Learning in a Complex World
Insights from an OLG lab experiment

Cars Hommes[†] Stefanie J. Huber[⋆] Daria Minina[§] Isabelle Salle[‡]

Version: 2022-11-26

Abstract

This paper brings novel insights into group coordination and price dynamics in complex environments. We implement an overlapping-generation model in the lab, where the output dynamics are given by the well-known chaotic quadratic map. This model structure allows us to study previously unexplored parameter regions where the perfect-foresight dynamics exhibit chaotic dynamics. This paper highlights three key findings: First, the price converges to the simplest equilibria, namely the monetary steady state or the two-cycle in all markets. Second, we document a novel and intriguing finding: we observe a non-monotonicity of the behavior when complexity increases. Convergence to the two-cycle occurs for the intermediate parameter range, while both the extreme scenarios of a simple stable two-cycle and highly nonlinear dynamics (chaos) lead to coordination on the steady state in the lab. All indicators of coordination and convergence significantly exhibit this non-monotonic relationship in the learning-to-forecast experiments—this finding also persists in the learning-to-optimize design. Finally, convergence in the learning-to-optimize experiment is more challenging to achieve: coordination on the two-cycle is never observed, although the two-cycle Pareto dominates the steady state.

Keywords: OLG models, Complexity, Learning, Equilibria selection, Learning-to-forecast experiment, learning-to-optimize experiment.

JEL Classifications: C62, C68, C91, C92, E13, E70, G12, G41.

We dedicate this paper to the memory of Jasmina Arifovic. May her work and friendship continue to inspire us. We would like to thank participants at the Theory and Experiments in Monetary Economics Conference in 2022 at George Mason University, the 2022 Experimental Finance Conference in Bonn, the 12th International Workshop on Theoretical and Experimental Macroeconomics in Berlin, the 4th Behavioral Macroeconomics Workshop in Bamberg, and the CREED Brown Bag Seminar at the University of Amsterdam for valuable comments and suggestions. The views and opinions expressed in this article are those of the authors and do not necessarily reflect the view of the Bank of Canada.

[†]Financial Markets Department, Bank of Canada; CeNDEF, Amsterdam School of Economics, University of Amsterdam; Tinbergen Institute, Amsterdam. E-mail: c.h.hommes@uva.nl/CHommes@bank-banque-canada.ca

[⋆]CeNDEF, Amsterdam School of Economics, University of Amsterdam; Tinbergen Institute, Amsterdam.
E-mail: s.j.huber@uva.nl.

[§]CeNDEF, Amsterdam School of Economics, University of Amsterdam. E-mail: d.minina2@uva.nl.

[‡]Department of Economics, University of Ottawa; Macro and International Economics, Amsterdam School of Economics, University of Amsterdam; Tinbergen Institute, Amsterdam.
E-mail: isalle@uottawa.ca/I.L.Salle@uva.nl.

1 Introduction

This paper provides new evidence of coordination of human subjects in a lab environment with previously unexplored types of complex dynamics. We find non-monotonic equilibrium selection as a function of the underlying complexity of the environment.

Indeterminacy and equilibrium selection are difficult but important questions in macroeconomics and finance that bear both modeling and practical implications. Pinning down a unique equilibrium is an essential prerequisite for using a model as a framework for policy analysis or for estimating it against empirical data.

The question of equilibrium selection can be approached theoretically and empirically. The theory is useful for defining the set of equilibria in a model, but theoretical selection criteria—typically learning mechanisms—are usually not selective enough to establish their empirical relevance. In particular, the precise specification of the learning rule (used by the agents) already restricts which equilibria may be achieved. In other words, any equilibrium can be reached in theory by designing an appropriate learning function. What type of learning rules economic agents use is clearly an empirical question. Hence, the theoretical approach has limits for understanding which equilibria would prevail in the real world. Additionally, most learning schemes used in the theoretical literature are designed under the representative-agent paradigm. And hence, these learning rules do not allow to modeling nor observing interactions between heterogeneous agents and the possible coordination outcomes.[1]

Laboratory experiments are a research method that has been proven beneficial to empirically tackle the problem of equilibrium selection.[2] This method conveniently offers

