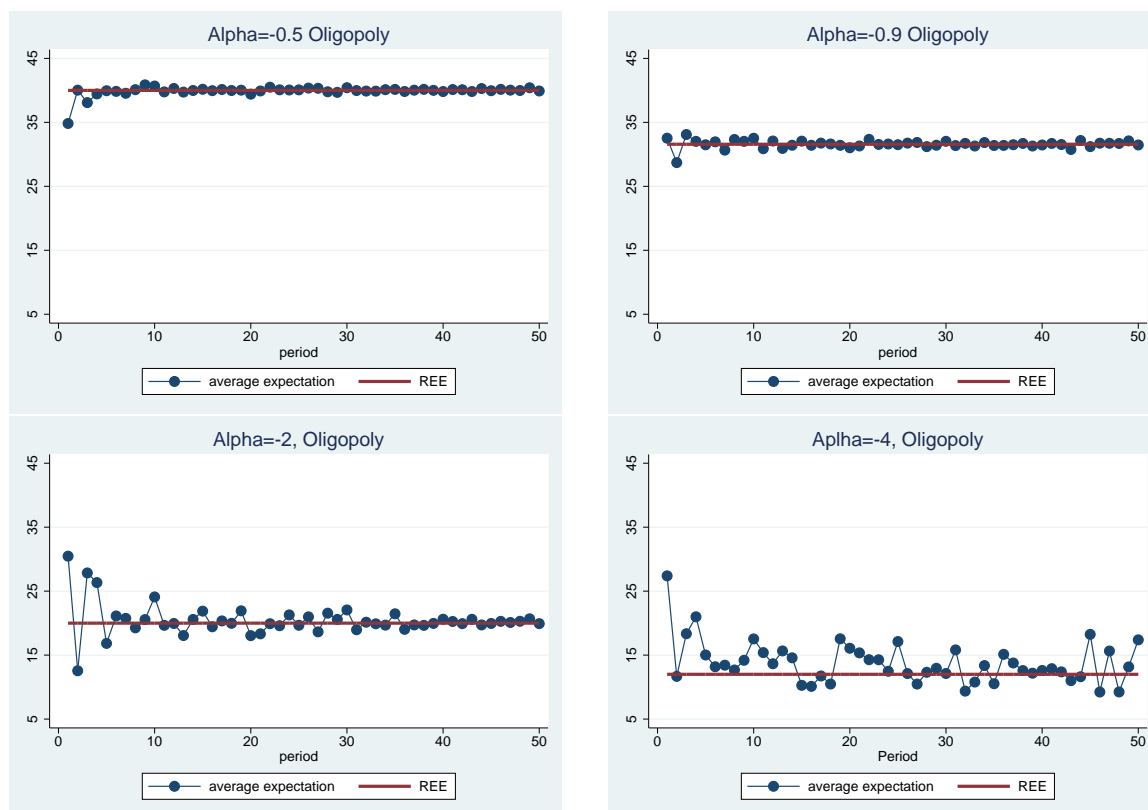


Figure 5: The average expectation against the REE price when $\alpha = -0.5, -0.9, -2$ and -4 (from top to bottom) in the oligopoly design.

Table 3 reports the mean and variance of price expectations across all markets of each of the four oligopoly treatments for the entire 50 period sample and for the first 25 and last 25 periods of the sample. Consistent with Figures 5-6, we observe that, on average, price expectations converged to the REE prediction for each treatment and that the variance in price expectations increases with α , with $T1$ having the lowest variance at 1.14 over all 50 periods and treatment $T4$ having the greatest at 77.91 over all 50 periods. These rankings do not change if attention is restricted to the last 25 periods. Note that under adaptive learning, the value of α should not matter for

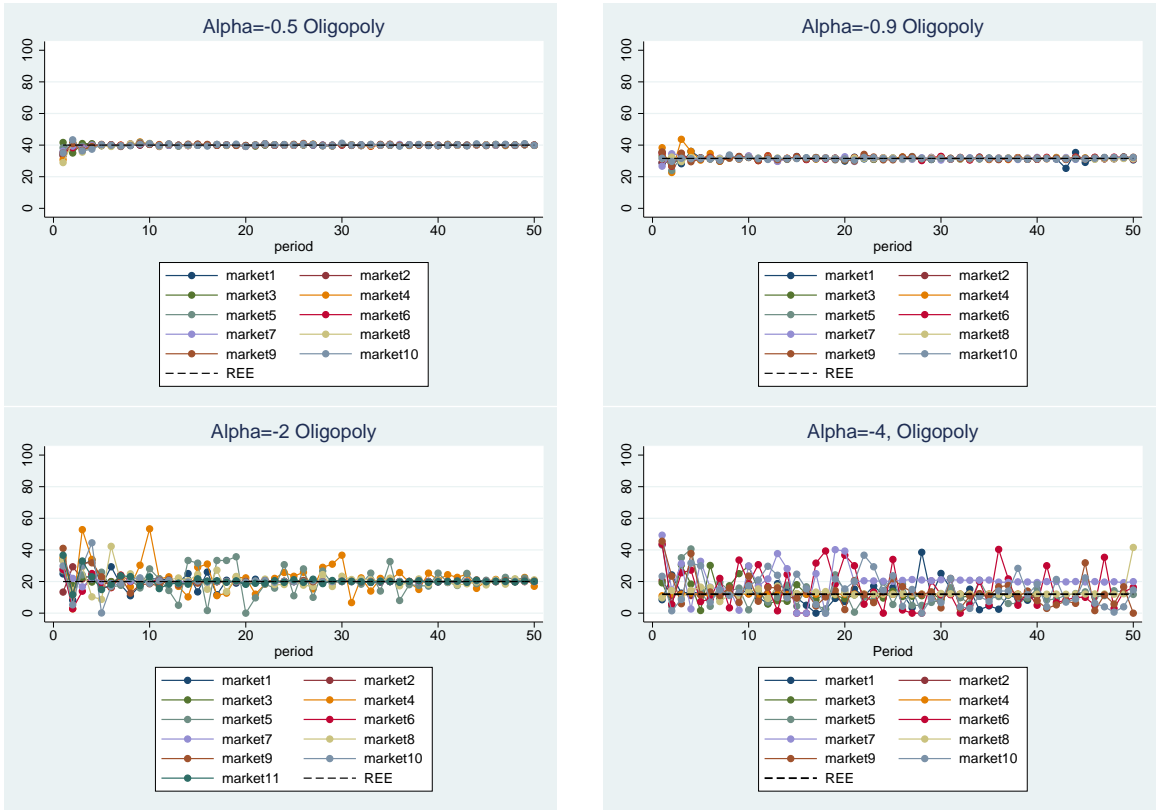


Figure 6: Disaggregated average expectations against the REE price when $\alpha = -0.5, -0.9, -2$ and -4 (from top to bottom) in the oligopoly design.

the variance of price expectations while under educative learning, REE is either stable or unstable.

Treatment	REE	Periods 1-50		Periods 1-25		Periods 26-50	
		Mean	Variance	Mean	Variance	Mean	Variance
$\alpha = -0.5$	$p^* = 40$	39.90	1.14	39.74	2.18	40.05	0.08
$\alpha = -0.9$	$p^* = 31.58$	31.60	1.61	31.61	2.82	31.58	0.45
$\alpha = -2.0$	$p^* = 20$	20.46	25.01	20.75	44.37	20.19	6.31
$\alpha = -4.0$	$p^* = 12$	13.98	77.91	15.29	97.66	12.67	55.00

Table 3: Mean price and variance of price expectations in each treatment ($\alpha = -0.5, -0.9, -2, -4$) of the oligopoly setting.

