UvA-DARE (Digital Academic Repository)

A strategy experiment in dynamic asset pricing

Hommes, C.H.; Sonnemans, J.; Tuinstra, J.; van de Velden, H.

Publication date
2002

[Link to publication](#)

Citation for published version (APA):

Hommes, C. H., Sonnemans, J., Tuinstra, J., & van de Velden, H. (2002). *A strategy experiment in dynamic asset pricing*. EPRINTS.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

A strategy experiment in dynamic asset pricing

Cars Hommes^a, Joep Sonnemans^{a,b}, Jan Tuinstra^a and Henk van de Velden^a

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

^b Center for Research in Experimental Economics and political Decision-making (CREED)

Department of Economics and Econometrics

University of Amsterdam

Roetersstraat 11, 1018WB Amsterdam

The Netherlands

<http://www.fee.uva.nl/cendef/>

Version November 2002

Abstract

This study presents a strategy experiment in asset pricing. In a simple dynamic asset pricing model the price in the present period is determined by the expectations of next period's price. After participating in an introductory laboratory experiment on expectation formation participants formulate a complete forecasting strategy. These strategies are programmed and markets are simulated. Participants receive feedback from the results of these simulations and can adapt their strategy. Four rounds are played. A final laboratory experiment compares predictions of participants with the predictions of the submitted strategy.

We find that most of the participants submit complicated strategies and that strategies become more complicated over the rounds. Most markets converge to a steady state price only after many periods, if at all. The number of converging price sequences increases over the rounds. These results suggest in general slow convergence and learning of the subjects over the rounds. Even in a stationary environment it turns out to be difficult to learn the correct fundamental price level. An important part of the non-convergence seems to be caused not by individual strategies but by the interaction of several strategies together. From the final experiment we conclude that the strategies are a good representation of what participants do in a laboratory experiment.

JEL classification codes: C91, C92, D84, G12, G14

1. Introduction

Asset markets can be considered as expectation feedback systems. Investors and traders form expectations about future prices and they trade based upon these expectations. The prices realized by these trades are public information and serve as an input for a new expectation formation process. The relation between (average) expectations and realized prices is positive: if a majority of investors expect prices to rise (decrease), the demand will increase and there will be an upward (downward) pressure on prices. Insights in these expectation formation processes and in the dynamics caused by these processes are pivotal for an understanding of the working of asset markets.

In the last decade several behavioral asset market models with boundedly rational traders having heterogeneous expectations about future asset prices have been introduced, e.g. Arthur et al. (1997), Brock and Hommes (1998), Gaunersdorfer (2000), LeBaron (2000), Lux and Marchesi (1999) among others. In these models two important classes of traders can be distinguished, namely fundamentalists and trend followers. Fundamentalists believe that asset prices converge to their fundamental value, given by the expected discounted sum of all future dividends. In contrast, trend followers believe that prices may deviate from their fundamental value and expect that prices move in trends. They extrapolate past price trends in forecasting future prices. Although theoretically appealing, little work has been done to investigate whether real investors actually use such behavioral rules as fundamentalist and trend following rules in forecasting prices.

In the real world data on prices are widely available but, unfortunately, reliable data on expectations or forecasting are rare. Economic experiments can be very useful in this respect. In Hommes et al (2002a, 2002b) we describe a laboratory environment in which subjects predict the next price in a standard asset market model. In these experiments participants act as advisors to a pension fund. The task of the advisors is to give an accurate prediction of the price of a risky asset. The pension fund bases its investment decision with respect to the risky asset upon the price forecast of the advisor/participant. The market price of the risky asset is determined by equilibrium of demand and supply. The length of these experiments is 50 periods; the data is well suited to study the short-term dynamics. Prices typically do not converge to the fundamental rational expectation (RE) price. To study long-term dynamics, the expectations strategies of the participants have to be interpreted (e.g. estimated as an AR (2) model) and simulated (see also Sterman 1989). This procedure is not very satisfactory because the interpretation of the researcher will probably not exactly coincide with the actual strategy of the subject. Another disadvantage is that the subjects get only limited experience, they would probably learn to predict differently if they could participate in more markets.

In the present experiment we solved interpretation problems by letting the subjects provide a complete forecasting strategy. A strategy has to be a complete representation of the subject's prediction in all possible states. The strategies are programmed and 1296 markets are simulated. Each market is in fact a deterministic dynamical system. Participants were financially motivated to submit the best strategy possible. This procedure is not new, these so called “strategy experiments” have become quite popular in experimental economics (for example, see Selten, Mitzkewitz and Uhlich (1997), Keser (1992), Offerman, Potters and Verbon (2001), Sonnemans (1998); and also the classic work of Axelrod (1984)).

Four rounds are played and at the end of each round the subjects receive feedback about the performance of their strategy. The feedback for each participant consists of the price and the prediction of the strategy in 5 randomly selected simulations of fifty periods. The ranking of the strategies as well as a programmed version of their own strategy are also included in the feedback. The ranking of the strategies is based on the average quadratic prediction error. Participants have to submit a revised strategy for the next round within a week.

After the subjects submit their fourth round strategies but before they receive the results of the fourth round we conduct a final laboratory experiment. The goal of this final experiment is to study the relationship between the actual behavior of subjects and the strategy they submit. We inform the subjects that they are in a market with five strategies from the third round. Since the subjects, in the final experiment, have information up till round 3 we can compare the prediction of the participant in the final experiment with the prediction of their fourth round strategy.

The main research questions we try to answer in this paper are: (1) what kind of strategies do subjects use? (2) Does the market price eventually converge to the fundamental price? (3) How does learning affect the price dynamics in the consecutive rounds? (4) How do individual strategies attribute to market stability or instability? (5) Is the submitted strategy an accurate description of the decisions subjects make in a regular experimental environment?

The paper is organized as follows. Section 2 describes the asset-pricing model. In section 3 the experimental design is explained. In section 4 we study the similarities and differences between the submitted strategies, the short, medium and long run price dynamics and the results of the final experiment. We summarize and conclude in section 5.

