

The Instability of a Heterogeneous Cobweb economy: a Strategy Experiment on Expectation Formation

Joep Sonnemans^a, Cars Hommes^b, Jan Tuinstra^b and Henk van de Velder^b

^a Corresponding author, Center for Nonlinear Dynamics in Economics and Finance (CeNDEF) and CREED

^b Center for Nonlinear Dynamics in Economics and Finance (CeNDEF)

Department of Economics and Econometrics
University of Amsterdam
Roetersstraat 11, 1018WB Amsterdam
The Netherlands
Phone +31 20 5254249
Fax +31 20 5255283
Email J.H.Sonnemans@uva.nl

January 2003, *Journal of Economic Behavior and Organization*, forthcoming

Abstract

Which strategies do agents use when forming expectations about future prices, and how often does this lead to stable or unstable outcomes? We performed a four-round strategy experiment in a cobweb economy with expectations feedback. Subjects received feedback about their performance, and could revise their strategy. Over the rounds forecasting errors decrease and realized market prices move close to the rational expectations steady state, but the complexity of the price fluctuations also increases. Convergence to the unique RE steady state occurs in less than 10% of all cases. In the final round 60% of the price fluctuations appears to be chaotic. Heterogeneous interaction of simple prediction strategies seems to be the main source of the endogenous price fluctuations, frequently leading to a boundedly rational equilibrium of 'close to the steady state chaos'.

Keywords: heterogeneous agents, bounded rationality, strategy experiment, cycles and chaos.

JEL-codes: E32, E37, D84.

*Financial support by the Netherlands' Organization for Scientific Research (NWO) by a NWO-MaG Pionier grant is gratefully acknowledged. We thank Abdolkarim Sadrieh and Claudia Keser for their suggestions in designing the strategy experiment. An earlier version of this paper was presented in a seminar at the Universität Bonn and we thank Reinhard Selten and his coworkers for their comments. We also thank participants of the Annual Meeting of the Economic Science Association in Lake Tahoe, May 27-30, 1999, and participants of the Seventh Viennese Workshop on Optimal Control, Dynamic Games and Nonlinear Dynamics: Theory and Applications in Economics and OR/MS, May 24-26, 2000 in Vienna for their comments.

1. Introduction

Expectations play a key role in modeling dynamic phenomena in economics and finance. Since the pioneering papers by Muth (1961) and Lucas (1971), the Rational Expectations Hypothesis (REH) has become the “currently still” dominating paradigm in expectation formation. According to the REH, agents’ subjective expectations equal the objective mathematical expectations conditional upon available information. In economic modeling the REH usually assumes perfect knowledge of the underlying market equilibrium equations and agents are assumed to use these equations to compute their rational expectations (RE) forecast.

The bounded rationality literature has recently put forward two important criticisms concerning the REH. The first is that it is unrealistic to assume that agents have perfect knowledge of underlying market equilibrium equations. It would be more reasonable to assume that agents derive their expectations from time series observations; see e.g. Sargent (1993,1999) for an overview of and many references to recent work on bounded rationality. The second criticism is that in a heterogeneous world, realized market prices depend upon beliefs of *all* market participants. Even if agents would have perfect knowledge about market equilibrium equations, rational expectations can only be achieved under the unrealistic assumption that the agents also have perfect knowledge about the *beliefs* of *all* other agents in the market. As an alternative to a world where all agents are perfectly rational, several authors have proposed *heterogeneous agents* models with boundedly rational agents using simple strategies; see for example computationally oriented work by Arthur et al. (1997) on the Santa Fe artificial stock market, theoretical work on evolutionary dynamics in the cobweb model by Brock and Hommes (1997), work on herd behavior and mimetic contagion in speculative markets by Topol (1991), Kirman (1993), Lux (1995) and Brock and Hommes (1998), work on noise traders in finance by De Long et al. (1990) and recent work in evolutionary game theory as surveyed e.g. by Fudenberg and Levine (1998).

It is hard to observe or obtain detailed information about individual expectations in real markets. One approach is by survey data analysis, as done for example by Frankel and Froot (1987) and Allen and Taylor (1990) on exchange rate expectations and Shiller (1989,2000) on stock market data. Economic experiments are well suited for a detailed investigation of expectation formation in a controlled dynamic environment. Unfortunately, as for example pointed out in Sunder (1995), only little experimental work on expectation formation has been done. Some exceptions are the well known 'bubble experiments' of Smith et al. (1988), the overlapping generations experiments by Marimon, Spear and Sunder (1993) and the inflationary economy experiments in Marimon and

Sunder (1993). Recently, Hommes, Sonnemans and van de Velden (2000a) and Hommes, Sonnemans, Tuinstra and van de Velden (2000b) investigated expectation formation in an experimental cobweb economy. In most of their experiments, prices do *not* converge to the unique rational expectations (RE) steady state, but keep fluctuating irregularly in a neighbourhood of the (unstable) steady state, suggesting expectations driven excess price volatility. In the present paper we are particularly concerned with individuals who form expectations *repeatedly* and have ample opportunity to *learn* and change prediction strategies. One of the main questions of this study is whether these learning opportunities enforce convergence to the unique RE steady state or whether a market with repeated learning may still exhibit excess volatility.

Most economic experiments take less than 2 hours and subjects only make a couple of dozens decisions. In such experiments not only are learning possibilities limited (especially in complicated situations), but it is also hard to detect exactly which expectation formation rules subjects use (because of the relatively few decisions made by the subjects). A strategy experiment in the spirit of Selten et al. (1997) seems to be well suited for our purposes. Our complete four-round strategy experiment lasted seven weeks. The strategy method is becoming a popular tool in experimental economics: see e.g. Brandts and Schram (2001), Keser (1992), Offerman et al. (2001), Sonnemans (1998), as well as the classic work of Axelrod (1984).

In our cobweb strategy experiment, subjects are asked to formulate a complete strategy, that is, a description of all their forecasts in all possible states of the world (e.g. history of prices). In each period all strategies that participate in the market forecast the next price. The realized market equilibrium price is then determined by a fixed, but unknown, (linear) demand curve and (nonlinear) supply, depending upon individual expected market prices, aggregated over all producers. The realized market price thus depends on all individual strategies. Subjects gain experience in forecasting next period's price in an introductory experiment before submitting their first strategy. These strategies are then programmed and simulated. After each round, subjects receive feedback about the relative performance of their strategy, and the outcomes of five randomly selected simulations in which their strategy is included. Subjects had one week to revise their strategy for the next round. In each of the four rounds of the strategy experiment (as well as in the introductory experiment), financial incentives, based upon prediction performance, were used to motivate the subjects.

The role of expectations is particularly important in speculative markets. Unfortunately, a complicating feature of dynamic asset pricing models is the existence of *multiple* RE equilibria in the

form of so-called speculative bubble solutions (see e.g. Cuthbertson 1996). Therefore, we focus on a simpler dynamic environment, namely the cobweb or 'hog cycle' model. A convenient feature of the cobweb model is that it has a *unique* RE equilibrium: the steady state price where demand and supply intersect. Since its introduction in the thirties (see e.g. Ezekiel 1938), the cobweb model has become one of the classical examples in economic dynamics. Nerlove (1958) introduced adaptive expectations into the cobweb model, whereas Muth used the cobweb model to introduce rational expectations. More recently, in the bounded rationality learning literature, the cobweb model has been used as a benchmark example to show that adaptive learning by ordinary least squares (Bray and Savin 1986), by genetic algorithms (Arifovic 1994) or by sample average and sample autocorrelations (Hommes and Sorger 1998) enforces convergence of prices to the unique RE equilibrium, even when demand and supply are unknown and agents only observe past prices. In general however, adaptive learning may also be destabilizing as emphasized, for example, by Grandmont and Laroque (1991) and Grandmont (1998). In particular, in a cobweb model with nonlinear, but monotonic demand and supply curves, adaptive expectations (which is in fact just an adaptive learning scheme with a constant gain factor) can lead to higher order stable periodic cycles or even chaotic price fluctuations (Chiarella 1988 and Hommes 1994). Furthermore, Brock and Hommes (1997) use the cobweb model to show that evolutionary competition between heterogeneous forecasting rules can destabilize the RE steady state and can lead to periodic or chaotic price fluctuations.

The main research questions of the present study are (1) What kind of strategies do subjects use? (2) Will prices in markets with heterogeneous agents converge to the unique RE steady state, or will market instability, price fluctuations and excess volatility prevail in a heterogeneous world? (3) How does learning affect the strategies and the price dynamics in the consecutive rounds? (4) Can market stability or instability be attributed to characteristics of individual strategies, or is heterogeneity the fundamental cause?