5.2 Convergence to REE

We declare that convergence to REE occurs in the first period for which the absolute difference between the (average) price expectation and the REE price prediction is less than 3 and stays below 3 forever after that period for the remainder of the 50 period market. We choose a threshold of 3 for two reasons: (1) ideally, one would like to declare convergence only if the price expectation was exactly equal to the REE, but if such a criteria were used, almost no market would satisfy that criterion so that it would not be possible to make any distinctions among our treatments; (2) the threshold should not be so large that it allows for substantial deviations from the REE. We choose the two sided range $[-3, +3]$ because it is 10% of the rationalizable price range, $[0, 60]$, and one-sided deviations from REE larger than 3 (5%) of this range may be regarded as substantial. We further categorize markets according to whether convergence happens immediately in period 1, or between periods 2 and 5, between periods 6 and 10, between periods 11 and 20, between periods 21 and 50 or the market is *non-convergent* according to our criterion as of the final period 50 (> 50). In calculating the average number of rounds to convergence, we use the first round in which convergence occurred according to our criterion, or if convergence did not occur within the 50 periods of the experiment, we declared the period of convergence to be period 51. Our tests for treatment differences using this convergence criterion (as discussed below) do not depend on the latter assumption because all of our tests are non-parametric so that only rank matters.

The results from applying our convergence criterion to each market of each treatment are reported in Table 4. In the final rows of this same table we also report the mean number of periods to convergence over all markets in each treatment as well as the variance. Table 4 reveals that it takes fewer periods, on average, for price forecasts to converge to the REE in treatment $T1$ as compared with treatments $T2$ - $T4$ for both the monopoly and oligopoly settings. There is less of a difference in the mean time to convergence between the $T2$ and $T3$ treatments of the monopoly setting, though convergence is slightly faster on average in $T2$ than in $T3$. Surprisingly, the number of periods before convergence obtains in treatment $T4$ of the monopoly design is smaller than in $T2$ and $T3$, though it remains larger than for $T1$.

As for the oligopoly setting, it is clear that as α becomes more negative in moving from treatments $T1$ - $T4$, the mean time to convergence steadily increases. A Wilcoxon Mann-Whitney test on market-level data (independent observations) suggests that the differences in the mean time to convergence between treatments $T1$ and either the $T2$ or $T3$ treatments is statistically significant at the 5% level for both the monopoly and oligopoly treatments while the differences in the mean time to convergence between treatments $T2$ and $T3$ is statistically significant at the 5% level for the oligopoly treatment but not statistically significant for the monopoly treatment. The mean time to convergence in treatment $T4$ is not significantly different at the 5% level from that of all other treatments in the monopoly design, and is significantly different from all other treatments of the oligopoly design. The mean time to convergence is smaller in the oligopoly design than in the monopoly design for the eductively stable treatments ($T1$ and $T2$), where the difference is significant at the 5% for $T1$ but not for $T2$, while it is larger in the oligopoly design than in the monopoly design in treatments $T3$ and $T4$, where the differences are significant at the 5% level for both treatments.

For both the monopoly and oligopoly markets, the variance in the number of periods before convergence is smallest in treatment $T1$. In the monopoly market treatment, the variance in the number of periods required for convergence is larger in treatment $T2$ than in treatment $T3$, but this is due to the random behavior of just a few subjects in $T2$ who inexplicably began experimenting with high/low price predictions after they had converged to the REE for more than 10 periods. For the oligopoly treatment, the variance in the number of periods required for convergence is also smallest in treatment $T1$ and is higher and similar in treatments $T2$ and

$T3$.¹⁷ These results generally support the notion that convergence is more difficult as the absolute value of the coefficient α becomes larger, as larger values of α make the market more unstable. We have verified, using simulations of adaptive learning agents (available on request) that the number of periods required for convergence to the REE is indeed increasing with increases in the value of $|\alpha|$.

Table 4 also reveals that for all four monopoly market treatments, there is at least 1 market (and often more) that converges to the REE beginning with the very first period. As noted earlier, convergence to the REE in the very first period may be regarded as support for the eductive learning approach. If this eductive learning criterion is relaxed to allow for convergence within the first 5 periods of the experiment, then more than 60% of all markets in all four monopoly treatments can be said to be consistent with eductive learning. A similar finding obtains for the eductively stable treatments, $T1$ - $T2$, of the oligopoly design. By contrast, in treatments $T3$ and $T4$ of the oligopoly treatment, there are *no* instances of convergence to the REE in the very first period of a session and in the $T4$ oligopoly treatment, 60% of markets (6/10) failed to satisfy our convergence criterion within the 50 periods allowed by our experiment. These differences in outcomes between the eductively stable treatments $T1$ and $T2$ and the eductively unstable treatments $T3$ and $T4$, suggest that the eductive stability criterion is indeed useful in understanding differences in the behavior of subjects in our experiment. Furthermore, the significant frequency of non-convergent outcomes in the $T4$ oligopoly treatment compared with the $T3$ oligopoly treatment (where all markets converged) provides support for Gaballo's (2013) general eductive stability requirement based $N = 3$ firms.

Figure 7 displays cumulative distribution functions (CDFs) of the percentage of markets in each treatment that have met our convergence criterion by each of the 50 periods of our experiment. For the monopoly treatment (left panel of Figure 7), these CDFs are closely aligned and intersect one another suggesting that there is not much difference in the distribution of convergence times across these four monopoly treatments. Indeed, a non-parametric Kolmogorov-Smirnov test indicates that there is no significant difference in these distributions at the 5% level for all pairwise comparisons of the four monopoly treatments. By stark contrast, in the oligopoly markets (right panel of Figure 7) it is clear that markets in treatment $T1$ converge the fastest

¹⁷We performed a Siegel-Tukey test and found that none of the differences in variances across treatments is significant at 10% level.

Convergence in period(s)	monopoly				oligopoly			
	$\alpha = -0.5$	$\alpha = -0.9$	$\alpha = -2$	$\alpha = -4$	$\alpha = -0.5$	$\alpha = -0.9$	$\alpha = -2$	$\alpha = -4$
1	21.4% (3)	41.7% (5)	30.8% (4)	7.1% (1)	20.0% (2)	30.0%(3)	0.0% (0)	0.0% (0)
[2, 5]	57.1% (8)	25.0% (3)	30.8% (4)	78.6% (11)	80.0% (8)	60.0%(6)	27.3% (3)	0.0% (0)
[6, 10]	14.3% (2)	16.7% (2)	7.7% (1)	7.1% (1)	0.0% (0)	0.0%(0)	18.2% (2)	0.0% (0)
[11, 20]	0.0% (0)	0.0% (0)	0.0% (0)	7.1% (1)	0.0% (0)	0.0%(0)	18.2% (2)	0.0% (0)
[21, 50]	7.1% (1)	16.7% (2)	30.8% (4)	0.0% (0)	0.0% (0)	10.0%(1)	36.4% (4)	40.0% (4)
> 50	0.0% (0)	0.0% (0)	0.0% (0)	0.0% (0)	0.0% (0)	0.0%(0)	0.0% (0)	60.0% (6)
Average	4.4	8.2	8.8	5.5	2.1	7.1	17.3	40.9
Variance	29.5	179.8	87.0	47.3	1.1	188.8	289.2	171.7
# Obs	14	12	13	14	10	10	11	10

Table 4: Frequency distribution of the number of periods required for convergence to the REE in each treatment. Convergence period ranges are given in the left-most column as bins []. The numbers of observations per bin are indicated in ().

followed by treatment T2, then treatment T3 and lastly by treatment T4 where only 40% of markets had satisfied our convergence criterion by the final period 50. A Kolmogorov-Smirnov test indicates there is no significant difference between the distribution of periods before convergence in Treatments T1 and T2. However, this same test indicates that both the oligopoly treatments T3 and T4 are significantly different from the other treatments (namely, T3 differs from T1, T2, T4, and T4 differs from T1-T3) at the 5% level.




Figure 7: The fraction of markets that have converged by each period under the different treatments. The vertical axis measures the percentage of markets that converge while the horizontal axis reports the period number.