2. An asset-pricing model

In the experiment we used a simple dynamic asset-pricing model with heterogeneous beliefs (as in Brock and Hommes 1998, Gaunersdorfer 2000). Consider an asset market with H different trader

types or strategies, indexed with h . Traders can choose between a risk free asset that pays a fixed return r or a risky asset that pays an uncertain dividend y_t in period t . Let p_t denote the price of the risky asset in period t . Dividends are assumed to be independently and identically distributed (IID) with mean $\bar{y}$ and variance s_y^2 . Denote with z_{ht} the number of shares of the risky asset purchased by trader h in period t and let $R=1+r$. The realized wealth for trader type h in period $t+1$ then is

$$W_{h,t+1} = RW_{h,t} + (p_{t+1} + y_{t+1} - Rp_t)z_{ht}.$$

Traders subjective beliefs about the evolution of their wealth are characterized by their subjective conditional mean E_{ht} and their subjective conditional variance V_{ht} . Traders are mean-variance optimizers and their demand for shares corresponds to the solution of

$$\max_{z_{ht}} \left\{ E_{ht}(W_{h,t+1}) - \frac{1}{2} a V_{ht}(W_{h,t+1}) \right\} \quad (1)$$

where a measures the degree of risk aversion (assumed to be the same and constant for all traders). We assume $V_{ht}(W_{h,t+1}) = V_{ht}(p_{t+1} + y_{t+1} - Rp_t) = s^2$ for all h , that is, traders believe the variance of excess returns to be (the same) constant. The solution of (1) gives the demand for the risky asset by trader type h

$$z_{ht} = \frac{E_{ht}(p_{t+1} + y_{t+1} - Rp_t)}{a s^2}. \quad (2)$$

The supply of stocks z^s is fixed, and without loss of generality we assume that outside supply of shares is zero. In that case (2) can be considered as the excess demand for trader type h . The market equilibrium condition is

$$\sum_{h=1}^H z_{ht} = \frac{1}{a s^2} \sum_{h=1}^H E_{ht}(p_{t+1} + y_{t+1} - Rp_t) = 0.$$

Solving for the market equilibrium, price then yields

$$p_t = \frac{1}{RH} \sum_{h=1}^H E_{h,t}(p_{t+1} + y_{t+1}). \quad (3)$$

Because it is common knowledge that dividends are IID with $\bar{y}$ in the experiment, the market equilibrium price simplifies to

$$p_t = \frac{1}{R} \left(\bar{y} + \frac{1}{H} \sum_{h=1}^H E_{h,t}(p_{t+1}) \right). \quad (4)$$

Equation (4) is the basic equation of the asset-pricing model. The asset price in period t depends on the traders' subjective expectations of the price in period $t+1$. Notice that the traders, when forming their expectation about p_{t+1} , only know the prices p_{t-1} , p_{t-2} , etc. A key feature of the expectation feedback structure is its self-fulfilling character: if all traders predict high (low) future prices the

market equilibrium price will also be high (low). This is a well-known feature of speculative asset markets.

It is well known that, when expectations are homogeneous and all traders have rational expectations (with speculative bubble solutions excluded by the so-called transversality condition), the *rational expectations* or *fundamental* price p_t^f is given by the discounted sum of expected future dividends, that is,

$$p_t^f = \sum_{k=1}^{\infty} \frac{E_t(y_{t+k})}{R^k}.$$

Given that the dividend process is IID with mean $\bar{y}$ the *fundamental* price p^f is constant and given by

$$p^f = \frac{\bar{y}}{R-1} = \frac{\bar{y}}{r}.$$

If all agents predict the fundamental price, this belief is self-fulfilling and the realized price will be constant and equal to p^f .

3. Experimental design and procedures

The experiment lasted for eight weeks. The subjects were recruited from a course “Dynamical Systems”, a mathematical introduction to dynamical systems in the undergraduate econometrics program, and from the course “Micro Economics”, a course in the undergraduate economics program. Participation was not a course requirement. Students had no prior knowledge about dynamic economic systems and the asset-pricing model is not taught in the courses.

Introductory experiment

The goal of the introductory experiment was to give the subjects some experience with their forecasting task. The experiment was completely computerized and took place in the CREED experimental laboratory. The participants understanding of the instructions was checked by control questions. The forecasting task was presented as follows (this is a summary, complete instructions are available from the authors):

- You are a **financial advisor** to a pension fund that wants to optimally invest a large amount of money. The pension fund has two investment options: a risk free investment and a risky investment. In each time period the pension fund has to decide which fraction of its money to put on the bank account and which fraction of the money to spend on buying stocks. In order to make an optimal investment decision the pension fund needs an accurate prediction

of the price p_t of the stocks. As their financial advisor, you have to predict the stock market price p_t (in guilder) during 51 subsequent time periods.

- **Information about the stock market**

The price p_t of the stocks is determined by market equilibrium, that is, the stock market price p_t in period t will be the price for which aggregate demand equals supply. The supply of stocks is fixed during the experiment. The demand for stocks is mainly determined by the aggregate demand of a number of different pension funds active in the stock market.

- **Information about the investment strategies of the pension funds**

The precise investment strategy of the pension fund that you are advising and the investment strategies of the other pension funds are unknown. The return in the stock market per time period is uncertain and depends upon (unknown) dividend payments. As the financial advisor of a pension fund you are **not** asked to forecast dividends, but you are only asked to forecast the price of the stock in each time period. Based upon your stock market price forecast, your pension fund will make an optimal investment decision. The higher your price forecast the larger will be the fraction of money invested by your pension fund in the stock market, so the larger will be its demand for stocks.