We find that subjects use a wide variety of strategies and that convergence to the RE steady state is relatively rare (less than 10%). Over rounds the amplitude of price fluctuations decreases, but at the same time the price dynamics becomes more irregular. In the final round about 60% of the price sequences are chaotic. Instability and excess volatility may be explained by a boundedly rational heterogeneous agents equilibrium.

The paper is organized as follows. Section 2 describes the design of the strategy experiment, and Section 3 discusses the main results. Finally, Section 4 concludes.

2. Design

Subjects

Since the experiment lasted about seven weeks many subjects might have been lost during the course of the experiment if we would have recruited in the normal way (by advertisements and bulletin boards). Therefore we recruited subjects in a course "Dynamical Systems," a mathematical introduction to dynamical systems in the undergraduate econometrics program. Participation to the experiment was on a voluntary basis and unrelated to the course itself. Students had no prior knowledge about dynamic economic systems. The cobweb model was only treated (briefly) as an economic example in the last week of the course, *after* the experiment had finished; this example was *not* treated in the textbook of the course. Subjects could hand in their strategies after class.

In the introductory experiment 29 subjects participated, earning on average 51 Dutch guilders (23 Euro) in approximately 2 hours. All 29 students handed in a first strategy. One of the students left the course after the first week, so that 28 students handed in their second strategy. The third and fourth (final) strategies were handed in by respectively 21 and 24 subjects.¹

The cobweb economy

The (unknown) cobweb model underlying the experiment has a *nonlinear*, but monotonically increasing supply curve, such that under simple forecasting rules the model can generate stable price cycles. In particular, we have chosen the nonlinear supply curve such that, if all producers would have naive price expectations (i.e. would expect tomorrow's price to equal today's price), the cobweb economy is unstable and prices converge to a stable 2-cycle (the well known 'hog-cycle' with constant up and down price oscillations). Parameters have been fixed such that under adaptive expectations a stable 4-cycle is the most complicated dynamical behaviour.²

¹ Originally we also planned an experiment after subjects submitted their final, fourth strategies, but before they would receive the final results. The main goal of this planned experiment was to study the relationship between actual behavior of subjects and the strategy they submitted (e.g. Sonnemans (2000)). We announced this experiment in the class at which the students submitted their final strategies, and we asked the subjects not to talk about their strategies yet. Unfortunately some students who did not attend that class but had handed in their final strategies already before the class were not informed about this experiment, and we had to cancel the experiment when we found out that many of these students, after handing in their fourth strategy, already heard about the successful strategies of the first three rounds.

² In general, as the weight factor of the adaptive expectations scheme changes, for a nonlinear, S-shaped supply curve bifurcation routes to chaos may arise (see Hommes (1994)). However, for our choice of the parameters in the experiment only the first two bifurcations from a stable steady state to a stable 2-cycle and to a stable 4-cycle arise, and fully developed chaos does not arise. An important motivation for this setup was whether in a heterogeneous world the strategies would be able to detect the regularities along the stable cycles and stabilize the system and enforce prices to converge to the RE steady state.

A market consists of six subjects (or strategies). There is no actual trade going on, but the realized market price depends upon the (unknown) demand and supply curves and individual expectations. The *price forecast* P_i of subject i determines the supply of that subject as follows:

$$S(P_i) = \frac{35 + 35 \tanh(0.25(P_i - C))}{6},$$

where C is a parameter determining the inflection point of the nonlinear S-shaped supply curve. It is important to note that the S-shaped supply curve is consistent with producers' expected profit maximization, since it can be derived from an increasing and convex cost function.³ The demand curve is linear and given by

$$D(P) = 20 - P + C$$

The realized price is determined by market clearing, that is, by demand equals aggregate supply

$$D(P) = \sum_{i=1}^6 S(P_i)$$

or equivalently by

$$P = 20 + C - \sum_{i=1}^6 S(P_i).$$

Figure 1 shows the demand and supply curves. In order to generate different levels of steady state equilibrium prices the parameter C is chosen randomly in each market, with C drawn from a uniform distribution over the interval $[50,80]$. This means that participants have to learn a different steady state equilibrium price for each market they enter. Notice, however, that as C changes *both* the demand and supply curves are shifted horizontally. Hence, for simple forecasting rules, the dynamics around the steady state remain qualitatively the same in all markets and only the price level differs.

Incentive structure

A good measure of the quality of a forecast is its quadratic forecasting error. In the introductory experiment subjects gained experience with this measure. In each market of 20 periods, subjects started with 25 Dutch guilders (approximately 11.35 Euro), and after each period an amount of 0.1 times the squared forecasting error (in cents) was subtracted from this amount. The payoff of a subject in a market of 20 periods was the amount left after period 20 if it was positive and 0

³ See e.g. Hommes (2000), where a similar S-shaped supply curve is derived from an increasing, convex polynomial cost function of degree 4 or higher.

otherwise. Three markets were played with different (randomly drawn) parameter values C . Total earnings of a subject were the sum of the earnings in the three markets.⁴




Figure 1. The supply and demand curves of the cobweb model. The parameter C varies between 50 and 80, so that the corresponding RE steady state varies between 48.39 and 78.39.

We recognized that a similar incentive structure in the strategy experiment might stimulate the subjects to cooperate and share their strategies and feedback results to optimize total earnings. For example, all subjects would benefit from a fast convergence to the equilibrium steady state. Therefore in the strategy experiment we employed a tournament incentive structure: payment was based upon relative performance of the strategies, and the performance was based upon the average squared prediction error of the strategy over all simulations in that round. The strategy with the smallest average quadratic forecasting error received 50 guilders (approximately 22.70 Euro) in rounds 1, 2, and 3, and in the final round, three prizes of 250, 150, and 50 guilders (113.60, 68.20, and 22.70 Euro respectively) were awarded. In addition to this students received a flat fee of 5 guilders (2.25 Euro) if they submitted a strategy.

⁴ It should be noted that minimizing the squared forecasting error is essentially the same as maximizing producers' profits, since the forecasting error function has its minimum at the steady state and the producers' profit function has its maximum at the steady state.

A possible disadvantage of payments based upon relative performance is that subjects may try to destabilize markets to make it harder for the other market participants to forecast prices. However, it is easy to see that this cannot work. If all strategies predict the equilibrium price except for one strategy that tries to destabilize the market by predicting a higher (lower) price, the realized price will be lower (higher) than the equilibrium price. The destabilizing strategy will thus end up with a larger quadratic error than the other strategies in that market. Even more importantly, one can only affect realized prices in the market in which the strategy participates; an increasingly unstable market will cause a comparative advantage of the strategies active in the other markets. No subject ever mentioned (in the questionnaires or in class after the experiment ended) that he or she had tried to destabilize markets.

Procedures

Introductory experiment. The goal of the introductory experiment was to give the subjects some experience in their forecasting task. The experiment was completely computerized and took place in the CREED experimental laboratory. Understanding of the instructions was checked by control questions. It was explained to the subjects that they only had to predict prices and that producers would decide how much to produce, based upon their predictions. The computer program would play the role of both producers and consumers and would calculate the realized market prices. See the appendix for the precise instructions.

Subjects played 3 consecutive markets of 20 periods with different parameters C . After the third market they received instructions for the strategy part of the experiment (see appendix 1).

Strategies. The subjects formulated their first strategy in the laboratory, immediately after the introductory experiment. The experimenters checked these strategies for clarity, completeness (the strategy provides a prediction in all possible situations), uniqueness (the strategy always provides exactly *one* prediction), and informational correctness (the strategy does not use information that is not available, such as future prices or previous predictions of other strategies). An example of a strategy form is included in appendix 2.

Questionnaires. Subjects filled in a small questionnaire every round. They were asked about (among other things) their considerations when changing their strategy, the effect of the feedback upon their new strategy, whether they had talked with other subjects about the experiment, and how effective they thought their new strategy would be.

Feedback was provided one week after the strategies were submitted. The feedback consisted of a general ranking of the strategies by mean quadratic forecasting error (1 page) and two personal pages containing for each subject 5 randomly chosen simulations in which his or her strategy participated. An example of the feedback information is included in Appendix 3.

Handing in strategies. After students received feedback about their latest strategy, they had one week's time to hand in a new strategy. Students could hand in their strategies during class (twice a week) or give them directly to one of the organizers. All strategies were checked immediately.

3. Results

This section starts with a description of some characteristics of the submitted strategies. Thereafter the short, medium, and long run dynamics will be discussed. We find complicated price fluctuations and evidence of chaotic behavior. In the final part of this section an attempt is made to find the cause of this unstable behavior: is this instability caused by *individual* strategies or is it due to *interactions between* strategies?

3.1 Characteristics of the strategies

In most empirical studies in market dynamics, researchers have only access to the sequence of realized prices since the underlying exact expectations rules used by the market participants cannot be observed. One of the nice features of the present study is the availability of the explicit strategies. Therefore we will first turn to the question: "what kind of strategies do subjects use?"