We summarize the findings in the above sections as Results 1-2 addressing Hy-

potheses 1-2:

Result 1. *We do not reject Hypothesis 1 for the monopoly markets as convergence to the REE obtains in all four treatments of the monopoly design. We do reject Hypothesis 1 for the oligopoly treatment, in particular for treatment T4 where 60% of markets (observations) failed to converge to the REE within 50 periods.*

Result 2. *We reject Hypothesis 2, that when the REE is eductively stable convergence to the REE occurs immediately as it is more frequently the case that convergence requires more than a single period.*

Results 1-2 suggest, as Evans (2001) posited, that the learning process may be regarded as a mixture of both adaptive and eductive approaches when the REE is not learnable under eductive learning.

5.3 Fit of the Two Learning Approaches to the Experimental Data

We next consider the fit of the two different learning approaches to our experimental data. For the adaptive learning model, we assume that the model's predictions coincide with the actual (average) price prediction in the experimental data. To initialize a simulation of the adaptive learning model we set the initial price prediction, p_1^e , equal to the individual (monopoly) or average (oligopoly) forecasts made by subjects in period 1 (and period 1 only). Thereafter, the adaptive learning model specifies how all subsequent simulated prices and price predictions are determined. That is, given p_1^e , the price for period 1, p_1 , is determined by equation (1). Given p_1 the adaptive learning model predicts the price for period 2 according to equation (3) and thus generates a simulated actual price for period 2 again via equation (1). In period 3, the adaptive model take the average of the simulated prices for periods 1 and 2 and makes a price prediction for period 3, which is then used to generate the simulated price for period 3 via equation (1), and so on. Thus, the adaptive learning model uses its own simulated prices as input to generate simulated market price predictions in each period. Importantly, this simulation only loads experimental data from period 1, and makes simulated prices and predictions for the remaining 49 periods, so there are no degrees of freedom in the predictions of the adaptive learning model for each market observation.

For the educative learning model, we assume that for each period, the simulated price prediction is $p_{edc,t}^e = \frac{\mu}{1-\alpha}$ and thus the actual market price equals the REE price, $\frac{\mu}{1-\alpha} + \nu_t$ in the treatments where educative learning predicts that the REE is learnable, namely, all treatments of the monopoly setting, and treatments $T1 - T3$ of the oligopoly setting. Note that for treatment $T4$ in the monopoly setting, the educative learning model predicts that the market price equals the REE price because educative learning assumes that individuals are perfectly capable of solving for the REE. Since in the monopoly case, there are no other firms (agents) forming expectations, there is no strategic uncertainty regarding the decision of others and so it follows that the market price should equal the REE price even in treatment $T4$ of the monopoly setting. By contrast, in treatment $T4$ under the oligopoly setting, the educative learning model does not exclude any combination of price predictions and market prices, strictly speaking for treatment $T4$ only, given our finite number of firms. Therefore, the MSE from educative learning for the $T4$ treatment of the oligopoly setting are essentially undefined, or 0 if we consider that the model predicts that “anything can happen.” For these reasons we did not construct MSEs between the data and the educative learning model prediction for the $T4$ treatment of the oligopoly setting. However, for the other seven treatments, we can calculate the MSE between the data and the educative learning model predictions. We note that, as was the case for the adaptive learning model predictions, there are again no degrees of freedom in the predictions of the educative learning model.

The mean squared error (MSE) between the simulated and experimental data for each market (independent observation) are presented in Table 5. The results suggest that the adaptive learning model is generally a better fit to the experimental data, as it results in a smaller MSE relative to the experimental data than does the educative learning model using that same experimental data; the average MSE for the adaptive learning model is lower than for educative learning model in all of the treatments where this MSE can be calculated. A Wilcoxon signed rank test suggests that the difference between the MSEs for the adaptive and educative learning models is significant (in favor of adaptive learning) at the 5% level for the $\alpha = -0.5$ and $\alpha = -0.9$ treatments of both the monopoly and oligopoly settings, and for the $\alpha = -4$ treatment of the monopoly setting. However, there is also some heterogeneity across the different markets/observations. For example, in the monopoly market with $\alpha = -2$, the adaptive learning model generates a higher MSE relative to the educative learning

model in markets 2, 6, 8, 9 and 13 but a lower MSE relative to the educative learning model in markets 1, 3, 4, 5, 7, 10, 11 and 12. This finding suggests that it is very likely that some markets are dominated by subjects using adaptive learning, while others are dominated by subjects using educative learning. We will provide evidence for such heterogeneity of types later on in section 5.6.

5.4 Payoff Efficiency

Table 5.4 reports average payoffs and payoff efficiency (payoffs divided by 25 euros, which was the maximum amount that each subject could earn when they made no forecasting errors) for each treatment. Payoff efficiency is more than 90% when $\alpha = -0.5$ or $\alpha = -0.9$, and lower, between 50%-90% when $\alpha = -2$ and $\alpha = -4$. Efficiency is higher in the $T1$ oligopoly treatment as compared with the $T1$ monopoly treatment and lower in the $T3$ - $T4$ oligopoly treatments than in the $T3$ - $T4$ monopoly treatments. We performed a Wilcoxon Mann-Whitney Test using individual earnings data in the monopoly design and average earnings for each market of the oligopoly design. The results indicate that for the monopoly treatment, there is no difference in payoff efficiency between the $T1$, $T2$ and $T4$ treatments at the 5% level, but that payoff efficiency in these treatments is significantly greater than payoff efficiency in the $T3$ treatment at the 5% level. In the oligopoly treatment, payoff efficiency is monotonically decreasing in α , and the difference between each pair of treatments is statistically significant at the 5% level. The considerably lower payoff efficiency found in treatment $T4$ of the oligopoly design suggests that the forecasting task is more cognitively demanding when the REE is unstable under educative learning. We also compare, for the same treatment (same α), whether there is a significant earnings difference between the monopoly and oligopoly settings. It turns out that in four pairwise comparisons, the difference is only significant at 5% level for the $T4$ treatment ($\alpha = -4$), where earnings in the oligopoly setting are substantially lower than in the monopoly setting.

5.5 Decision Time

We collected data on the time it took subjects to make their decisions. Specifically, we measured the time, in seconds, from the start of each new period to the time at which

Learning		Monopoly				Oligopoly			
Model	Market	$\alpha = -0.5$	$\alpha = -0.9$	$\alpha = -2$	$\alpha = -4$	$\alpha = -0.5$	$\alpha = -0.9$	$\alpha = -2$	$\alpha = -4$
Adaptive	1	0.94	0.15	377.75	0.40	0.10	2.21	11.88	18.14
	2	0.33	0.38	78.71	9.60	0.18	0.50	1.26	7.15
	3	0.17	23.46	42.25	0.31	0.35	0.30	6.98	3.71
	4	109.73	0.06	0.93	9.94	0.26	4.63	69.94	182.94
	5	0.97	5.15	10.57	98.77	0.12	0.26	62.46	37.64
	6	0.16	5.40	0.13	0.45	0.27	0.73	5.40	147.64
	7	23.76	8.73	8.03	0.49	0.63	0.57	12.87	129.85
	8	0.08	0.70	223.79	0.98	0.57	0.33	25.50	20.32
	9	0.79	0.17	31.01	6.27	0.26	0.97	10.19	65.50
	10	0.34	3.94	0.67	6.00	0.63	0.28	21.88	88.17
	11	0.49	9.67	371.13	1.72			5.51	
	12	1.49	1.43	1.19	127.85				
	13	0.03		2.19	13.24				
	14	9.46			0.66				
	average	10.62	4.94	88.34	19.76	0.34	1.08	21.29	70.11
Eductive	1	2.52	0.00	384.14	0.27	0.89	2.41	16.70	-
	2	0.91	0.09	76.26	18.01	1.04	0.63	3.12	-
	3	2.17	52.14	81.08	1.72	0.68	0.42	3.88	-
	4	141.14	0.06	1.02	8.82	2.02	6.34	91.21	-
	5	3.29	30.14	13.41	121.62	0.19	1.62	63.38	-
	6	2.02	5.75	0.00	0.04	0.67	1.30	9.29	-
	7	22.05	23.43	45.04	0.28	0.46	1.14	3.96	-
	8	0.19	0.69	230.52	0.21	3.12	0.19	26.19	-
	9	3.02	0.21	10.81	0.64	0.90	1.60	18.03	-
	10	0.15	3.79	0.71	39.17	1.34	0.33	28.67	-
	11	0.49	9.76	387.83	15.49			12.75	
	12	3.62	1.33	4.00	137.22				
	13	5.01		2.05	50.63				
	14	13.01			8.59				
	average	14.26	10.62	95.14	28.77	1.13	1.60	25.20	-

Table 5: MSE between the experimental data and the two learning model predictions when these models predict convergence to the REE.