- **Forecasting task of the financial advisor**

The only task of the financial advisors in this experiment is to forecast the stock market index p_t in each time period as accurate as possible. The price of the stock will always be between 0 and 100 guilders in the first two periods. The stock price has to be predicted **two** time periods ahead. After all participants have given their predictions for the first two periods, the stock market price p_1 in the first period will be revealed and based upon your forecasting error $p_1 - p_1^e$ your earnings for period 1 will be given.

- **Summary of information**

- ❖ Interest rate r and the mean dividend $\bar{y}$
- ❖ Past *prices* up to period $t-2$: $\{p_{t-2}, p_{t-3}, \dots, p_1\}$
- ❖ Past *predictions* up to period $t-1$: $\{p_{t-1}^e, p_{t-2}^e, \dots, p_1^e\}$
- ❖ Past *earnings* up to period $t-2$

- **Earnings**

Earnings will depend upon forecasting accuracy only. The better you predict the stock market price in each period, the higher your aggregate earnings. Earnings will be according to the following earnings table. If you are the best advisor in the experiment, you can earn an **additional bonus of 50 guilders**. This bonus will be given to the participant with the smallest average prediction error during the experiment.

Given the six individual predictions, the realized price was determined by the market equilibrium equation (4) with $H=6$. Subjects participated sequentially in two different markets of 50 periods with different composition of subjects and different mean dividend $\bar{y}$ and risk free interest rate r . At the beginning of the markets the exact values of the mean dividend and interest rate were announced, in the first market $\bar{y} = 3$ and $r=5\%$, and in the second market $\bar{y} = 2.4$ and $r = 6\%$. The corresponding fundamental prices for the two markets are thus given by: $p^f = 3/0.05 = 60$ and $p^f = 2.4/0.06 = 40$. In both markets participants could earn 1300 points each period. The earnings were negatively related to the size of the quadratic forecasting error. The number of points earned in period t by participant h is given by the quadratic scoring rule

$$e_{h,t} = \max \left\{ 1300 - \frac{1300}{49} (p_t - p_{ht}^e)^2, 0 \right\}$$

where 1300 points is equivalent to 0.65 Dutch guilders. Notice that earnings in period t were zero when $|p_t - p_{ht}^e| = 7$. The average earnings of the 21 participating subjects, in approximately 1.5 hours, were 31 Dutch guilders (14 Euro), where the maximum total earnings are 65 Dutch guilders (30 Euro).

Strategy experiment

The strategy experiment consisted of four rounds. In each round the subjects had to submit a strategy. Participants could submit their strategy anytime before a biweekly deadline. The number of participants decreased over the rounds. In round 1 all 21 participants submitted a strategy and in rounds 2, 3 and 4 respectively 19, 17 and 16 participants submitted strategies¹. The first round started at the end of the introductory experiment when all participants had to submit their first strategy. The participants were given a folder with all necessary information about the procedures and the forecasting task (available upon request). The experimenters checked these strategies for clarity, completeness (whether the strategy provides a prediction in all possible situations), uniqueness (whether the strategy always provides exactly one prediction) and informational correctness (whether the strategy does not use information that is not available, such as the price p_{t-1} when forecasting price p_t , future prices or previous predictions of other strategies). Every participant had his own personal code and we used this code to identify the subjects over the rounds. The participants did not know the personal codes of other participants.

In every round we simulated the experimental asset market 1296 times. Each market consisted of six different, randomly drawn, strategies. For the simulations we used 9 different values of the mean dividend ($\bar{y} = 2, 2.25, 2.5, \dots, 4$) and 9 different values of the interest rate ($r = 4\%, 4.25\%, 4.5\%, \dots, 6\%$). The fundamental price was therefore always in the interval $[33 \frac{1}{3}, 100]$. Every combination of mean dividend and interest rate was simulated 16 times, resulting in a total of 1296 simulations. The average number of simulations for a particular strategy is then around 400. The subjects were informed that the exact values of the mean dividend and the interest rate would not be the same for every simulation but would be available to the participating strategies at the beginning of each simulation. At the end of each round the subjects received private information about how their strategy predicted in the simulations. The private information

---

¹ Submitting a strategy seems not to be related to the ranking in the previous round. The participants who did not submit a new strategy had an average percentile score of 55% (a little better than average).

consisted of a page with the programmed version of their strategy (programmed in Borland Pascal) and five pages containing for each subject five randomly chosen simulations in which his or her strategy had participated. They also received public information, namely information given to all participants, about the ranking of the strategies by mean quadratic forecasting error. One week after participants received this information a new strategy had to be submitted.

We recognized that a similar incentive structure as in the introductory experiment might stimulate the subjects in the strategy experiment to cooperate and share their strategies and private information. For example, all subjects benefit from a fast convergence to the fundamental price. In the strategy experiment we therefore employed a tournament incentive structure: payment was based upon relative performance (average squared prediction error of the strategy). The participant with the strategy with the smallest average quadratic forecasting error received 50 guilders (approximately 22.70 Euro) in rounds 1, 2 and 3. 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) for each strategy they submitted.

A possible disadvantage of payments based upon relative performance is that subjects may try to destabilize markets in order 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 fundamental price except for one strategy that tries to destabilize the market by predicting a higher (lower) price, the realized price will be higher (lower) than the fundamental price but since the price is determined by the average of all predictions the price will always be closer to the fundamental price than to the prediction of the strategy that tries to destabilize the market. A single destabilizing strategy will thus end up with a larger quadratic prediction 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 questionnaires² or in class after the experiment ended) that he or she had tried to destabilize markets.

Final experiment

After subjects submitted their final, fourth round strategy, but before they received the final results we conducted another laboratory experiment. The main goal of this experiment is to investigate whether the strategies are a good representation of the actual behavior of the subjects in a

---

² Subjects filled in a small questionnaire every round. In this questionnaire 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 well they thought their new strategy would perform.

laboratory experiment. In the final experiment the subjects were informed that they were in a market with five strategies from the third round. Since the participants had already submitted their fourth round strategy but did not receive the results of the fourth round simulations yet, we can compare their fourth round strategy with their behavior in the final experiment³. For every participant we conducted two sequential markets of 50 periods. A market consisted of one participant and five third round strategies that were randomly drawn from the submitted third round strategies. The fundamental prices in the two markets differed but were the same for all participants. For the final experiment we also used the tournament structure for the earnings. The winner, that is, the subject with the lowest average quadratic error over the two markets, earned 100 guilders, the runner up earned 95 guilders etc.