As can be seen in table 1 a total of 102 strategies were submitted. The strategies are all different, although some subjects made only minor changes between rounds. It is impossible to describe all strategies in detail; therefore we will focus on some general characteristics.

	<i>Round 1</i>	<i>Round 2</i>	<i>Round 3</i>	<i>Round 4</i>	<i>Total</i>
Number of strategies	29	28	21	24	102
Continuous	16	12	11	12	51
Simple adaptive (+conditionally simple adaptive)	4 (+4)	3 (+3)	2 (+3)	2	11 (+10)
More complicated adaptive (+conditionally complicated adaptive)	3 (+1)	6	5 (+1)	5 (+5)	19 (+7)
Do not include prediction(s) of previous period(s) at all	15	15	8	12	50
Includes (weighted) average of previous prices	16	12	9	12	49
Complexity: Number of code lines	18.5	23.8	26.6	28.1	23.9
Simulation time (round 1=100)	100	100	173	295	

Table 1: Characteristics of the strategies

All strategies use the same format (see Appendix 2). All start with a prediction for the first period. Predictions in the subsequent periods can be conditional on the period (many subjects use a strategy that differs in the first few periods from later periods) and can be conditional on the history (previous realized prices and own predictions). Many strategies list conditions under which specific sub-strategies are to be used. An example of such a strategy is “if the last two realized prices differ more than 50, I will take the average of these prices; otherwise I will take the last price as my prediction.” Note that on the border between the two conditions (where the last two prices differ exactly 50), this strategy is discontinuous. Strategies are classified as continuous if arbitrary small changes in the history always result in small changes in the prediction. Only half of the strategies are continuous (see Table 1). However, note that some discontinuities may have little effect on the dynamics because in many markets some of the conditions are never satisfied. (See also Section 3.3 where we look at the relation between stability of markets and characteristics of the participating strategies.)

In a simple form of adaptive expectations, the predicted price is a weighted average of the previous predicted price and the previous realized price. Given that subjects have no information about the underlying model, such a simple adaptive strategy seems natural. However, only about 10% of the strategies are adaptive in this sense, whereas another 10% are conditionally adaptive (that is, the strategy is adaptive only if the sequence of past prices and/or predictions fulfils certain conditions). Another 25% of the strategies seem to use a more complicated kind of adaptive

expectations (some of them conditionally)⁵. Half of the strategies do not include previous predictions at all.

Approximately half of the strategies use a weighted average of previous prices somewhere in the strategy. Some of these strategies use all previous prices, whereas others only use recent ones. Not all previous prices have the same weight if subjects try to anticipate cycles. For example, a subject anticipating two-cycles may overweight the realized price of two periods back.⁶

Finally, strategies have a tendency to become more complicated during the experiment. A rough but simple measure of the complexity of a strategy is the number of lines used in the simulation program. The average number of lines increased from 18 in round 1 to 28 lines in the final round.⁷ Another measure of complexity is the time needed for the computer simulations. Both complexity measures are displayed in the last row of table 1.

It would be interesting to know the relationship between the submitted strategies and the actual behavior of subjects in an experiment. When the subjects submitted their first strategy, we asked in a questionnaire to what extent and in what way their strategy differed from their behavior in the experiment. On a 7-point scale from completely different to exactly the same the average score was 5.13⁸. Some subjects were not satisfied with the results of the experiment and adapted their strategy accordingly, but the most common remark was that in the experiment decisions were based upon 'feeling' while the strategies are (necessarily) based upon exact calculations. For example, a subject who uses an average in the strategy indicated that in the experiment she estimated but did not calculate the average exactly. A strategy may thus be seen as an attempt to quantify the more intuitive habitual rule of thumb used during the experiment.

⁵ For example, the first strategy of subject 8 was as follows (for $t > 2$). IF $((P(t-2) \leq P_e(t-2) \text{ AND } P(t-1) \leq P_e(t-1)) \text{ OR } (P(t-2) \geq P_e(t-2) \text{ AND } P(t-1) \geq P_e(t-1)))$ THEN $P_e(t) = (P(t-1) + P_e(t-1))/2$ ELSE $P_e(t) = (P(t-1) + P(t-2) + 2P_e(t-1))/4$.

⁶ Note that a strategy that uses a weighted average of previous prices can behave very much like an adaptive strategy. In an adaptive strategy the prediction in period t for period $t+1$ will be $P_e(t+1) = (1-w) * P_e(t) + w * P(t) = (1-w)^2 * P_e(t-1) + (1-w) * w * P(t-1) + w * P(t)$, etc, which is exactly the same as a weighted average of past prices with exponentially decreasing weights $w * (1-w)^k$ for $P(t-k)$. However, one would expect a subject who wants to use an adaptive strategy to use the simple adaptive rule instead of a complicated weighted average rule.

⁷ This measure may underestimate the increasing complexity because the programming may have become more efficient during the experiment due to increased experience of the programmer.

⁸ Only two subjects indicated that their strategy was very different from their decisions in the experiment (score 2 and 3 on the 7-point scale). The frequencies are: 2 (1), 3 (1), 4 (4), 5 (12), 6 (9) and 7 (2).

3.2 Dynamics

Short run dynamics: Success of the strategies

The incentives of the subjects are based upon the behavior of their strategies in the first 20 periods. Therefore the first 20 periods are of special interest. How close do the realized prices come to the RE steady state? Figure 2 shows the quadratic distance to the RE steady state over the periods. Recall that a different parameter C is used in each market and therefore each market has a different RE steady state that has to be learned again. Two important characteristics are seen in Figure 2. Firstly, the distance between the realized market price and the RE steady state is largest in period 1 and decreases afterwards (the exception is in round two, where the largest distance is found in period 3). Almost no improvement is observed after the seventh period. Hence, in each round a short learning phase of about 7 periods is needed to learn the new 'price level,' after which the average distance to the RE steady state remains approximately constant. The second important characteristic is that there is clear evidence of learning between rounds; the prices in the third and fourth round move much

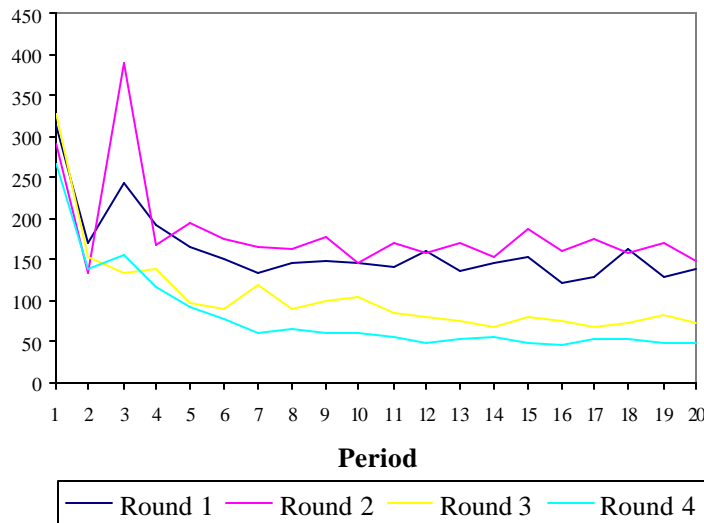


Figure 2. *The mean quadratic distance to RE steady state in each period (620 simulations per round).*

	<i>Round 1</i>	<i>Round 2</i>	<i>Round 3</i>	<i>Round 4</i>
Mean quadratic forecasting error over all strategies	6448	8878	4608	3814
Mean quadratic error winner	3194	2946	2628	2016
Mean variance first 20 prices	158.59	179.75	102.57	78.27
Winning strategy	Subject 15 $P_e(t)=0.75P_e(t-1)+0.25P_{\text{mean}}$	subject 14 $P_e(t)=1.5P_{\text{mean}}-0.5P(t-1)$	subject 15 $P_e(t)=0.75P_e(t-1)+0.25P_{\text{mean}}$	subject 18 $P_e(t)=P_e(t-1)+(P(t-1)-P_e(t-1))^{1/(2t-1)}$

Table 2: Results of the simulations (620 simulations per round). In each simulation a market is formed with 6 strategies that are submitted in that round. The quadratic error is the total error in one market (20 periods), the mean quadratic error of a strategy is calculated by averaging the total error in all simulations in which that strategy participated. The last row presents the prediction formulas of the winning strategies in the later periods, the predictions in the early periods are not displayed to limit the size of the table.

closer to the RE steady state price than in the first two rounds. More experience thus leads to price fluctuations closer to the RE steady state.

Table 2 summarizes the main results of the simulations in each of the four rounds. The mean quadratic forecasting error over all strategies decreases after round 2. Subjects thus learn to make better forecasting strategies during the experiment⁹. The mean quadratic error of the winner decreases over the rounds from about 3200 to 2000.