Market Structure	α	Avg. Payoff	Efficiency
Monopoly	-0.5	22.9	91.6%
	-0.9	22.7	90.8%
	-2	20.1	80.4%
	-4	21.8	87.0%
Oligopoly	-0.5	23.7	94.6%
	-0.9	22.6	90.6%
	-2	16.7	65.6%
	-4	12.6	50.2%

Table 6: Payoffs and payoff efficiency across the eight treatments.

each subject clicked “send” to submit their price forecast for that same period. Such data can be useful in understanding possible variations in the cognitive difficulty of decision-making tasks. In particular, Rubinstein (2007) provides evidence that choices requiring greater cognitive activity are positively correlated with a longer decision response time. In our experiment, subjects face a more difficult task in treatment $T4$ as compared with treatments $T1 - T3$ and so they may be expected to take more time to make their decisions in treatment $T4$ than in treatments $T1 - T3$. At the suggestion of a referee, we compare the average decision time in the first period only, since our computer program only advances to the next period when all subjects have made a forecast, the decision time in later periods may be influenced by subjects’ experience with waiting times in earlier periods, and thus these observations may not be independent beyond the first period. Table 7 provides descriptive statistics of the decision time in the first period in each treatment. For the monopoly treatment, the average decision time is 53.1 seconds in $T1$, 49.6 seconds in $T2$, 103.8 seconds in $T3$ and 148.5 seconds in $T4$. In the oligopoly treatment, the average decision time is 55.2 seconds in $T1$, 62.3 seconds in $T2$, 98.1 seconds in $T3$ and 89.2 seconds in $T4$.

Figure 8 shows the empirical cumulative distribution function (CDF) of decision time in the first period for treatments $T1 - T4$ of the monopoly setting. We find that the difference between each of $T1$ and each of $T3 - T4$, and $T2$ and $T4$ is significant at the 5% level according to a Wilcoxon Mann-Whitney test, while other differences across treatments are not statistically significant.

In the oligopoly treatment, the difference between each of $T1 - T2$ and each of

Monopoly				
Treatment	T1	T2	T3	T4
Average	53.1	49.6	103.8	148.5
Variance	2956.8	393.7	10569.3	11237.2
Number of Obs.	14	12	13	14
Oligopoly				
Treatment	T1	T2	T3	T4
Average	55.2	62.3	98.1	89.2
Variance	1246.6	1562.9	14595.2	5507.8
Number of Obs.	30	30	33	30

Table 7: The decision time in the first period in each treatment.

$T3 - T4$ is significant at the 5% level according to a Wilcoxon Mann-Whitney test, while the differences between $T1$ and $T2$ or $T3$ and $T4$ are not significantly different from one another. These findings support the notion that subjects face a more difficult task in $T4$, and therefore require more time to make a decision.

The findings in this section are summarized by Result 3 which addresses Hypothesis 3:

Result 3. *We reject Hypothesis 3. When common knowledge of rationality is an issue as in our oligopoly setting, decision time is always significantly greater in treatment $T4$ relative to the other two treatments.*

We note additionally that the cognitive cost of decision-making in treatment $T4$ of our monopoly setting is often significantly larger than in $T1 - T2$.

5.6 Categorization of Subjects into Adaptive or Eductive Learners

In this section we focus on individual subject price forecasts and we attempt to categorize each subject in our experiment as one of three types: adaptive learner, eductive learner or neither. We do this using two different approaches and we examine the consistency between these two approaches.

The first approach is to make categorizations based on the definition of the two types of learning. This categorization is performed as follows:




Figure 8: The empirical CDF of the time taken to complete decision tasks in treatments $T1 - T4$ of the monopoly (solid line or markers) and oligopoly (long dash line or hollow markers) settings. The unit of time is seconds, as measured on the horizontal axis. The vertical axis measures the cumulative frequency.

1. Eductive learners: We consider all subjects who predict the REE in the very first period to be eductive learners. Since the REE in treatment $T2$ where $\alpha = -0.9$ is 31.58, and not an integer, taking into account that some subjects may use $\alpha = -1$ as a proxy, we consider all subjects making predictions in the range $[30, 32]$ to be eductive learners in $T2$. For the other three α treatment values, the REE is an integer value so to be categorized as an eductive learner, subjects must correctly predict a price of 40 in $T1$, a price of 20 in $T3$ and a price of 12 in $T4$.
2. Adaptive learners: For each subject we use their first period price forecast to initialize the adaptive learning model as given in equation (3) and we then calculate the mean squared error between the simulated predictions of that model and each individual subject's actual price predictions. If the mean squared error between the actual and predicted price forecasts is smaller than 1, then the subject is classified as an adaptive learner. We choose a threshold of 1 as we wanted the threshold to be as low as possible, but at the same time to allow for subjects to engage in some rounding of numbers to integer values. Since adaptive learning does not make assumptions on the initial price prediction, the probability that one happens to come up with the REE is infinitely close to 0 under adaptive learning.

If a subject meets our criteria for being categorized as *both* an adaptive and an eductive learner, then we classify him/her as an eductive learner. If a subject meets neither criteria, then he/she is classified as “neither”.

Our second approach to type classification makes use of answers that subjects gave to a post-experimental questionnaire (see the Appendix for details). The questionnaire asked subjects to characterize the type of price prediction strategy they used during the experiment. We provided them with four options and we asked them to choose the option that best described how they made their predictions in the experiment. Specifically, the four options were:

1. I refer to information about past prices.
2. I make calculations based on the value of α .
3. I eliminate unlikely numbers iteratively.

4. None of the above.

A subject is classified as an adaptive learner if he chooses option 1, and is classified as an educative learner if he chooses option 3. If the subject chooses option 2, it is likely that he solves the REE directly, and we also classify this type as a an educative learner, since, as discussed in section 3.2, the educative learning model typically starts with the assumption that agents solve the REE directly from equation (1).¹⁸ Subjects choosing option 4 are classified as “neither”. Due to a technical problem, we lost some data on self-reported strategies in the first, and relatively larger session of our monopoly market treatments, (9 markets for treatment 1, 8 markets for each of treatments 2 and 3). Nevertheless, we do have data on self-reported strategies for many of our subjects and for all eight treatments.

Table 5.6 shows the number of participants who can be categorized as adaptive or educative learners in each of our eight treatments using Approach 1 or Approach 2, as well as the overall frequency of each type classification for each treatment. Tables 10 and 11 in the Appendix report more disaggregated information on each individual subject’s type using both approaches (where possible). We observe that using approach 1, classification becomes more difficult as α becomes more negative, as indicated by the frequency of the ‘neither’ category. Among those subjects who can be classified using approach 1, overall about 55% can be classified as either adaptive or educative learners, with a roughly equal split between the two types. Using approach 2 subjects are more likely to be classified as educative learners than adaptive learners in the monopoly treatment, but the reverse generally holds for the oligopoly treatment where overall, a majority of subjects (55%) can be classified as adaptive learners. An important exception is the $\alpha = -4$ oligopoly treatment where a majority of subjects are classified as educative learners. We note further that there is a good level of consistency between the categorizations based on our two different approaches. For 65 subjects for which both approaches yield a classification of either adaptive or educative learners, the two approaches agree on the type assignment in 43 cases, which means that the two approaches assign the same category with a probability of $43/65 = 66.2\%$.