4. Results

This section reports the results of the strategy experiment. We first look at the specific characteristics of the strategies that are submitted. After that we will investigate the short run price dynamics in section 4.2, the medium run price dynamics in section 4.3 and the long run price dynamics in section 4.4. In section 4.5 we investigate the behavior of particular strategies in the simulations and in particular the effect of individual strategies on the (non-) convergence of the market price. Finally in section 4.6 we investigate whether the submitted strategies are a good description of the behavior of the participants in an experimental setting by looking at the final experiment.

4.1 What kind of strategies do the participants use?

In most studies on economic dynamics researchers make assumptions about the expectations of the agents in the model. Usually the agents are assumed to be rational or boundedly rational using a simple prediction rule. In this study we are interested in what kind of strategies/prediction rules subjects actually use. This section focuses on the general characteristics of the submitted strategies without describing every individual strategy in detail.

---

³ The participants were informed about how their third round strategy predicted in round 3. They have no information about how their fourth round strategy predicts in the fourth round. Therefore their predictions in the final experiment are based upon the same information as their fourth round strategy. The only difference is that in the experiment the participants are more flexible since they can update and change their prediction strategies during the final experiment.

	<i>Round 1</i>	<i>Round 2</i>	<i>Round 3</i>	<i>Round 4</i>	<i>Total</i>
Number of strategies	21	18	16	16	71
Continuous	8	5	3	3	19
Conditional (+conditional on error)	7 (+6)	7 (+7)	6 (+7)	6 (+7)	26 (+27)
Mean dividend included	14	12	12	13	51
Interest rate included	13	12	11	12	48
Includes prediction(s) of previous period(s)	7	9	8	8	32
Includes weighted average of previous prices	6	6	3	3	18
Includes latest observed price	21	18	16	16	71
Trend following	7	7	9	6	29
Mean number of code lines	19.71	26.35	30	26.31	25.59

Table 1: Characteristics of the strategies⁴.

Table 1 summarizes the main characteristics of the strategies. In total 71 strategies are submitted. About one quarter of all submitted strategies is continuous. The strategies are continuous if a small change in the realized price will also result in a small change of the prediction. If a strategy is not continuous it is in all cases conditional. The conditions on which these strategies depend differ from past prices, past predictions, past prediction errors to the mean dividend or the interest rate. An example of a conditional strategy is “if the price in period $t-3$ is smaller than the price in period $t-2$ my prediction is ... if not my prediction is...”. About 75% of all strategies are conditional. In 51 strategies the mean dividend ($\bar{y}$) is included; in most of these cases the interest rate is then also included. In only seven strategies the fundamental price $p^* = \bar{y} / r$ is used for the prediction; in six of these strategies the fundamental price is only used in predicting the first and/or second price. Not a single participant had the fundamental price as its prediction for periods 3-50. All strategies include the last observed price. Almost 40% of all strategies are following a “trend”. An example is the winning strategy of round 4 (subject 19): $p_t^e = p_{t-2} + 2(p_{t-2} - p_{t-3})$. There are two strategies that try to detect a cycle, anticipating cyclic behavior. If the strategy detects a cycle, say of period 8, the strategy predicts the price eight periods back. The strategies get more complicated (measured by the total lines of program needed) until round 3; in round 4 there is a (slight) decrease of the complexity.

⁴ In rounds 2 and 3 we had to remove one strategy from the simulations (strategy of subject 9 and 21 respectively). These strategies often predict extreme values because of what seem to be mistakes by the participants (for example, strategy 9 in round 2 increases (decreases) the prediction when the last known price was lower (higher) than the prediction in that period). Both strategies perform extremely bad and have a large effect on the prices. We decided not to change the strategies ourselves, but to exclude these strategies from the simulations. Van de Velden (2001) contains the results when these strategies are included.

4.2 Short run dynamics

This section discusses predictions and price behavior in the first 50 periods (the *short run*). Because the incentives of the subjects and the ranking of the strategies are based upon the first 50 periods these periods are of special interest. The top panel of figure 1 shows the average absolute distance, $|p_t - p^f|$, between the realized market price p_t and the fundamental price p^f for the first 50 periods for each of the four rounds averaged over 1296 simulations. The average prices in the first two rounds are relatively close to the fundamental price. In rounds 3 and 4 prices are on average farther away from the fundamental price than in the first two rounds. Here the average distance to the fundamental price increases until about period 13, thereafter the mean distance gradually decreases over time. Recall that the fundamental price p^f ranges from 33 to 100 so an absolute distance of 20 or more is considerable.

The bottom panel of figure 1 shows how well the strategies predict on average (the absolute distance between prediction and realized price $|p_t - p_t^e|$). Note the difference in scale between the two plots; realized prices are much closer to the predictions than to the fundamental price (on average). This means that the strategies (sometimes) correctly expect prices to be away from the fundamental price. In rounds 3 and 4 where prices in the first 13 periods on average move away from the fundamental (top panel figure 1) the prediction errors do *not* increase in these periods. This illustrates the positive feedback structure of the asset market model (realized prices follow average forecasts). Apparently the fundamental price is not a good forecast of the realized market price. Participants seem to be aware of this and their forecasts are closer to the realized price.