We can compare the numbers of table 2 with the results of homogeneous naive or adaptive players. If all players in a market started with a prediction of 50 and for the next periods always predicted the previous price, they would very fast end up in a two-cycle (C-50,C+20), the mean quadratic error would be 94,308, and the variance over the prices would be 1290. If all market participants start with a prediction of 50 and use the adaptive rule $P_e(t)=(P(t-1)+P_e(t-1))/2$, the result would be a four-cycle (C-14.6,C+18.2,C-30.7,C+20). In that case the mean quadratic error

⁹ The number of strategies is lower in rounds 3 and 4 than in the first two rounds. The relatively bad performance in rounds 1 and 2 is not caused by subjects who left the experiment after round 2. The missing subjects in round 3 and 4 did equally well in round 1 and 2 as the other subjects (based upon the mean quadratic error).

would be 24,899 and the variance of the prices 734. The adaptive rule $P_e(t)=(P(t-1)+3*P_e(t-1))/4$ results in a two-cycle (C-12.7, C+19) with mean quadratic error 4361 and price variance of 144¹⁰. Compared to these numbers the (winning) strategies of our subjects do very well.

Both Figure 2 and Table 2 show that the strategies in round 1 perform better than the strategies in round 2. Additional analyses were done to study this surprising increase in forecasting errors from round 1 to 2. Only 5 subjects out of 28 had a lower average prediction error in round 2 than in round 1. Of course, the quality of a strategy depends also on the other strategies involved. Therefore we investigated for each subject whether the results would have been better if that subject would not have changed the strategy from round 1 to 2 (while the others would). For each subject we ran simulations in which the strategy of the first round was coupled with 5 strategies of the second round (of other subjects). For 15 subjects the first round strategy had a lower average prediction error than the second round strategy, whereas for 13 subjects it is the other way around. In this sense the strategies of round 2 were not (at least not statistically significant) of a lower quality than the strategies of the first round. The errors in these simulations were also compared with the errors of the first round simulations. All (29) first round strategies do better when they are coupled with other first round strategies than when they are coupled with second round strategies (Wilcoxon test $p<0.000$). We conclude that the increasing errors from round 1 to round 2 are apparently due to *interaction of strategies* of round 2, making prediction harder for any strategy.

The third row of table 2 shows the volatility of the prices, as measured by the variance, over the rounds. Prediction is easier if volatility is low, and indeed, we see the same pattern as in the first row (quadratic errors). Volatility is much smaller in rounds three and four. Note that the volatility of the prices is very low compared with the case of a homogeneous population of naive or adaptive players and the corresponding stable 2- or 4-cycles (see above).

The last row of table 2 shows the formulas of the winning strategies in the later periods. All winning strategies tend to be relatively simple. The winning strategy in round 1 and 3 is an adaptive strategy in which the prediction is adapted in the direction of the mean price. The winning strategy of round 2 anticipates 2-cycles, and the winner of the final round uses an adaptive strategy that is much more adaptive in early periods than later on.

¹⁰ The exact mean quadratic errors and variance of the prices depend on C because of the first few forecasts (e.g. the first forecast of 50 is better if C=50 than if C=80). The variances and quadratic errors presented here are for the case C=65.

Medium run dynamics

In many economic applications the short term is as least as important as the long term. The horizon of the subjects was 20 periods, and we sometimes observe convergence to a steady state or to cycles even in this very short term¹¹. However, within the first 20 periods it is hard to observe cycles of a period longer than two or three, and it may be hard to see whether what looks like a cycle will eventually converge to a steady state. Therefore the medium run dynamics, periods 51 to 100, were studied. By our definition, a sequence of prices converges to a steady state if all prices in periods 51-100 are within a range of one point. A similar criterion is used for cycles; for example in a 2-cycle, the prices in all odd periods are within a 1-point range, and the prices for all even periods are within a 1-point range (and we do not observe a steady state).

Figure 3 shows the percentages of simulations that converge to a steady state or a cycle. The first thing to notice is that convergence to a steady state price is relatively rare, occurring only in about 10% of all cases. Many low period cycles are observed, but the number of cycles decreases over the rounds. The percentage of simulations that do not converge to a low period cycle at all increases from about 45% in round 1 to about 75% in round 4.




Figure 3: Medium term convergence to a steady state or a cycle. Based upon 620 simulations per round and periods 51-100.

¹¹ For example, in almost 5% of the simulations the prices of periods 16-20 are within a range of 1 point. The percentages for the rounds 1 to 4 are 5.5%, 1.9%, 5.5% and 6.1% respectively.

Long Run dynamics

Our results suggest that in most cases, prices do *not* converge to the RE steady state, not even after 100 periods. There are two possible scenarios, both consistent with this observed behavior. Either the RE steady state is *stable* but convergence of the learning process is *slow*, or the RE steady state is intrinsically *unstable*. Stated in dynamical systems terminology, either price oscillations are a *transient* phenomenon and the underlying system has a stable steady state, or the steady state is locally unstable and the long run dynamics is characterized by a *stable cycle* or even by chaotic price fluctuations on an underlying complicated *strange attractor*. In order to investigate which of the two scenarios explains the observed fluctuations in the short and medium run dynamics, it is thus important to study the properties of the long run dynamics and, in particular, to study the *attractors* of the system.

In order to obtain an accurate picture of the underlying attractors, we have focused upon periods 951 to 1000. Figure 4 shows the results for the long run dynamics. Convergence to a steady state price or a cycle is defined in the same way as in the medium run analyses. As in the medium run analysis, convergence to a steady state price is rare and less than 10% over all rounds. Compared with the medium run analysis, convergence to a low order stable cycle is observed more frequently, especially in round 3.

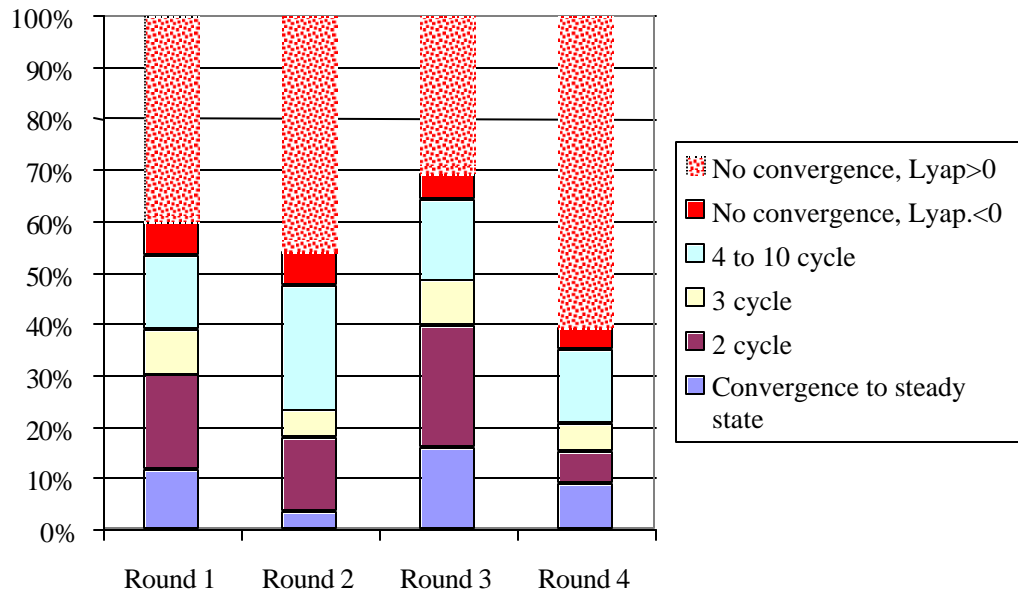


Figure 4: Long term convergence to a steady state or a cycle. Based upon 620 simulations per round. The classification of steady states and cycles is based upon periods 951-1000. Two

different kinds of nonconvergence simulations are distinguished based upon the estimated largest Lyapunov exponent (1000 periods were used to estimate the Lyapunov exponent).

The case of non-convergence and possible occurrence of chaos and strange attractors is of special interest. In order to investigate whether the price fluctuations are chaotic, the well-known Wolf algorithm (Wolf et al. (1985)) was used to estimate the largest Lyapunov exponent (using prices in periods 1-1000)¹². A positive Lyapunov exponent implies that the system exhibits sensitive dependence upon initial conditions and is *chaotic*. We find that almost 90% of the non-convergent price series have a positive Lyapunov exponent.¹³ It is important to note that each simulation, given the 6 predictions strategies, is a completely deterministic system without any influences of external noise. The fact that in the final round more than 60% of all simulations yield a positive Lyapunov exponent may thus be interpreted as strong evidence for chaos in our strategy experiment.

In summary, one may characterize these results by saying that the forecasting errors decrease significantly over the rounds and that prices converge to some neighborhood of the RE steady state, while at the same time the price fluctuations become more complicated and the fraction of chaotic price sequences increases.