In addition to classifying subjects as adaptive or educative learners, we further

¹⁸Here we do not consider the case where subjects use a mixture of observation and calculation as in Evans and Ramey (1992), though that would be an interesting extension

Approach 1										
Treatment	$\alpha = -0.5$		$\alpha = -0.9$		$\alpha = -2$		$\alpha = -4$		All	
Monopoly										
Adaptive	8	57.14%	2	16.67%	2	15.38%	5	35.71%	17	32.08%
Eductive	3	21.43%	5	41.67%	3	23.08%	1	7.14%	12	22.64%
Neither	3	21.43%	5	41.67%	8	61.54%	8	57.14%	24	45.28%
Total	14	100.00%	12	100.00%	13	100.00%	14	100.00%	53	100.00%
Oligopoly										
Adaptive	17	56.67%	8	26.67%	9	27.27%	0	0.00%	34	27.64%
Eductive	7	23.33%	10	33.33%	8	24.24%	10	33.33%	35	28.46%
Neither	6	20.00%	12	40.00%	16	48.48%	20	66.67%	54	43.90%
Total	30	100.00%	30	100.00%	33	100.00%	30	100.00%	123	100.00%
Approach 2										
Treatment	$\alpha = -0.5$		$\alpha = -0.9$		$\alpha = -2$		$\alpha = -4$		All	
Monopoly										
Adaptive	1	7.14%	1	8.33%	2	15.38%	2	14.29%	6	11.32%
Eductive	3	21.43%	2	16.67%	3	23.08%	10	71.43%	18	33.96%
Neither	10	71.43%	9	75.00%	8	61.54%	2	14.29%	29	54.72%
Total	14	100.00%	12	100.00%	13	100.00%	14	100.00%	53	100.00%
Oligopoly										
Adaptive	20	66.67%	23	76.67%	15	45.45%	10	33.33%	68	55.28%
Eductive	10	33.33%	6	20.00%	10	30.30%	16	53.33%	42	34.15%
Neither	0	0.00%	1	3.33%	8	24.24%	4	13.33%	13	10.57%
Total	30	100.00%	30	100.00%	33	100.00%	30	100.00%	123	100.00%

Table 8: Number and percentage of subjects who can be categorized as adaptive or eductive learners or neither in each treatment. Approach 1 is the approach based on first period price predictions and the mean squared error of individual price predictions from the adaptive learning model. Approach 2 is the approach based on self-reported strategies.

considered whether any of our subjects were following naïve expectations, i.e., forecasting $p_t^e = p_{t-1}$. We do not consider such expectations to constitute a real learning model as the forecast rule does not update over time. Nevertheless we investigate the existence of naïve types because, as noted earlier in section 3.2, stability under naïve expectations is a potentially confounding criterion (with eductive stability) in our model set-up. We classify a naïve forecaster as follows: if the MSE between the naïve expectations prediction and a subject’s actual price forecast is less than 1 and this MSE is also less than the MSE of adaptive learning model as well, then that subject can be regarded as having naïve expectations. Using this criterion, only 2 subjects (3.77%) in the monopoly treatment and 7 subjects (5.69%) in the oligopoly treatment can be categorized as users of naïve expectations. These numbers are far smaller than the number of subjects classified as adaptive or eductive learners. We thus conclude that while there are indeed some followers of naïve expectations in our experiment, most of our subjects are forming expectations in a more sophisticated manner.

6 Conclusion

The process by which agents might learn a REE has been the subject of a large amount of theorizing but surprisingly there has been little empirical assessment of the leading theories of this learning process. To address this gap, we have designed and implemented a learning-to-forecast experiment in the context of a simple cobweb economy with negative feedback where expectations matter and where subjects are informed about the law of motion for prices. We are particularly interested in knowing which of two leading approaches to modeling learning – adaptive learning or eductive learning – provides the better characterization of human learning behavior in this setting. In particular, we vary the slope parameter of the price determination equation, α , in such a way that in some of our treatments the REE may not be learnable (stable under learning) if agents are eductive learners but should always be learnable if agents are adaptive learners. We further investigate different predictions between the two learning theories with regard to the speed of convergence. Finally, our experimental design includes both monopoly and oligopoly settings in order to better understand the role played by common knowledge of rationality.

In most (7/8) of our treatments, we always observe convergence of prices to the REE within the 50 period time frame of our experiment. However, the variance in market prices is much greater as the α parameter becomes more negative. In the oligopoly treatment where $\alpha = -4$ and the REE is unambiguously eductively unstable, even adjusting for the finite number of firms, we observe that 60 percent of our markets fail to achieve convergence to the REE within the 50 periods. The latter finding is supportive of eductive learning as a characterization of subject behavior. Further, there are many instances of markets that satisfy our criteria for convergence to the REE in the very first period, which is more in line with eductive rather than adaptive learning. On the other hand, the observation that most markets take some time to converge to the REE, and convergence is observed for at least some markets in all eight of our treatments, including the oligopoly treatment where $\alpha = -4$ and the REE is unambiguously eductively unstable, favors adaptive learning as a characterization of subject behavior. Indeed, our efforts to classify subjects as adaptive or eductive learners reveal a mix of both learning types in all treatments (as well as many subjects who are unclassifiable). Perhaps, as Evans (2001) suggests, individuals or populations of individuals use a mixture of both adaptive and eductive learning approaches.

The cobweb economy that we study is a very simple economic model involving negative feedback. Our experimental examination of forecasting behavior in this model provides subjects with complete information about the data generating process. In this sense, our experiment can be viewed as providing very favorable conditions for the rational expectation hypothesis and for the eductive learning approach in particular. Our findings confirm that the rational expectation hypothesis and rational expectation equilibrium provide a reasonable characterization of the market outcome in this setting after a period of learning, provided that the REE is both adaptively and eductively stable. Further comparisons between adaptive and eductive learning approaches should be conducted in environments where subjects face a more complicated, forward-looking dynamic economic model where forecasts matter for realizations of future state variables, as in dynamic, stochastic general equilibrium models. Another extension would be to consider positive feedback systems, as opposed to the negative feedback system of the Cobweb model. Previous research, e.g., Hommes et al. (2005, 2008) shows that positive feedback systems tend to generate oscillatory bubbles and crashes when the slope parameter in an asset pricing model is positive

and close to 1. Future research might consider the extent to which adaptive and eductive learning approaches predict *individual* learning behavior when $0 < \alpha < 1$ and both learning theories predict convergence to REE. We leave these extensions to future research.

References

- Akiyama, E., N. Hanaki, and R. Ishikawa. "Effect of Uncertainty about Others' Rationality in Experimental Asset Markets: An Experimental Analysis." Aix-Marseille School of Economics Discussion Paper, (2012), DP2012-34.
- Akiyama, E., N. Hanaki, and R. Ishikawa. "It Is Not Just Confusion! Strategic Uncertainty in an Experimental Asset Market." Aix-Marseille School of Economics Discussion Paper, (2013), DP2013-40.
- Bao, T., J. Duffy, and C.H. Hommes. "Learning, Forecasting and Optimizing: An Experimental Study." *European Economic Review* 61, (2013), 186-204.
- Bao, T., C.H. Hommes, J.H. Sonnemans, and J. Tuinstra. "Individual Expectations, Limited Rationality and Aggregate Outcomes." *Journal of Economic Dynamics and Control* 36, (2012), 1101-1120.
- Bernheim, D. "Rationalizable Strategic Behavior." *Econometrica* 52, (1984), 1007-1028.
- Brock, W.A., and C.H. Hommes. "A Rational Route to Randomness." *Econometrica* 65, (1997), 1059-95.
- Beshears, J., J. Choi, A. Fuster, D. Laibson, and B.C. Madrian. "What Goes Up Must Come Down? Experimental Evidence on Intuitive Forecasting." *American Economic Review* 103(3), (2013), 570-74.
- Bosch-Doménech, A., and N.J. Vriend. "Imitation of Successful Behaviour in Cournot Markets." *Economic Journal* 113, (2003), 495-524.
- Bray, M.M. "Learning, Estimation and the Stability of Rational Expectations Equilibria." *Journal of Economic Theory* 26, (1982), 318-339.
- Bray, M.M., and N.E. Savin. "Rational Expectations Equilibria, Learning, and Model Specification." *Econometrica* 54(5), (1986), 1129-1160.
- Duffy, J. "Macroeconomics: A Survey of Laboratory Evidence." to appear in J.H. Kagel and A.E. Roth (eds.), *Handbook of Experimental Economics Volume 2*, Princeton: Princeton University Press (2016).