Figure 1: Absolute distance from the fundamental (top panel) and absolute prediction error (bottom panel) averaged over 1296 simulations. All 81 pairs of mean dividend, $\bar{y}$, and interest rate, r , were simulated 16 times

	<i>Sample mean SD prices</i>	<i>Winner</i>	<i>Main part of winning strategy</i>	<i>Characterization</i>
Round 1	4.96	subject 12	$p_t^e = (p_{t-2} + p_{t-3} + p_{t-4})/3$	Mean price last 3 periods
Round 2	6.73	subject 17	$p_t^e = (p_{t-2} + p_{t-2}^e)/2$	Adaptive expectations
Round 3	21.77	subject 11	$p_t^e = p_{t-2} - r/10$	Naive minus small constant
Round 4	22.34	subject 19	$p_t^e = p_{t-2} + 2(p_{t-2} - p_{t-3})$	Trend following

Table 2: The sample mean standard deviation of the first 50 prices and the prediction strategies of the winners. To limit the size of the table we only report the most relevant part of the winning strategies⁵.

For each of the simulations the standard deviation of the prices in the first 50 periods is calculated (second column table 2). Prices are on average more volatile in later rounds. Table 2 also shows that the winning strategies are quite different in the different rounds. An (weighted) average of previous prices (round 1) and adaptive expectations (round 2) are more successful in the first two rounds where volatility is relatively low. A naive strategy or a trend following strategy is more successful in the last two rounds where volatility is high.

4.3 Medium run dynamics

In the short run only very few simulations (overall about 4%) have converged to a steady state (based upon periods 26-50). The question arises whether the low percentage of convergence is due to an inherently unstable steady state, or to a slow speed of convergence to a stable steady state. Table 3 summarizes percentages of convergence to a steady state, convergence to a cycle and non-convergence in the *medium run*, that is, the first 200 periods. A price sequence is defined to converge to a steady state if all prices in the periods 151-200 are within a range of 0.1 point⁶. A similar criterion is used for cycles. For example, in a 2-cycle all prices in the odd periods are within a 0.1 range and also all prices in the even periods are within a 0.1 range (and we do not observe a steady state). In the medium run the convergence of the price to a steady state is highest in the third round and lowest in the first round while the non-convergence is highest in the first round and lowest in the second round. In rounds 1 and 2 the percentage of convergence to periodic cycles is still above 20%, in rounds 3 and 4 this fraction has decreased to about 5% and 1.6%, respectively. In the final round about 50% of all simulations has settled down to a steady state but the other half

⁵ All winning strategies are conditional and therefore too complicated and too long to report in the table. We report that part of the strategy for which the condition (for example an "if" condition) was satisfied most of the time.

⁶ The results do not change much with a different measure of convergence. For example if, for convergence, the difference between the highest and lowest price must be smaller than 1 (instead of the 0.1) we find the same amount of

does not converge at all.

<i>Convergence</i>	<i>Round 1</i>	<i>Round 2</i>	<i>Round 3</i>	<i>Round 4</i>
steady state	24.3%	47.1%	57.2%	49.3%
2 cycle	14.1%	17.4%	0.3%	0.2%
3 cycle	0.2%	0.0%	0.1%	0.1%
4 to 10 cycle	7.0%	20.6%	4.9%	1.4%
non-convergence	54.3%	15.0%	37.7%	49.1%

Table 3: Medium run dynamics. Percentage of convergence to a steady state, a period 2-cycle, cycles of periods 3-10 and non-convergence (i.e. non-convergence and convergence to periods 11 or higher-order cycles) of 1296 simulations, based upon the periods 150-200.

4.4 Long run dynamics

In the medium run the percentage of convergence to a steady state is increasing over the rounds, but there are still about 50% non-converging price sequences in round 4. One explanation for this phenomenon is that the steady state is *stable* but the learning process converges only slowly. Another explanation can be that the steady state is intrinsically *unstable* and that prices will never converge to the steady state. To investigate which of the scenarios explains the properties in the medium and short run dynamics we study the long run dynamics. The long run dynamics is defined here as 1000 periods.

Table 4 shows the results for the long run dynamics. Convergence is defined in the same way as in the medium run analysis of section 4.3, but now based upon the periods 951-1000. Table 4 also makes a distinction between convergence to the fundamental steady state or to a steady state different from the fundamental price. We find that as the convergence to a steady state increases over the rounds the convergence to the fundamental price also increases. In the long run the majority of the simulations converges to a steady state. While in round 1 only approximately 40% of the simulations converge to a steady state, in round 4 this has doubled to 80%. Compared with the results from the medium run, in the long run there is more convergence to a steady state and less non-convergence. Many of the non-converging sequences in the medium run have dissolved and converged in the long run. Hence, the speed of convergence is slow. Notice however that there is a considerable fraction (60%) that does not converge to the fundamental steady state, but either converges to a steady state different from the fundamental (40%) or does not converge at all (20%)

---

convergence to a steady state and cycles with a small period (smaller than 4). On the other hand we find more convergence to higher periods and therefore less non-convergence is observed.

<i>Convergence</i>	<i>Round 1</i>	<i>Round 2</i>	<i>Round 3</i>	<i>Round 4</i>
Fundamental	10.8%	18.3%	19.8%	39.7%
non-fundamental	27.8%	40.3%	44.6%	40.7%
Steady state	38.6%	58.6%	64.4%	80.4%
2 cycle	18.8%	21.4%	0.2%	0.1%
3 cycle	0.2%	0.1%	0.0%	0.0%
4 to 10 cycle	10.7%	11.9%	1.5%	0.6%
non-convergence LE =0	7.0%	2.5%	8.2%	5.9%
non-convergence LE >0	24.9 %	5.6%	25.7%	13.0%

Table 4: Percentage of convergence to a steady state, a period 2-cycle, cycles of periods 3-10 and non-convergence (i.e. non-convergence and convergence to periods 11 or higher-order cycles) of the prices within 1000 periods of 1296 simulations. Two different kinds of non-convergence simulations are distinguished based upon the estimated largest Lyapunov Exponent (LE)(1000 periods are used to estimate the Lyapunov exponent). A positive Lyapunov exponent implies chaotic price fluctuations.