3.3 What causes chaos?

The fact that so many price sequences are chaotic raises the question what exactly causes this chaotic behavior. First we focus on the role of heterogeneity. Next we study whether there are *specific* strategies that are the ‘rotten apples’ preventing convergence (to a steady state or a cycle). Finally we look whether continuity of the strategies plays a role.

¹² In applying the Wolf algorithm several parameters have to be selected, such as the embedding dimension, the maximum allowable distance between initial points and the separation time. In the results presented below we used an embedding dimension of 3, a maximum allowable distance of 0.5 and a separation time of 4, which are in the order of magnitude of what is commonly used; see the discussion in Wolf et al. (1985). For other values of these algorithm parameter values, similar results were obtained, and in particular the fraction of positive largest Lyapunov exponents was roughly the same.

¹³ The remaining 10% typically has a slightly negative Lyapunov exponent close to 0, indicating quasi-periodic behaviour or periodic behaviour with long period.

Homogeneous vs. Heterogeneous agents

We study the behavior of the individual strategies in a representative agent framework. Simulations were run in which all 6 strategies in the market are the same (31 simulations were run per strategy, one for each C from 50 to 80). In the nonconvergent simulations Lyapunov exponents were calculated.

Figure 5 shows the results of these simulations. Compared to the simulations with heterogeneous agents, convergence to a steady state price occurs more often (around 30% versus 10%), whereas the percentage of nonconvergence simulations is relatively small

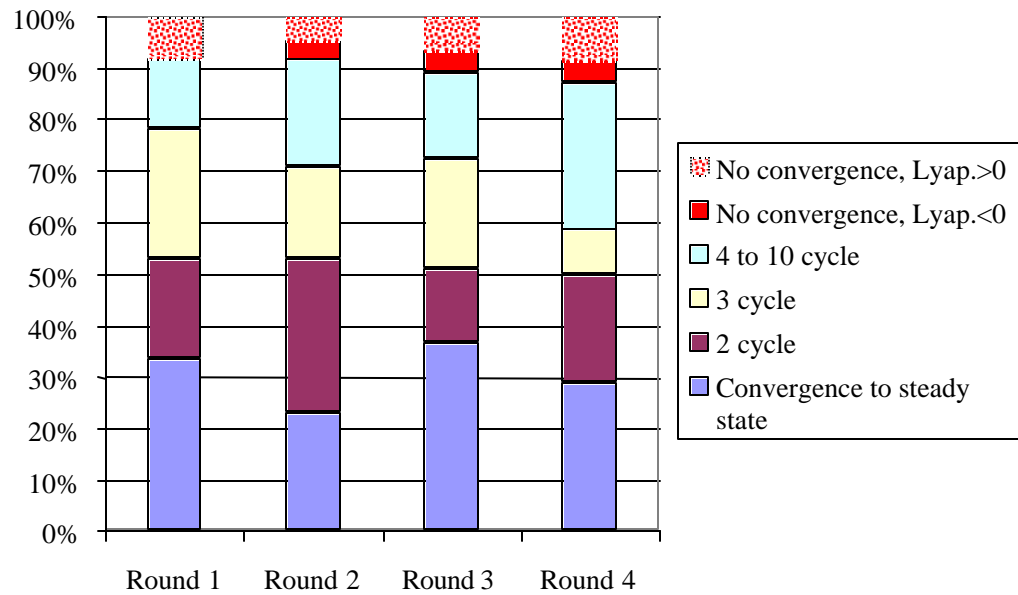


Figure 5: Long term convergence to a steady state or a cycle of simulations with individual strategies in the homogeneous agent case. Based upon 31 simulations per individual strategy ($C=50$ to 80) The classification of steady states and cycles is based upon periods 951-1000. Two different kinds of nonconvergent simulations are distinguished based upon the estimated largest Lyapunov exponent (1000 periods were used to estimate the Lyapunov exponent).

(approximately 10% versus 50% in the heterogeneous agents simulations). In more than 50% of the cases, an individual strategy leads to a stable cycle. No clear pattern of changes over the rounds is observed. Apparently, the same strategies that so often lead to chaotic price dynamics in a heterogeneous agent situation typically cause prices to converge either to a steady state or to a stable periodic cycle in a homogeneous agent situation.

Interestingly, for 99 of the 102 strategies the mean quadratic error in the homogeneous markets is larger than in the heterogeneous markets. Strategies differ much more in predicting quality

in the homogeneous markets than in the heterogeneous markets. For example, in round 1 the mean quadratic error of strategies in the homogeneous market is between 2635 and 86486 (SD=23230) and in the heterogeneous market between 3194 and 10964 (SD=2515). A similar result is found in the cobweb laboratory experiments of Hommes et al. 2000a where single-agent experiments yield a much higher mean quadratic forecasting error than multi-agent experiments. The mean variance of the prices is also much larger in the homogeneous markets (593, 736, 609, and 593 in rounds 1 to 4 respectively, versus 159, 180, 103, and 78 in the heterogeneous markets). In contrast to the 'close to the steady state chaos' of a typical heterogeneous market, many homogeneous agent cases are characterized by 'far from the steady state stable cycles.'¹⁴

Do specific strategies cause chaos?

Next we study whether there are *specific* strategies that are the 'rotten apples' that prevent convergence (to a steady state or a cycle). Prime candidates are the strategies that lead to nonconvergent price sequences in the homogeneous market. If we discard all heterogeneous situation simulations in which one of the participating strategies would not converge (to a steady state or a stable cycle) in a homogenous situation, the percentage of nonconvergence decreases only from 50% to 39%¹⁵. Apparently, chaos cannot be attributed (only) to these strategies. In other words, even if all participating strategies are stable (in the sense that in the homogeneous situation prices would converge to a steady state or a cycle), the *interaction* of strategies leads to an unstable, non-converging outcome in almost 40% of the cases.

Another way to look for the 'rotten apples' is to determine how often the prices converge to a steady state, a cycle, or not at all when a specific strategy participates in the market, and to compare these numbers with the overall results. A strategy is defined as a '*stabilizer*' (a '*destabilizer*') if the price sequences of markets in which this strategy participates converge more (less) often to a steady state and does less (more) often not converge at all¹⁶. In each round we find 2 destabilizers and between 1 and 3 stabilizers. After removing the destabilizers from the simulations, we find (of course, almost by construction) more convergence to a steady state (14.4% instead of 10.1%) and less non-convergence (38% instead of 50%), but these differences are not spectacular¹⁷.

¹⁴ Note that the strategies, as designed by the subjects, are intended for use in heterogeneous markets, which may (partly) explain the lower quality of the forecasts in the corresponding homogeneous markets.

¹⁵ The percentages of nonconvergence decrease in round 1 to 4 from 46.5 to 33.7, from 52.4 to 37.1, from 35.6 to 33.2, and from 65.0 to 54.7 respectively. These numbers are based upon the 1000 periods simulations.

¹⁶ Based upon a 5% statistical significance (binomial tests).

¹⁷ New 1000-periods simulations were performed without the destabilizers. The percentages of nonconvergence decrease in round 1 to 4 from 46.5 to 29.0, from 52.4 to 47.4, from 35.6 to 27.1, and from 65.0 to 48.5, respectively.

Although the destabilizers apparently are not the (only) source of non-convergence and chaos, it may be interesting to take a closer look at some characteristics of these strategies. First, the destabilizers are typically stable in the homogeneous situations,¹⁸ which shows that it is the *interaction* with other strategies that causes the destabilizing force of these strategies. Second, the destabilizers are not significantly less successful (in the heterogeneous markets) than the other strategies (the average ranking is somewhat under the median, at the 67th percentile)¹⁹. Third, returning to the general description of the strategies at the beginning of this section, we find that 7 of the 8 destabilizers are not continuous, while 7 of the 8 stabilizers are continuous. This leads to the more general question of which characteristics of the strategies are important for the price dynamics.

Characteristics of individual strategies and price dynamics

Over all, half of the strategies are not continuous (see table 1). Continuity of a strategy appears to be unrelated to own average quadratic prediction errors and also unrelated to the average quadratic prediction errors of the other market participants. This last point means that a discontinuous strategy in a market does not make prediction harder for the other participants. Most discontinuous strategies do not seem to be strange or unreasonable. The discontinuity arises because the strategy tries to distinguish different situations that call for a different kind of prediction rule, and the discontinuity occurs only at the border of different situations. A good example is strategy 9 from round 4 (one of the destabilizers in that round):

$$\begin{aligned}
 t=1: & \quad P_c(1)=55, \\
 t=2..4: & \quad P_c(t)=0.5*P_c(t-1)+0.5*P(t-1), \\
 t>4: & \quad \text{if } (|P(t-1)-P(t-3)|<10 \text{ and } |P(t-2)-P(t-4)|<10) \text{ then} \\
 & \quad P_c(t)=P(t-2) \text{ else } P_c(t)=0.5*P_c(t-1)+0.5*P(t-1)
 \end{aligned}$$

This strategy checks for two-cycles and is adaptive otherwise. Prices in a homogeneous market converge to a four-cycle, but in the heterogeneous simulations, prices do not converge in 78% of the simulations when this strategy is one of the participants (compare with 65% in all simulations in round 4). Apparently, the interaction of this strategy with some other strategies leads to instability.