- Duffy, J., and R. Nagel. "On the Robustness of Behaviour in Experimental Beauty Contest Games." *Economic Journal* 107, (1997), 1684-1700.
- Evans, G.W. "Expectations in Macroeconomics. Adaptive versus Eductive Learning." *Revue Économique* 52(3), (2001), 573-582.
- Evans, G.W., and S. Honkapohja. *Learning and Expectations in Macroeconomics*. Princeton: Princeton University Press, (2001).
- Evans, G.W., and S. Honkapohja. "Convergence for Difference Equations with Vanishing Time-dependence, with Applications to Adaptive Learning." *Economic Theory* 15(3), (2000), 717-725.
- Evans, G.W., and G. Ramey. "Expectation Calculation and Macroeconomic Dynamics." *American Economic Review* 82(1), (1992), 207-224.
- Fehr, E., and J.R. Tyran. "Individual Irrationality and Aggregate Outcomes." *Journal of Economic Perspectives* 19(4), (2005), 43-66.
- Fehr, E., and J.R. Tyran. "Money Illusion and Coordination Failure." *Games and Economic Behavior* 58(2), (2007), 246-268.
- Fehr, E., and J.R. Tyran. "Limited Rationality and Strategic Interaction: the Impact of the Strategic Environment on Nominal Inertia." *Econometrica* 76(2), (2008), 353-394.
- Gaballo, G., "Eductive Learning and the Rationalizability of Oligopoly Games." *Economics Letters* 120(3), (2013), 401-404.
- Grandmont, J.M. "Expectations Formation and Stability of Large Socioeconomic Systems." *Econometrica* 66(4), (1998), 741-781.
- Grosskopf, B., and R. Nagel. "The Two-Person Beauty Contest." *Games and Economic Behavior* 62(1), (2008), 93-99.
- Guesnerie, R. "An Exploration of the Eductive Justifications of the Rational-Expectations Hypothesis." *American Economic Review* 82(5), (1992), 1254-1278.
- Guesnerie, R. "Anchoring Economic Predictions in Common Knowledge." *Econometrica* 70(2), (2002), 439-480.

- Heemeijer, P., C.H. 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(5), (2009), 1052-1072.
- Ho, T.H., Camerer, C., and K. Weigelt. “Iterated Dominance and Iterated Best Response in Experimental “p-beauty Contests”. *American Economic Review* (1998): 947-969.
- Hommes, C.H., J.H. Sonnemans, J. Tuinstra, and H. Van de Velden. “Coordination of Expectations in Asset Pricing Experiments.” *Review of Financial studies* (2005), 18(3), 955-980.
- Hommes, C.H., J. Sonnemans, and H. van de Velden. “Expectation Formation in a Cobweb Economy: some one Person Experiments.” In: Delli Gatti, D., Gallegati, M., and Kirman, A.P. (Eds.) *Interaction and Market Structure*, Springer Verlag Berlin (2000), 253-266.
- Hommes, C.H., J.H. Sonnemans, J. Tuinstra, and H. Van de Velden. “Expectations and bubbles in asset pricing experiments.” *Journal of Economic Behavior and Organization* (2008): 116-133.
- Hommes, C.H. “The Heterogeneous Expectations Hypothesis: Some Evidence from the Lab.” *Journal of Economic Dynamics and Control* 35(1), (2011), 1-24.
- Hommes, C.H., J. Sonnemans, J. Tuinstra, and H. van de Velden. “Learning in Cobweb Experiments.” *Macroeconomic Dynamics* 11 (Supplement 1), (2007), 8-33.
- Hommes, C.H., and F. Wagener. “Does Eductive Stability Imply Evolutionary Stability?” *Journal of Economic Behavior and Organization* 75(1), (2010), 25-39.
- Huck, S., H.T. Normann, and J. Oechssler. “Learning in Cournot Oligopoly—An Experiment.” *Economic Journal* 109, (1999), C80-C95.
- Marimon, R., and S. Sunder. “Indeterminacy of Equilibria in a Hyperinflationary World: Experimental Evidence.” *Econometrica* 61(5), (1993), 1073-1107.
- Muth, J.F., “Rational Expectations and the Theory of Price Movements.” *Econometrica* 29(3), (1961), 315-335.

- Nagel, R. "Unraveling in Guessing Games: An Experimental Study." *American Economic Review* 85(5), (1995), 1313-1326.
- Offerman, T., J. Potters, and J.H. Sonnemans. "Imitation and Belief Learning in an Oligopoly Experiment." *Review of Economic Studies* 69(4), (2002), 973-997.
- Pearce, D. "Rationalizable Strategic Behavior and the Problem of Perfection." *Econometrica* 52, 1029-1050.
- Rubinstein, A. "Instinctive and Cognitive Reasoning: A Study of Response Times." *Economic Journal* 117, (2007), 1243-1259.
- Sargent, T.J. *Bounded Rationality in Macroeconomics*. Oxford: Oxford University Press, (1993).
- Sonnemans, J., C.H. Hommes, J. Tuinstra, and H. van de Velden. "The Instability of a Heterogeneous Cobweb Economy: A Strategy Experiment on Expectation Formation." *Journal of Economic Behavior and Organization* 54, (2004), 453-481.
- Sonnemans, J., and J. Tuinstra. "Positive Expectations Feedback Experiments and Number Guessing Games as Models of Financial Markets." *Journal of Economic Psychology* 31(6), (2010), 964-984.
- Sutan, A., and M. Willinger. "Guessing with Negative Feedback: An Experiment." *Journal of Economic Dynamics and Control* 33(5), (2009), 1123-1133.
- Vriend, N.J. "An Illustration of the Essential Difference between Individual and Social Learning, and its Consequences for Computational Analyses." *Journal of Economic Dynamics and Control* 24(1), (2000): 1-19.
- Woodford, M. "Macroeconomic Analysis without the Rational Expectations Hypothesis." *Annual Review of Economics* 5(1), (2013), 303-346

A Experimental Instructions

A.1 Experimental Instructions (Monopoly)

Experimental Instructions

Welcome to this experiment in economic decision-making. Please read these instructions carefully as they explain how you earn money from the decisions you make in today's experiment. There is no talking for the duration of this session. If you have a question at any time, please raise your hand and your question will be answered in private.

General information

Imagine you are an advisor to a farm that is the only supplier of a product in a local market. In each time period the owner of the farm needs to decide how many units of the product he will produce. To make an optimal decision each period, the owner requires a good prediction of the market price of the product in each period. As the advisor to the farm owner, you will be asked to predict the market price, p_t of the product during 50 successive time periods, $t=1,2,\dots,50$. Your earnings from this experiment will depend on the accuracy of your price predictions alone. The smaller are your prediction errors, the greater will be your earnings.

About the determination of the market price p_t

The actual market price for the product in each time period, t , is determined by a market clearing condition, meaning that it will be the price such that demand equals supply for that period.

The amount demanded for the product depends on the market price for the product. When the market price goes up (down) the demand will go down (up). The supply of the product on the market is determined by the production decision of the farm owner. Usually, a higher (lower) price prediction by you causes the farm owner to produce a larger (smaller) quantity of the product which increases (decreases) the supply of the product on the market. Therefore, the actual market price p_t in each period depends upon your prediction, p_t^e , of the product's market price. More precisely, equating demand and supply, we have that the market price of the product is

determined according to:

$$p_t = \max(60 - \alpha p_t^e + \eta_t, 0)$$

This means that the price cannot be below 0. The parameter α is different for different local markets. You will see the α value for your own local market on your decision page during the experiment. This α parameter will remain the same for your local market for all 50 periods of the experiment. η_t is a small random shock to the supply caused by non-market (demand/supply) factors, for example, weather conditions. This small shock is randomly drawn each period and is sometimes positive, sometimes negative and sometimes zero. It is not correlated across periods. This small shock is normally distributed. The long term mean value of this small shock is 0, and the standard deviation is 1.