To study the possible occurrence of chaos and strange attractors the Wolf algorithm (Wolf et al. 1985) is applied to estimate the largest Lyapunov exponent over the period 1-1000. A positive Lyapunov exponent implies that the system exhibits sensitive dependence upon initial conditions and is chaotic⁷. Almost 75% of the non-converging price sequences have a positive Lyapunov exponent, i.e. 17.3% of all simulations have a positive Lyapunov exponent (13% in round 4).

Most of the simulations converging to some steady state do *not* exactly converge to the fundamental steady state (a steady state is considered as fundamental if the distance to the fundamental price is less than 1). However, after 1000 time periods, most of the price sequences converge to a price close to the fundamental price, as figure 2 shows.

Summarizing, in the long run a large fraction of the simulations converges to the fundamental price. The fraction of convergence increases over the rounds and in round 4 even 80% of all simulations converge to a neighborhood of the fundamental price. Nevertheless, a small fraction of the price sequences does not converge and even in round 4 13% is chaotic with a positive Lyapunov exponent.

⁷ 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. 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 (Wolf et al. 1985). For other algorithm parameter values similar results were obtained.

Figure 2: Frequency distribution of the difference between the steady state price, at $t=1000$, and the fundamental price. In round 1 500 simulations converged to a steady state, in round 2, 3 and 4 755, 829 and 1038 simulations converged to a steady state. The distance is zero if the price converges to the fundamental price $p^f = \bar{y}/r$.

4.5 Contribution of individual strategies to non-convergence

The previous section reported that only approximately 10 to 40% of all simulations converge to the fundamental steady state price and about 25% of the price sequences do not converge at all within 1000 periods. One rotten apple can spoil the whole crate: it is possible that a single strategy (or a few strategies) cause the observed non-convergence. This section investigates whether there are specific strategies that ‘destabilize’ the market and cause non-convergence.

The first four columns in table 5 show for each strategy the percentages of simulations that strategy was in that converged to a steady state or a cycle. In order to find ‘destabilizing’ strategies the following step-wise procedure was applied. In each round the strategy with the lowest percentage of convergence was excluded and new simulations were run. A strategy was excluded if this would result in more convergence to a steady state and less non-convergence. New simulations were run and the procedure was repeated. This process was stopped when the results did not change significantly anymore. This way eight ‘destabilizing’ strategies (one, two or three in each round)

were found. These strategies belong to four participants⁸. Table 6 shows the percentages of convergence and non-convergence in the long run, with the exclusion of these ‘destabilizing’ strategies. Of course, by construction these markets are (much) more stable. The percentage of convergence to a steady state increases over the rounds from 62% in round 1 to 90% in round 4. However, there is only a modest improvement in convergence to the *fundamental* price in comparison with table 4 (overall 22% to 25%).

By construction, the exclusion of these ‘destabilizing’ strategies makes the markets more stable (compare table 6 with table 4). The number of simulations that do not converge to a steady state decreases with about 40%, but the effect on convergence to the fundamental price is much more modest. The question is whether these ‘destabilizing’ strategies destabilize by making wild and stupid predictions, or whether the instability is caused by the interaction of these strategies with other strategies. If these strategies are unstable ‘by nature’, one would expect also non-convergence in homogeneous markets (markets where the prediction of one strategy determines the price or, equivalently, all six prediction strategies are the same). For all strategies 81 homogeneous markets are simulated (9 different values of the interest rate and 9 different values of the mean dividend). The second part of table 5 shows the convergence in the homogeneous simulations per strategy. Of the 71 strategies, 43 always converge. There are also a number of strategies that almost never converge in the homogeneous simulations (e.g. the strategies of subjects 1, 13 and 20). Note that the strategy of participant 9 in round 3 and 4 has a 100% convergence in the homogeneous markets, but is one of the ‘destabilizing’ strategies in the heterogeneous markets. The 28 strategies with less than 100% convergence (in the homogeneous markets) are not necessarily bad predictors. Although 23 of these 28 are in the bottom half of the ranking in the original strategy environment (heterogeneous, 50 periods) two strategies are winners: strategy 19 in round 1 and strategy 11 in round 3 (and strategies 11 and 18 are in the third and fourth places in round 1). The bold printed figures in the table refer to the “destabilizing strategies”, discussed earlier. Six out of eight are also unstable in the homogeneous simulations (statistically significant at 5%, binominal test with $n=8$ and $p=28/71=.394$). However, the relationship is not very strong; some strategies that are unstable in the homogeneous simulations are quite stable in the heterogeneous simulations (e.g. strategy 1 and 20) while the “destabilizing” strategies of subject 9, round 3 and 4, are very stable in the homogeneous simulations. The Pearson correlation between the percentages convergence in homogeneous and heterogeneous simulations is only 0.26.

---

⁸ In round 1 we excluded three strategies, 9, 11 and 17. Strategy 13 was excluded in round 2. In round 3 and 4 strategies 9 and 13 were excluded.

<i>Subject</i>	<i>Heterogeneous markets</i>				<i>Homogeneous markets</i>			
	<i>Round 1</i>	<i>Round 2</i>	<i>Round 3</i>	<i>Round 4</i>	<i>Round 1</i>	<i>Round 2</i>	<i>Round 3</i>	<i>Round 4</i>
1	74	90	68	80	0	0	0	12
2	72	93	72	83	100	100	100	100
3	72	x	x	x	100	x	x	x
4	78	92	67	81	100	100	100	100
5	72	93	x	86	99	99	x	99
6	75	95	x	x	100	100	x	x
7	72	93	69	82	100	100	100	100
9	31	x	47	73	0	x	100	100
10	72	x	x	x	88	x	x	x
11	64	90	72	79	28	63	63	63
12	73	93	67	77	100	100	100	100
13	72	83	54	75	100	5	5	5
14	70	95	69	x	100	100	100	x
15	71	92	75	85	100	100	100	100
16	69	95	77	84	100	100	100	100
17	56	93	66	82	4	100	100	100
18	69	92	76	82	0	25	25	53
19	61	92	61	80	36	35	35	38
20	72	91	65	84	100	0	0	0
21	70	94	x	84	100	100	x	100
22	62	93	50	x	100	100	100	x
Average	68	92	66	81	74	77	70	73

Table 5: For each strategy percentages of convergence (to steady state or cycle) are presented in heterogeneous (columns 2-5) and homogeneous markets (columns 6-9). For every strategy a homogeneous market is simulated 81 times (all combinations of dividend yield and interest rate). The bold printed numbers refer to “destabilizing strategies”, see main text.