It is an interesting exercise to redo the heterogeneous simulations with only continuous strategies. We find much more convergence to a steady state (41.5% instead of 10.1%), fewer

The percentages of convergence to a steady state increases in round 1 to 4 from 11.8 to 15.6, from 3.4 to 3.7, from 16.1 to 22.3, and from 9.2 to 16.1, respectively.

¹⁸ Two of the 8 destabilizers converge for all C's to a steady state in the homogeneous situation, one to a two-cycle, four to a four-cycle, with only one typically not converging.

¹⁹ The rankings of the destabilizers are in round 1, 7th and 17th out of 29, in round 2, 14th and last out of 28, in round 3, 12th and 18th out of 21 and in round 4, 16th and last out of 24.

cycles (only 23.5% instead of 40%), but still a considerable number of non-convergent price sequences (35% instead of 50%)²⁰. Again, we can only conclude that the main source of instability is *not* to be found in individual strategies, but in the *interaction* of different strategies.

4. Concluding remarks

Summary of the results.

The strategy method has proven to be a successful tool in studying expectation formation. It provides information about what kind of rules individuals use in forecasting prices and it enables an analysis of the stability and instability of dynamic economic systems with expectations feedback. Subjects use a wide variety of strategies and have a tendency to use more complicated strategies when they gain more experience. In our simulations with heterogeneous agents, only in about 10% of all cases the market settles down to the unique RE equilibrium. After round 2, the mean quadratic distance between the realized market price and the RE steady state decreases. However, at the same time, the complexity of the price fluctuations increases. In the final round, in more than 60% of the cases apparently chaotic price fluctuations around an unstable RE equilibrium price arise.

In order to investigate the causes of chaotic behavior, simulations of the same cobweb model with homogeneous, individual strategies are run. In these simulations, convergence to the RE equilibrium price occurs in roughly 30% of all cases and stable periodic price fluctuations in about 60% of the cases, whereas chaos only arises in 10% of the cases. Moreover, in the homogenous market simulations the mean squared forecasting errors and the variance of the price fluctuations are significantly higher than in the heterogeneous markets. Apparently, in a majority of cases the homogeneous markets converge to a *'far from the steady state stable cycle,'* whereas a heterogeneous market converges most likely to *'close to the steady state chaos.'*²¹ If we discard all heterogeneous situation simulations in which one of the participating strategies would not converge in a homogenous situation, the percentage of non-convergence decreases only from 50% to 39%. This means that, even if all participating strategies are stable in the homogeneous situation (i.e., prices converge to a steady state or a cycle), in the heterogeneous markets the *interaction* of these

²⁰ Large differences exist between rounds in these simulations. The results of the 3rd round, with only 11 continuous strategies, have a large impact on the overall results. In these simulations 84.4% of the price sequences converges to a steady state, only 4.4% to a cycle, and 11.3% does not converge.

²¹ 'Close' to the steady state has a relative meaning here, compared to the large amplitude cycles in atypical homogeneous market. Figure 2 for example shows that the mean quadratic distance to the RE price in the last 10

different strategies still leads to an unstable outcome in almost 40% of the cases. Finally, characteristics of the participating strategies, such as continuity, seem to influence the stability. However, no characteristic of the *individual* strategies seems crucial for the stability or instability of the market.

A boundedly rational heterogeneous agents equilibrium.

A tentative explanation of these results is as follows. The S-shaped supply curve (see figure 1) is characterized by high marginal supply in a neighborhood of the RE equilibrium price, which may cause (local) instability, and by low marginal supply far away from the RE equilibrium price, which may cause stable periodic motion. Consider, for example, the case where all strategies coincide with naive expectations or some other simple forecasting rule. An individual forecast below (above) the RE steady state leads to a realized market price far above (below) the RE steady state, and the system converges to a regular, large amplitude stable 2-cycle far from the RE steady state, with large and systematic forecasting errors. This can not be an equilibrium, however, since individual agents would have a strong incentive to improve their forecasts and the cyclic pattern is sufficiently regular to learn from their systematic mistakes and explore other, better prediction strategies. In a heterogeneous market individual forecasts will typically be distributed over an interval containing the RE steady state price leading to a realized market price not too far away from the steady state. In the early stages of a heterogeneous market this may lead to large amplitude (regular) price cycles around the steady state, but as agents become more experienced they should be able to reduce their forecasting errors to a reasonable level, thus pushing realized market prices closer to the RE steady state. However, as prices move closer to their steady state level, the cobweb system moves to the steep part of the supply curve and enters the local instability region. A decrease of the amplitude of the price fluctuations due to more experience is thus accompanied by an increase of local instability apparently leading to small amplitude chaotic price oscillations. This effect may be intensified by the increasing complexity of the submitted strategies over the rounds. As prices get closer to the steady state and start fluctuating chaotically, it becomes increasingly difficult for individuals to discover regularities in the observed patterns and to improve forecasts further. A heterogeneous market may thus end up in a boundedly rational chaotic equilibrium in a neighborhood of the (unstable) RE

periods of round 4 is about 50, showing that deviations of at least +7 or -7 to the RE steady state are no exception.

steady state, with agents using simple strategies with forecasting errors that are both of reasonable size and non-systematic, their structure hard to detect from time series observations.

The *interaction* of the different strategies together with the local instability of the steady state is the main source of the complicated price fluctuations in the case with heterogeneous agents. These results seem to be in line with the *Adaptive Rational Equilibrium Dynamics (ARED)* introduced in Brock and Hommes (1997). The ARED is an evolutionary competition, based upon predictive success, between simple predictions strategies that can lead to bifurcation routes to chaos and strange attractors.²² In the simplest of all dynamic economic expectations feedback system, the cobweb model, with a unique but locally unstable RE steady state and no exogenous random shocks, the interaction of boundedly rational agents can lead to instability and chaos in a neighborhood of the RE steady state.

Future perspective

Expectations and learning may play a key role and may affect asset prices significantly in speculative markets. In this paper, we have focussed on the simplest dynamic expectations feedback system, the textbook case of the hog-cycle model. Even in this simplest case having a unique RE solution, which is in theory stable under various learning schemes, our strategy experiment does *not* converge to RE. Instead our experiment is characterized by a boundedly rational heterogeneous agents equilibrium exhibiting excess volatility and moderate, but unpredictable price fluctuations. In future work, we plan to follow a similar experimental approach to dynamic speculative markets for risky assets. Such a dynamic framework is typically characterized by multiple RE solutions: a fundamental RE solution as well as rational bubble solutions. Whether in such a more complicated speculative market the fundamental RE solution can be learned or whether excess volatility will prevail is an important question for future experimental work.

²² Brock and Hommes (1997) focus their analysis mainly on the case with a sophisticated forecasting strategy, such as rational expectations which can be obtained at positive information costs, versus a simple strategy, such as naive expectations which is freely available. Hommes (2000) presents an example without any information costs where evolutionary competition between simple strategies, driven by the predictive success of the strategies, leads to chaotic price fluctuations in the cobweb economy. Brock and Hommes (1998) demonstrate the possibility of chaotic fluctuations in evolutionary competition without information costs in an asset pricing model.