Here is an example: Suppose the parameter α is 0.8 in your local market. Suppose further that your price prediction for the period is 35, and the realization of the shock η_t is -0.2. Using the equation given above, the market price is then determined as:

$$p_t = 60 - 0.8 * 35 - 0.2 = 31.8$$

Note that in this case your forecast error, $|p_t^e - p_t|$, is $35 - 31.8 = 3.2$. This forecast error of 3.2 would determine your points for the period as discussed below.

Please also note that this example is for illustration purposes only. The value of the parameter α in your local market may be different from 0.8. The precise value of alpha and the equation for the determination of the market price in your local market is given on your decision page.

About your task

Your only task in this experiment is to correctly predict the market price in each time period as accurately as possible. The only constraint on your predicted price is that it cannot be less than zero (negative), since the actual price itself can never be less than zero. At the beginning of the experiment you are asked to give a prediction for the price of your farm's product in period 1. Note that, while there are several farms being advised by a forecaster like you in each period, these different local markets are totally separate from your own so what happens in other markets does not have any influence on your market. After all forecasters have submitted their predictions

for the first period, the local market price for period 1 will be determined and will be revealed to you. Based the accuracy of your prediction in period 1, your earnings will be calculated. Subsequently, you are asked to enter your prediction for period 2. When all forecasters have submitted their predictions for the second period, the market price for that period in your local market will be revealed to you and your earnings will be calculated, and so on, for all 50 consecutive periods.

Information

Following the first period, you will see information on your computer screen that consists of 1) a plot of all past prices together with your market predictions and 2) a table containing the history of your past forecasts and payoffs, as well as realized market price and the shock term η_t .

About your payoff

Your payoff depends on the accuracy of your price forecast. The earnings shown on the computer screen will be in terms of points. When your prediction is p_t^e and the market price is p_t your payoff is a decreasing function in your prediction error, namely the distance between your forecast and the realized price.

$$Payoff_t = \max[1300 - \frac{1300}{49}(p_t^e - p_t)^2, 0]$$

Recalling the example above, if your forecast error for the period t , $|p_t^e - p_t|$, was 3.2, then according to the payoff function you would earn 1028.33 points for the period.

Notice that the maximum possible payoff in points you can earn from the forecasting task is 1300 for each period, and the larger is your prediction error, $|p_t^e - p_t|$, the fewer points you earn. You will earn 0 points if your prediction error is larger than 7. There is a Payoff Table on your desk, which shows the points you can earn for various different prediction errors.

At the end of the experiment your *total points* earned from *all 50 periods* will be converted into Euros at the rate of 1 Euro for every 2600 points that you earned. Thus, the more points you earn, the greater are your Euro earnings.

Questions?

If you have questions about any part of these instructions at any time, please raise your hand and an experimenter will come to you and answer your question in private.

A.2 Experimental Instructions (Oligopoly)

Welcome to this experiment in economic decision-making. Please read these instructions carefully as they explain how you earn money from the decisions you make in today's experiment. There is no talking for the duration of this session. If you have a question at any time, please raise your hand and your question will be answered in private.

General information

Imagine you are an advisor to a farm that is **one of the three** main suppliers of a product in a local market. In each time period the owner of the farm needs to decide how many units of the product he will produce. To make an optimal decision, the owner requires a good prediction of the market price of the product in each period. As the advisor to the farm owner, you will be asked to predict the local market price, p_t of the product during 50 successive time periods, $t = 1, 2, 3, \dots, 50$. Your earnings from this experiment will depend on the accuracy of your price predictions alone. The smaller are your prediction errors, the greater will be your earnings.

About the determination of the market price p_t

The actual market price for the product in each time period, t , is determined by a market clearing condition, meaning that it will be the price such that demand equals supply for that period.

The amount demanded for the product depends on the market price for the product. When the market price goes up (down) the demand will go down (up). The supply of the product on the market is determined by the production decision of the farm owners. Usually, a higher (lower) price prediction by the advisors causes the farm owners to produce a larger (smaller) quantity of the product which increases (decreases) the supply of the product on the market. Therefore the actual market price p_t in each period depends upon the average prediction, \bar{p}_t^e of the product's market price. For example, if the predictions made by the advisors are $p_{1,t}^e$, $p_{2,t}^e$ and $p_{3,t}^e$ respectively, $\bar{p}_t^e = \frac{1}{3}(p_{1,t}^e + p_{2,t}^e + p_{3,t}^e)$. Equating demand and supply, we have that the

market price of the product is determined according to:

$$P(t) = 60 - \alpha \bar{p}_t^e + \eta_t$$

This means that the price cannot be below 0. The parameter α will be shown on your decision page during the experiment. This α parameter will be the same for all three farms in your local market and for all 50 periods. Note also that η_t is a small random shock to the supply caused by non-market (demand/supply) factors, for example, weather conditions. This small shock is randomly drawn each period and is sometimes positive, sometimes negative and sometimes zero. It is not correlated across periods. This small shock is normally distributed. The long term mean value of this small shock is 0, and the standard deviation is 1.

Here is an example: Suppose the parameter α is 0.8 for all three farms in your market. Suppose further that your prediction for the price is 30 and the predictions by the other two advisors in your market are 35 and 40 respectively. Finally, suppose that the realization of the shock, η , is -0.2. The market price in your three farm local market is then determined as follows:

$$p_t = 60 - 0.8 \times \frac{1}{3}(30 + 35 + 40) - 0.2 = 31.8$$

Note that in this case your forecast error (the distance between your forecast and the market price), $|p_t^e - p_t|$, is $|30 - 31.8| = 1.8$. This forecast error would be used to determine your points for the period as discussed below.

Please also note that this example is for illustration purposes only. The value of the parameter may be different from 0.8. The precise value of α and the equation for the determination of the market price in your local market are given on your decision page.

About your task

Your only task in this experiment is to correctly predict the market price in each time period as accurately as possible. The only constraint on your predicted price is that it cannot be less than zero (negative), since the actual price itself can never be less than zero. At the beginning of the experiment you are asked to give a prediction

for the price in period 1. There are several markets of various products and each of them consists of 3 farms, and each of the farms is advised by a forecaster like you. These different local markets are totally separate from your own market so what happens in other markets does not have any influence on your market. After all forecasters have submitted their predictions for the first period, the local market price for period 1 will be determined and will be revealed to you. Based on the accuracy of your prediction in period 1, your earnings will be calculated. Subsequently, you are asked to enter your prediction for period 2. When all forecasters have submitted their predictions for the second period, the market price for that period in your local market will be revealed to you and your earnings will be calculated, and so on, for all 50 consecutive periods.

Information

Following the first period, you will see information on your computer screen that consists of 1) a plot of all past market prices together with your market price forecasts and 2) a table containing the history of your past forecasts and payoffs, as well as realized market prices and the shock term, η_t .

About your payoff

Your payoff depends on the accuracy of your price forecast. The earnings shown on the computer screen will be in terms of points. When your prediction is and the market price is your payoff is a decreasing function of your prediction error, namely the distance between your forecast and the realized price. Specifically:

$$payoff = \max\left[1300\left(1 - \frac{(p_t^e - p_t)^2}{49}\right), 0\right]$$

Notice that the maximum possible payoff in points you can earn from the forecasting task is 1300 for each period, and the larger is your prediction error, the fewer points you earn. You will earn 0 points if your prediction error is larger than 7. There is a Payoff Table on your desk, which shows the points you can earn for various different prediction errors.

At the end of the experiment your total points earned from all 50 periods will be converted into Euros at the rate of 1 Euro for every 2600 points that you earned. Thus, the more points you earn, the greater are your Euro earnings.

Questions?

If you have questions about any part of these instructions at any time, please raise your hand and an experimenter will come to you and answer your question in private.