<i>Convergence</i>	<i>Round 1</i>	<i>round 2</i>	<i>round 3</i>	<i>round 4</i>
Fundamental	24.1%	10.9%	22.2%	42.7%
non-fundamental	38.0%	54.1%	64.7%	47.1%
steady state	62.1%	65.0%	86.9%	89.8%
2 cycle	22.8%	21.6%	0.0%	0.1%
3 cycle	0.00%	0%	0%	0.2%
4 to 10 cycle	7.0%	10.9%	1.9%	0.9%
non-convergence	8.1%	2.5%	11.3%	9.0%

Table 6: Heterogeneous simulation without “destabilizing strategies”. Percentage of convergence to a steady state, a period 2-cycle, cycles of periods 3-10 and non-convergence (i.e. non-convergence and convergence to periods 11 or higher-order cycles) of the prices within 1000 periods of 1296 simulations without the “destabilizing” strategies. In round 2, 3 and 4 we excluded 2 strategies while in round 1 we excluded three strategies (see footnote 8).

Summarizing, there are some strategies (‘destabilizing strategies’) that contribute more than average to the non-convergence. By excluding these strategies from the simulations the percentage

of convergence to a steady state increases substantially but the percentage of convergence to the fundamental steady state increases only slightly. The ‘destabilizing’ strategies seem also to be less stable in homogeneous simulations, but the relationship is not very strong. We conclude that a large part of the non-convergence is caused by the interaction of the strategies.

4.6 Final Experiment

The final experiment was conducted after the participants submitted their fourth round strategy and before they received the fourth round results. The main goal of this final experiment is to study the relationship between actual behavior of subjects and the strategy they submit. Participants are in a market the only ‘human’ agent, together with 5 computer strategies of round 3. Since the participants did not yet receive the fourth round results, the information in the final experiment is the same as for the fourth round strategies. We can, therefore, compare their predictions in the experiment with the fourth round predictions of their strategies⁹. Just like in the introductory experiment we conduct two markets in the final experiment. The fundamental price in the first market is $p^f = 60$ while the fundamental price in the second market is $p^f = 40$.

Table 7 shows the average quadratic error of the participants (columns two and five) and their strategies (columns three and six). A striking feature of table 7 is that the strategies and the participants have a qualitatively similar average prediction error. In particular, large errors by the participants coincide with large errors by their strategies. The average prediction errors of the participants and their strategies are strongly correlated (Pearson correlation is 0.80).

We can compute the ranking of the participants and of the strategies based on the average quadratic prediction error. The difference between the ranking of the participant and the ranking of its strategy is never more than 2. That is, the strategies of the subjects who perform good/bad also perform good/bad. The Pearson correlation of the rankings is 0.96, i.e. the rankings after the final experiment of the strategies and of the subjects is high and almost perfect.

---

⁹ For comparison, we let the strategy of the participant predict the price in the same market experiment but the strategy does not participate in the market, i.e. the strategies prediction does not influence the realized market price.

	Part 1			Part 2				
	<i>Quadratic error participant</i>	<i>Quadratic error strategy</i>	<i>Quadratic distance participant and strategy</i>	<i>Quadratic error participant</i>	<i>Quadratic error strategy</i>	<i>Quadratic distance participant and strategy</i>	sign perc.	pred. corr.
Subject								
1	0.4	0.6	0.1	5.5	10.0	3.0	0.91	0.98
2	17.6	50.9	23.0	33216.3	30602.3	23090.2	0.85	0.90
4	0.9	3.6	1.2	1.0	7.5	6.2	0.95	0.72
7	1648.2	1399.8	2822.7	20.3	8.2	8.8	0.85	0.98
9	11.3	23.0	28.0	5.5	10.3	7.9	0.85	0.67
11	7.4	7.9	5.4	2.7	14.0	13.6	0.85	0.91
12	5.1	19.6	18.0	21171.3	86231.3	80115.0	0.66	0.68
13	11.4	10.1	1.9	12.7	17.7	17.1	0.90	0.93
15	2.0	1.7	0.7	11.3	14.5	1.4	0.71	0.91
16	9242.3	14283.6	14113.6	15266.6	76663.4	56956.1	0.79	0.90
17	2.0	3.3	0.4	22.5	73.2	38.1	0.77	0.92
18	6.0	6.2	5.4	3.6	18.4	7.9	0.86	0.57
19	13.3	13.2	7.2	6.8	13.5	12.3	0.74	0.97
20	15172.3	24413.1	9289.4	30.3	11.7	45.3	0.76	0.82
21	65.1	103.1	125.0	20.4	22.9	11.4	0.91	0.91

Table 7: The first column shows the participants number. Column 2 to 4 (5 to 7) show for the first (second) part of the final experiment the average quadratic error per period of the participant and the strategy and the average quadratic distance per period between the prediction of the strategy and the participant's prediction. The eighth row (sign perc.) shows the fraction that the participant adapts his prediction in the same direction as his strategy. Finally, the ninth row (pred. corr.) shows the average correlation (over the two parts) between the strategy's predictions and the subject's predictions over 50 periods.

The analyses above are based upon 50 periods averages. On the level of individual predictions, column 8 of table 7 shows the times that the strategy and the participant adapt their prediction in the same direction¹⁰. We find that 9 out of 15 participants adapt their prediction in the same direction as its strategies in at least 85% of the time. The last column shows that the correlation between the predictions of the strategy and the predictions of the participants is high. Participants 1 and 4 both have a correlation of 98%, i.e. almost perfect correlation.