References

- Allen, H. and Taylor, M.P., 1990. Charts, noise and fundamentals in the London foreign exchange market, *Economic Journal* 100 (Conference 1990), 49-59.
- Arifovic, J., 1994. Genetic algorithm learning and the cobweb model, *Journal of Economic Dynamics and Control* 18, 3-28.
- Arthur, W.B., Holland, J.H., LeBaron, B., Palmer, R. and Taylor, P., 1997, Asset pricing under endogenous expectations in an artificial stock market, In: Arthur, W.B., Lane, D. and Durlauf, S. (eds.), *The Economy as an Evolving Complex System II*, Addison-Wesley, Redwood City, CA, 15-44.
- Axelrod, R. 1984. *The Evolution of Cooperation*. New York: Basic Books.
- Brandts, J. and Schram, A. 2001. Cooperative gains or noise in public good experiments: applying the contribution function approach. *Journal of Public Economics*, 79, 399-427.
- Bray, M.M. and Savin, N.E., 1986. Rational expectations equilibria, learning and model specification, *Econometrica* 54, 1129-1160.
- Brock, W.A., and Hommes, C.H., 1997. A rational route to randomness, *Econometrica*, 65, 1059-1095.
- Brock, W.A. and Hommes, C.H. 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model, *Journal of Economic Dynamics and Control* 22, 1235-74.
- Chiarella, C., 1988. The Cobweb Model: Its Instability and the Onset of Chaos, *Economic Modelling* 5, 377-384.
- Cuthbertson, K., 1996. *Quantitative financial economics. Stocks, bonds and foreign exchange*, John Wiley & Sons, Chichester.
- DeLong, J.B., Schleifer, A., Summers, L.H. and Waldman, R.J., 1990. Noise trader risk in financial markets, *Journal of Political Economy* 98, 703-738.
- Ezekiel, M., 1938. The cobweb theorem, *Quarterly Journal of Economics* 52, 255-280.
- Frankel, J.A. and Froot, K.A., 1987. Using survey data to test standard propositions regarding exchange rate expectations, *American Economic Review* 77, 133-153.
- Fudenberg, D. and Levine, D.K., 1998. *The Theory of Learning in Games*, MIT Press, Cambridge.
- Grandmont, J.M., 1998. Expectations Formation and Stability in Large Socio-Economic Systems, *Econometrica*, 66, 741-781.
- Grandmont, J.-M., and Laroque, G., 1991. Economic dynamics with learning: some instability examples, In: *Equilibrium theory and applications, Proceedings of the sixth International Symposium in Economic Theory and Econometrics*, edited by Barnett et al. Cambridge University Press, Cambridge, 247-273.
- Hommes C. H., 1994. Dynamics of The Cobweb Model with Adaptive Expectations and Nonlinear Supply and Demand, *Journal of Economic Behavior and Organization* 24, 315-335.
- Hommes, C. H., 2000. Cobweb Dynamics under Bounded Rationality, In: Dockner, E.J. et al. (eds.) *Optimization, Dynamics and Economic Analysis. Essays in Honor of Gustav Feichtinger*, Physica-Verlag, Heidelberg, 134-150.

- Hommes, C.H., Sonnemans, J. and van de Velden, H. 2000a. Expectation formation in a cobweb economy: some one-person experiments, In: Delli Gatti, D., Gallegati, M. and Kirman, A., Interaction and market structure. Essays on heterogeneity in economics, Lecture Notes in Economics and Mathematical Systems 484, Springer Verlag, pp.253-266.
- Hommes, C.H., Sonnemans, J., Tuinstra, J. and van de Velden, H. 2000b. Expectations driven price volatility in an experimental cobweb economy, CeNDEF Working paper 99-07, University of Amsterdam.
- Hommes, C. H., and Sorger, G., 1998. Consistent Expectations Equilibria, *Macroeconomic Dynamics* 2, 287-321.
- Keser, C. 1992. Experimental Duopoly Markets with Demand Inertia, *Game-Playing Experiments and the Strategy Method*. Heidelberg: Springer Verlag.
- Kirman, A., 1993. Ants, rationality and recruitment, *Quarterly Journal of Economics* 108, 137-156.
- Lucas, R.E., 1971. Econometric testing of the natural rate hypothesis, In: *The econometrics of price determination Conference*, Board of Governors of the Federal Reserve System and Social Science Research Council, edited by O. Eckstein, 51-59.
- Lux, T., 1995. Herd behaviour, bubbles and crashes, *Economic Journal* 105, 881-896.
- Marimon, R., Spear, S.E. and Sunder, S., 1993. Expectationally driven market volatility: an experimental study, *Journal of Economic Theory* 61, 74-103.
- Marimon, R. and Sunder, S., 1993. Indeterminacy of equilibria in a hyperinflationary world: experimental evidence, *Econometrica* 61, 1073-1108.
- Muth, J.F., 1961. Rational expectations and the theory of price movements, *Econometrica* 29, 315-335.
- Nerlove, M., 1958. Adaptive expectations and cobweb phenomena, *Quarterly Journal of Economics*, 72, 227-240.
- Offerman, T., Potters, J. and Verbon, H. 2001. Cooperation in an overlapping generations experiment. *Games and Economic Behavior*, 264-275.
- Sargent, T.J., 1993. *Bounded Rationality in Macroeconomics*, Oxford: Clarendon Press.
- Sargent, T. J., 1999. *The Conquest of American Inflation*, Princeton: Princeton University Press.
- Selten, R., Mitzkewitz, M. and Uhlich, G. R. 1997. Duopoly strategies programmed by experienced players. *Econometrica* 65, 517-555.
- Shiller, R., 1989. *Market Volatility*, MIT Press, Cambridge.
- Shiller, R., 2000. Measuring bubble expectations and investor confidence. *Journal of Psychology and Financial Markets* 1, 49-60.
- Smith, V., Suchanek, G.L. and Williams, A.W., 1988. Bubbles, crashes and endogenous expectations in experimental spot asset markets, *Econometrica*, 56, 1119-1151.
- Sonnemans, J., 1998. Strategies of search. *Journal of Economic Behavior and Organization* 35, 309-332.
- Sonnemans, J., 2000. Decisions and strategies in a sequential search experiment. *Journal of Economic Psychology* 21, 91-102.

- Sunder, S., 1995. Experimental asset markets: a survey, In: Handbook of Experimental Economics, chapter 6, edited by Kagel, J. and A. Roth, Princeton University Press, 445-500.
- Topol, R., 1991. Bubbles and volatility of stock prices: effects of mimetic contagion, Economic Journal 101, 786-800.
- Wolf, A., J.B. Swift, H.L. Swinney and J.A. Vastano 1985. Determining Lyapunov exponents from a time series. Physica D 16, 285-317.

Appendix 1: Instructions for formulating a strategy (translated from Dutch)

Your strategy has to predict prices in a situation that is much like the experiment in which you participated. Therefore we first summarize the essential features of that situation.

The situation

In this experiment you are the adviser to a producer. The nature of the product that is being produced by this producer is not relevant in this experiment. At the start of each period you make a prediction of the price of the product in that period. The producer you are coupled with decides how much to produce, based upon your prediction of the price.

Several producers are active in one market. Every producer is coupled with exactly one adviser (participant in this experiment) and every adviser with exactly one producer. The realized price is determined by the total production of all producers in a market and the total consumer demand (the realized price is such that total supply equals total demand).

In this experiment all strategies have the role of adviser; a computer program plays the role of both producers and consumers.

After all predictions are collected, the computer calculates the realized market price. After that the next period starts.

Information

There is only a limited amount of information you (your strategy) can use. You do **NOT** know:

- The number of producers that are active in the market of your producer;
- The predictions of the other participants;
- How producers determine their production based upon your prediction;
- How the price is determined by total demand and supply.

You **DO** know the realized prices of the previous periods as well as how good your predictions have been in these periods.

The consumer demand and the way the production is determined by your prediction may differ between markets. Therefore realized prices may also differ considerably between markets. You can interpret this as follows: every simulation (of 20 periods) your strategy is coupled with a different producer (who may have a different technology) who is active on another market than the previous producer you were coupled with.

How to formulate a strategy

A strategy is a complete plan of action. If you would give your strategy to someone else, he or she should be able to make exactly the predictions that you yourself would have made.

Your strategy should comply with three requirements: your strategy should be complete, unambiguous and informational correct. The requirement of completeness means that your strategy should provide a prediction in all possible situations. The requirement of unambiguousness means that your strategy should provide exactly one prediction that is a real number between 0 and 100 in all possible situations. The requirement of informational correctness means that your strategy only uses information that is available at that moment.

Example of an incomplete strategy

“In the first period my prediction is 40. In the next periods my prediction is 60 if the previous price was larger than 50 and 40 if the previous price was lower than 50”. This strategy is not complete because it provides no prediction if the previous price was exactly 50.

Example of an ambiguous strategy

“In the first period my prediction is 70. In the next periods I will raise my prediction of the previous period with 10 if my previous prediction was lower than the realized price, I will lower my prediction with 10 if my previous prediction was higher than the previous price, and I will maintain my prediction if my prediction error in the previous period was less than 5.”

This strategy is ambiguous because it is unclear what the prediction should be if the previous prediction was (for example) 3 above the realized price: should the prediction be maintained or decreased by 10? By indicating which rule has priority this strategy can be made unambiguous.

Example of an informational incorrect strategy

“In the first period my prediction is 45. In the other periods my prediction depends on the price in period 5. If the price in period 5 was larger than 40 I predict 30 and otherwise I predict 70”

This (rather strange) strategy is informational incorrect because in period 2 through 5 it is unknown what the price in period 5 will be.

In the examples above the strategies are described in words. However, it is easier to check whether a strategy complies to all requirements if everybody uses the same notation.

The present period is indicated as t .

The information you can use in period t are the realized prices of the previous periods $P(i)$, $1 \leq i \leq t-1$, and the predictions of your strategy in the previous periods $V(i)$, $1 \leq i \leq t-1$.