B Payoff Table

Table 9 is the payoff table used in this experiment.

Payoff Table for Forecasting Task							
Your Payoff= $\max[1300 - \frac{1300}{49}(\text{Your Prediction Error})^2, 0]$							
2600 points equal 1 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	<i>error</i> \geq 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		

Table 9: Payoff Table for Forecasters

C Alternative Proof of the Stability Condition of Adaptive Learning based on Mathematical Induction

Without loss of generality, let $p_1^e = p^* + \Delta$, where Δ is the difference between the period 1 prediction and the REE. Substituting this forecast into equation (1), we obtain $p_1 = \mu + \alpha(p^* + \Delta)$. Since $p^* = \mu + \alpha p^*$, this expression simplifies to $p_1 = p^* + \alpha\Delta$. In period 2, the prediction is the price in period 1, $p_2^e = p_1 = p^* + \alpha\Delta$. Substituting this prediction into equation (1) and simplifying, yields $p_2 = \mu + \alpha p_2^e = p^* + \alpha^2\Delta$. In period 3, the prediction should be the average price in periods 1 and 2, $p_3^e = \frac{p_1 + p_2}{2} = p^* + \frac{1}{2}\alpha(\alpha + 1)\Delta$. Substituting this prediction into equation (1) and simplifying yields $p_3 = \mu + \alpha p_3^e = p^* + \frac{1}{2}\alpha^2(\alpha + 1)\Delta$. By iterating in this fashion it is not difficult to find that in general, for period t , $p_t^e = \frac{1}{t-1} \sum_{s=1}^{t-1} p_s = p^* + \frac{\alpha(\alpha+1)(\alpha+2)\dots(\alpha+t-2)}{1 \times 2 \times 3 \dots (t-1)} \Delta$ and so $p_t = \mu + \alpha p_t^e = p^* + \alpha \frac{\alpha(\alpha+1)(\alpha+2)\dots(\alpha+t-2)}{1 \times 2 \times 3 \dots (t-1)} \Delta$.

Clearly this system converges to the REE whenever the ratio $\frac{\alpha(\alpha+1)(\alpha+2)\dots(\alpha+t-2)}{1 \times 2 \times 3 \dots (t-1)}$ goes to 0. This ratio consists of $t - 1$ components in both the numerator and the denominator. We can pair the components in the numerator and the denominator according to the sequence, namely, let α be paired to 1, $\alpha + 1$ be paired to 2, ..., $\alpha + t - 2$ be paired to $t - 1$. When $\alpha > 1$, each component of the numerator is larger than its paired number in the denominator. Therefore $\frac{\alpha(\alpha+1)(\alpha+2)\dots(\alpha+t-2)}{1 \times 2 \times 3 \dots (t-1)}$ will increase over time with t , diverging away from 0. When $\alpha = 1$, the ratio is exactly equal to 1. When $-1 < \alpha < 1$, each component in the numerator has a smaller absolute value than its paired number, so the ratio will decrease with t , and goes to 0 as $t \rightarrow \infty$.

When $\alpha < -1$, we make a slightly different re-matching of the components in the numerator and the denominator. First, let m be an integer such that $\alpha + m - 1 < 0$ and $\alpha + m > 0$. We re-state the ratio as $\frac{\alpha(\alpha+1)(\alpha+2)\dots(\alpha+m-1)(\alpha+m)(\alpha+m+1)\dots(\alpha+t-2)}{1 \times 2 \times 3 \dots (t-m-1)(t-m)(t-m+1)\dots(t-1)}$. We then “cut” the numerator into two parts, $N_1 = \alpha(\alpha + 1)(\alpha + 2)\dots(\alpha + m - 1)$ and $N_2 = (\alpha + m)(\alpha + m + 1)\dots(\alpha + t - 2)$, and we also cut the denominator into two parts, $D_1 = 1 \times 2 \times 3 \dots (t - m - 1)$ and $D_2 = (t - m)(t - m + 1)\dots(t - 1)$. We pair N_2 to D_1 , namely, $\alpha + m$ to 1, $\alpha + m + 1$ to 2, ... $\alpha + t - 2$ to $t - m - 1$. It is not difficult to see that each item in N_2 is smaller than the paired item in D_1 ($\alpha + m < 1$, $\alpha + m + 1 < 2$, ... $\alpha + t - 2 < t - m - 1$), and therefore that $\frac{(\alpha+m)(\alpha+m+1)(\alpha+m+2)\dots(\alpha+t-2)}{1 \times 2 \times 3 \dots (t-m-1)}$ decreases with t , and goes to 0 as $t \rightarrow \infty$. There remain m extra components in both the numerator

and the denominator. In the numerator, $|N_1| = |\alpha(\alpha + 1)(\alpha + 2)\dots(\alpha + m - 1)| < |\alpha^m|$ is a finite number, while in the denominator, $D_2 = (t - m)(t + m + 1)\dots(t - 1)$ goes to infinity as $t \rightarrow \infty$. Therefore, the remaining fraction $\frac{\alpha(\alpha+1)(\alpha+2)\dots(\alpha+m-1)}{(t-m)(t-m+1)\dots(t-1)}$ also goes to 0 as $t \rightarrow \infty$. It follows that, under adaptive (least squares) learning, the system converges to the REE provided that $\alpha < 1$ and diverges from the REE only if $\alpha > 1$.

D Categorization of Subjects

	$\alpha = -0.5$		$\alpha = -0.9$		$\alpha = -2$		$\alpha = -4$	
	Categorized	Reported	Categorized	Reported	Categorized	Reported	Categorized	Reported
exp1	A		E				A	E
exp2	A		A					E
exp3	A				A			E
exp4			E		E		A	E
exp5	A							
exp6	A				E		A	A
exp7							A	E
exp8	E		A				E	E
exp9	A		E	E		A		E
exp10	E	E		E	E	A		E
exp11	E		E			E		E
exp12	A	A		A		E		
exp13	A	E			E	E		A
exp14		E					A	E

Table 10: Categorization of subjects into adaptive and eductive learners in the monopoly setting. “A” means adaptive learner. “E” means eductive learner. We leave the cell blank for subjects we cannot categorize into either of the two types. “Categorized” means categorization according to the first approach where we use the definition of the learning rules. “Reported” means categorization is done according to the second approach based on self-reported strategies.

	$\alpha = -0.5$		$\alpha = -0.9$		$\alpha = -2$		$\alpha = -4$	
	Categorized	Reported	Categorized	Reported	Categorized	Reported	Categorized	Reported
exp11	E	E	E					E
exp12	A	A		A		E		
exp13	A	A		A		E		E
exp21	A	E	A	A	E			A
exp22	E	E	A	A		E		A
exp23	A	A	A	A		A	E	E
exp31	A	E	A	A		A		E
exp32	A	A	E	A	E	E	E	E
exp33		A		A		A	E	
exp41	A	A		A	E	E		E
exp42		A		A		A	E	E
exp43	A	E		A		E		
exp51	E	E	A	A		A		E
exp52	A	A		A	E	A	E	A
exp53	A	A	E	E				A
exp61	E	A		A		E		E
exp62	A	A		A		E		A
exp63	A	E	A	E				A
exp71	A	A		A		A	E	E
exp72		A	E	E		A		A
exp73	A	A		A		E	E	E
exp81		A	A	E		E		E
exp82	E	E	E	A		A	E	E
exp83	A	E	E	E			E	E
exp91	A	A		A		E		E
exp92	E	A	E	E		A		
exp93		A	E	A		A	E	E
exp101		A	E	A		A		A
exp102	A	A	E	A				A
exp103	E	E	A	A		A		A
						A		
						A		
						E		

Table 11: Categorization of subjects into adaptive and eductive learners in the oligopoly setting. “A” means adaptive learner. “E” means eductive learner. We leave the cell blank for subjects we cannot categorize into either of the two types. “Categorized” means categorization according to the first approach where we use the definition of the learning rules. “Reported” means categorization is done according to the second approach based on self-reported strategies.