In summary, the final experiment indicates that in general the strategies give a reasonably good description of the prediction behavior of the participants.

¹⁰ We count the times that the prediction, p_t^e , is higher/lower than the previous prediction, p_{t-1}^e for both the participant and the strategies.

5 Concluding remarks

One of the goals of this study is to investigate the kind of prediction strategies used by participants in an experimental setting. The strategy method used in this study has some clear disadvantages: programming all the strategies is very labor intensive and it is hard to keep all participants in an experiment that lasts for several weeks. However, there are also some important advantages. It is the only way to get unambiguous complete strategies of individuals that can also be used in long-term simulations, and the learning possibilities may enhance the external validity. This study is a demonstration of the kind of answers a strategy experiment can provide.

We find that only few participants use simple (linear) strategies. Most of the participants submit complicated strategies and the strategies become more complicated over the rounds. In the short run (within 50 periods) the prices have not yet converged to either the fundamental price or another steady state price or periodic cycle. In the first two rounds the average distance to the fundamental is fairly small. In rounds 3 and 4 the average distance to the fundamental price increases until period 13, and then gradually decreases, but remains larger than in the first two rounds. Also the average prediction error is smaller in the first two rounds than in rounds 3 and 4. In other respects the learning over the rounds lead to an increase in stability. The number of converging price sequences increases over the rounds in both the medium and the long run. In the medium run the prices convergence to a steady state or periodic cycles but this percentage is not very high. In the long run in round 4 almost 40% of the price sequences converge to the fundamental value. Another 40% also converges to a steady state price but not to the fundamental value. These results suggest in general slow convergence and learning of the subjects over the rounds.

How much individual strategies attribute to market stability or instability is difficult to say. Eight ‘destabilizing strategies’ are identified (most markets in which these strategies participated didn’t converge). New simulations without these strategies showed much more stable markets. However, convergence to the fundamental price did only increase slightly. The ‘destabilizing strategies’ seem on average to be less stable in homogeneous simulations also, but the difference is not very large. A part of the non-convergence seems to be caused not by individual strategies but by the *interaction of several strategies* together.

From the final experiment we find that the correlation between the strategies prediction and the participants forecast is high. Also the ranking of the strategies and the ranking of the participants based on the quadratic forecast errors are almost the same. This suggests that the strategies prediction is ‘close’ to the forecast of the participant. Furthermore, the average quadratic

distance between the prediction of the strategy and participant is small. Together with the high percentages of adaptation in the right direction and the high correlation between the strategies and the predictions of the participants we conclude that the strategies are a good representation of what participants do in an experiment.

We conclude that it is hard to learn the correct ‘fundamental’ price level in a stationary speculative asset market. Deviations from the fundamental price are persistent and fundamentals do not yield accurate predictions of asset prices, even in this relatively simple laboratory environment. Real world markets are more complex, fundamentals like dividends and interest rate change over time and information is not always complete. The relation between real world markets assets prices and the underlying fundamentals may be even weaker than in our experiments.

References

- 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., Durlauf, S.N. and Lane, D.A. (eds.), *The economy as a complex evolving system* Vol. II, Addison-Wesley, Reading MA, pp. 15-44.
- Axelrod, R. (1984). *The Evolution of Cooperation*. New York: New York: Basic Books.
- Brock, William A.; Hommes, Cars H. (1997). A rational route to randomness *Econometrica* 65, 1059-1095.
- Brock, William A.; Hommes, Cars H. (1998). Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model. *Journal of Economic Dynamics Control* 22 (1998), 1235-1274
- Gaunersdorfer, Andrea (2000). Endogenous fluctuations in a simple asset pricing model with heterogeneous agents *Journal of Economic Dynamics and Control* 24, 799-831.
- Gode, Dan K., and Sunder, Shyam (2001). Double auctions dynamics: structural effects of non-binding price controls. Working paper.
- Hommes, Cars, Sonnemans, Joep , Tuinstra, Jan and van de Velden, Henk (2002a) Expectations and bubbles in asset pricing experiments. CeNDEF working paper.
- Hommes, Cars, Sonnemans, Joep , Tuinstra, Jan and van de Velden, Henk (2002b) Coordination of expectations in asset pricing experiments. CeNDEF working paper.
- Keser, C. (1992). *Experimental Duopoly Markets with Demand Inertia, Game-Playing Experiments and the Strategy Method*. Heidelberg, Germany: Springer Verlag.
- LeBaron, B., (2000), Agent based computational finance: suggested readings and early research, *Journal of Economic Dynamics and Control* 24, 679-702.
- Lux, T. and Marchesi, M. (1999), Scaling and criticality in a stochastic multi-agent model of a financial market, *Nature* 397, February 498-500.
- Offerman, T., Potters, J. and Verbon, H (2001). Cooperation in an overlapping generations experiment. *Games and Economic Behavior*; 36, 264-275.
- Selten, R., Mitzkewitz, M. and Uhlich, G.R. (1997). Duopoly strategies programmed by experienced players. *Econometrica*, 65, 517-555.
- Sonnemans, Joep (1998). Strategies of search. *Journal of Economic Behavior and Organization* 35, 309-332.
- Sonnemans, Joep, Cars Hommes, Jan Tuinstra and Henk van de Velden (1999). The instability of a heterogeneous cobweb economy: a strategy experiment on expectation formation. Cendef working paper.

- Sterman, John D. (1989). Deterministic Chaos in an Experimental Economic System *Journal of Economic Behavior and Organization* 12, 1-28.
- van Velden, Henk (2001). *An experimental approach to expectation formation in dynamic Economic systems* PhD dissertation University of Amsterdam.
- Wolf, A., Swift, J.B., Swinney, H.L. and Vastano, J.A. (1985). Determining Lyapunov exponents from a time series. *Physica D* 16, 285-317.