A strategy consists of two columns. In the first column you put the periods in which that part of your strategy is valid, in column 2 you put your strategy for these periods. You can use conditions when describing your strategy, as in the example of a (not necessarily successful) strategy below.

Periods:	Prediction:
$t=1$	$V(t)=70$
$t=2$	$V(t)=50$
$t>2$	If $ V(t-1)-P(t-1) < 10$ then $V(t)=P(t-1)$ else $V(t)=(P(t-2)+V(t-1))/2$

Explanation

The prediction of the first period should be a number between 0 and 100 (because no information is yet available). In this example a number is also given in period 2.

In periods 3 to 20 this strategy first checks whether the absolute value of the error of the previous period ($|V(t-1)-P(t-1)|$) was less than 10. If that is the case, the prediction is the previous period's price. If that is not the case the prediction is the mean of the price of two periods ago and the previous prediction.

Notation of more complicated strategies

If you want to construct more complicated conditional strategies, you have to use brackets. For example the strategy "If $|V(t-1)-P(t-1)| < 10$ then (If $P(t-1) > 60$ then $V(t)=P(t-1)$ else $V(t)=V(t-1)$) else $V(t)=(P(t-2)+V(t-1))/2$ ". In this strategy the prediction depends on the previous price, if the absolute error in the previous period was less than 10. If you have experience with programming in Pascal or Basic, you may also use the regular IF-THEN-ELSE statements of Basic or Pascal.

You may use all usual mathematical notations you need (like S). If you are not sure whether your strategy will be clear for our programmer you should tell us when you hand in your strategy, and we will check the strategy immediately.

How to check a strategy

Check the left column. Does the strategy predict a price in all periods? If not, your strategy is not complete.

Check the right column. For each cell in this column (each sub-strategy) you should check the completeness and unambiguousness: is exactly one prediction generated in each possible situation? Check also the information that is used in each column: is this information indeed available? In the example above, the sub-strategy in the bottom right cell uses the price of two periods ago, and such a strategy can only work from period 3 onwards.

The computer simulations

All submitted strategies will be programmed, and several thousands of simulations will be run. Each simulation starts with the random draw of some strategies; these strategies will form a market for 20 periods. Next some random parameters will be drawn that determine the demand and production curves. The market is run for 20 periods, and for each participating strategy the quadratic prediction errors are calculated. After thousands of simulations for each strategy the mean quadratic prediction error is calculated and a ranking is made. At the top of this list is the strategy with the smallest mean quadratic error, and strategies below have an increasing mean quadratic error.

Information about the simulations

As soon as the simulations are run, the ranking list will be made public on the website of CeNDEF. On this list all strategies are identified with a personal code.

A printed version of the ranking list will be distributed at the Monday classes. All participants will then also receive a printout of the results of 5 simulations in which their strategy participated. These 5 simulations are randomly chosen.

Appendix 2: example of a strategy form (translated from Dutch)

Strategy Form

Your personal code:

Your strategy:

Period:	Prediction:
t=1	V(t)=....

Please use the following notation:

t period number

$P(t)$ realized price in period t

$V(t)$ your prediction in period t

The information you can use in period t are the realized prices of the previous periods $P(i)$, $1 \leq i \leq t-1$ and the predictions of your strategy in the previous periods $V(i)$, $1 \leq i \leq t-1$.

We will do the best we can, but in case you later find out that the programmer did not program your strategy the way you meant it, there is nothing we can do about it. The results of the simulations are final. Therefore, be sure to make exactly clear what you want your strategy to be and please **write legible**.

Don't forget to fill in the questionnaire!

Appendix 3: example of feedback (translated from Dutch)

This is the feedback subject 1 (who used an adaptive strategy) received after round 2

Results Strategy Experiment Round 2

Below you find the ranking of the second round. The table displays for each student the mean quadratic error over 20 periods. The student with **code 14** has won the fifty guilders of this round!

Strategy	Mean Quadratic Error
214	2946.62
215	3876.26
227	4347.99
201	5077.21
213	5459.69
210	5718.21
216	5872.36
205	5909.91
209	6953.06
206	7620.94
229	7699.34
224	7752.06
211	7937.15
226	8253.44
223	8287.59
218	8377.60
225	8532.96
204	8615.77
221	8891.20
203	9532.42
219	9866.34
207	10216.90
228	10486.71
212	11126.55
208	14791.05
222	16355.44
202	17282.80
220	20784.03

On the next two pages you will find the results of 5 (randomly chosen) simulations. This may help you to get an impression about the situations in which your strategy performed well or poorly.

You can hand in your third strategy during the classes at Thursday February 4 and Monday February 8. We will then check your strategy immediately. Please also hand in the short questionnaire. Only if you hand in both a strategy and a questionnaire you will earn the 5 guilders fee.

You can also find this ranking and the ranking of the next rounds on internet: www.fee.uva.nl/cendef/

Strategy: 201

Period	Prediction	Realized Price	Quadratic Error
1	50.00	46.77	10.41
2	49.68	60.04	107.33
3	50.71	33.10	310.15
4	48.95	66.93	323.14
5	51.20	16.78	1184.83
6	46.90	71.06	583.97
7	49.92	27.37	508.53
8	47.10	62.90	249.62
9	49.07	34.38	215.94
10	47.24	65.40	330.07
11	49.51	21.49	784.95
12	46.01	70.10	580.65
13	49.02	24.19	616.45
14	45.91	65.94	400.99
15	48.42	30.44	323.00
16	46.17	65.70	381.28
17	48.61	25.25	545.84
18	45.69	68.81	534.58
19	48.58	23.66	620.98
20	45.47	67.43	482.37
			9095.11

Strategy: 201

Period	Prediction	Realized Price	Quadratic Error
1	50.00	96.66	2176.73
2	54.67	82.84	794.03
3	57.48	39.31	330.33
4	55.67	84.41	826.15
5	59.26	92.79	1124.14
6	63.45	48.94	210.65
7	61.64	80.47	354.73
8	63.99	93.56	874.42
9	67.69	52.10	242.95
10	65.74	73.28	56.95
11	66.68	94.30	762.80
12	70.13	60.78	87.40
13	68.96	54.15	219.37
14	67.11	93.58	700.40
15	70.42	74.83	19.43
16	70.97	51.10	394.99
17	68.49	93.11	606.15
18	71.57	88.36	282.04
19	73.67	45.85	773.68
20	70.19	85.69	240.41
			11077.76

Strategy: 201

Period	Prediction	Realized Price	Quadratic Error
1	50.00	93.00	1848.94
2	54.30	73.11	353.63
3	56.18	70.57	207.09
4	57.62	69.63	144.23
5	59.12	80.61	461.70
6	61.81	60.73	1.16
7	61.67	87.98	691.98
8	64.96	38.77	685.85
9	61.69	81.28	384.08

10	64.14	78.52	206.81
11	65.93	71.68	32.96
12	66.65	67.53	0.77
13	66.76	80.77	196.38
14	68.51	55.41	171.79
15	66.87	91.05	584.61
16	69.90	36.82	1093.82
17	65.76	81.02	232.77
18	67.67	80.04	153.01
19	69.22	63.75	29.92
20	68.53	78.48	98.91
			7580.40

Strategy: 201

Period	Prediction	Realized Price	Quadratic Error
1	50.00	76.37	695.12
2	52.64	62.26	92.57
3	53.60	63.62	100.42
4	54.60	81.52	724.46
5	57.97	39.44	343.02
6	55.65	84.79	848.89
7	59.29	60.07	0.60
8	59.39	63.66	18.21
9	59.92	85.03	630.52
10	63.06	36.82	688.41
11	59.78	82.76	527.90
12	62.65	63.41	0.57
13	62.75	61.64	1.23
14	62.61	85.00	501.56
15	65.41	38.45	726.50
16	62.04	81.65	384.41
17	64.49	63.40	1.18
18	64.35	61.84	6.30
19	64.04	84.51	419.04
20	66.60	37.75	832.43
			7543.34

Strategy: 201

Period	Prediction	Realized Price	Quadratic Error
1	50.00	83.39	1115.13
2	53.34	71.13	316.53
3	55.12	42.39	161.94
4	53.85	72.43	345.25
5	56.17	68.14	143.25
6	57.66	62.56	23.93
7	58.28	69.25	120.53
8	59.65	52.17	55.99
9	58.71	76.86	329.35
10	60.98	55.22	33.17
11	60.26	68.87	74.03
12	61.34	61.71	0.14
13	61.38	64.87	12.18
14	61.82	60.85	0.94
15	61.70	70.44	76.48
16	62.79	53.29	90.38
17	61.60	74.70	171.44
18	63.24	54.14	82.89
19	62.10	69.65	57.03
20	63.05	60.27	7.73

3218.32