[1]Evolutionary learning models, inspired by the concept of natural section, have been implemented based on genetic algorithms (Arifovic, 1994, 1995, 1998; Arifovic and Ledyard, 2018; Bullard and Duffy, 1995, 1998) or probability choice models (Brock and Hommes, 1997) and rely on heterogeneity. Adaptive learning, the most common form of learning in macroeconomic models, typically involves a representative agent (Evans, 1985; Evans and Honkapohja, 2001). The idea that empirically relevant equilibria need to be a stable outcome of an adaptive learning process may be traced back to the contributions of DeCanio (1979); Lucas Jr (1978, 1986). Note that this form of bounded rationality does not deviate from rationality in the broad sense: agents are just assumed to lack the necessary information to form rational expectations, see Sargent et al. (1993). See Evans and Honkapohja (2009) for a survey of learning in the field.

[2]See the pioneer contributions of Aliprantis and Plott (1992); Arifovic (1995); Heemeijer et al. (2009b); Lim et al. (1994); Marimon et al. (1993); Marimon and Sunder (1993, 1994, 1995). These OLG experiments may be considered the foundations of the field of experimental macroeconomics. Duffy (2016)

a controlled environment where information and fundamentals of the economy are set by the researchers, but the decisions and resulting economic behaviors are left to human subjects. In group experiments, in particular, one may investigate the coordination of a group of interacting and heterogeneous participants and hence equilibrium selection.

Arifovic et al. (2019) are the first to explore chaotic dynamics in general equilibrium models in the laboratory and point at the prevalence of simple equilibria, even when many more equilibria co-exist, including high-periodic cycles along with chaotic dynamics. Subjects learn to coordinate but always do so on steady states or two-period cycles. This result holds in so-called learning-to-forecast experiments (LtFEs) and learning-to-optimize experiments (LtOEs). In LtFs, subjects' beliefs (e.g., their expectations about future prices) are elicited, while their optimal decisions (conditional on their beliefs) are computerized. In contrast, subjects are asked to make these decisions in the LtOEs. Importantly, Arifovic et al. (2019) find that as complexity increases, coordination on two-cycles becomes increasingly likely, while coordination on the steady state disappears.

This paper advances this line of research by exploring a wider range of complex environments. In particular, we explore behaviors in parameter regions where the underlying model exhibits chaotic dynamics once the period-three cycle has lost stability in the backward perfect-foresight dynamics. To do so, we use an OLG model where production evolves according to a widely-studied quadratic map. This allows us to implement more complex and non-linear environments than in Arifovic et al. (2019)—by considering higher values of the complexity parameter.

This feature enables us to establish our main result: the relationship between the complexity of the chosen equilibrium in the lab and the model parameter that tunes the complexity of the environment is non-monotonic. The coordination on two-cycles only occurs in the LtFE for the intermediate range of the parameter values. In contrast, coordination on the steady state is systematic in the LtOE and happens for both low and high values of the complexity parameter in the LtFE. For low parameter values, the model dynamics are represented by the simple steady state. For high parameter values,

provides a comprehensive survey of this literature.

the model dynamics exhibit the other extreme—chaotic dynamics after the period-three cycle has lost stability.

Interestingly, indicators of aggregate price convergence, individual coordination of forecasts and production decisions or earnings efficiency also exhibit non-monotonic patterns. For the intermediary region of the model parameter, subjects have a harder time making sound decisions and coordinating on an equilibrium compared to simpler (i.e., low parameter values) or highly complex (i.e., large parameter values) environments. . These non-monotonic behaviors are a robust finding of our study: they hold both in the LtF and LtO versions of our experiment. Furthermore, it holds even though our design makes the two-cycle Pareto-dominant—as opposed to the design in Arifovic et al. (2019) where the two-cycle and the steady state yield the same payoff by design.[3]

We connect this non-monotonicity to subjects' decision rules that we estimate using the time series of individual choices. We find that participants' chosen strategies are complexity-dependent in a non-monotonic way. For the intermediary region of the model parameter, where the two-cycle may emerge in the LtFE and coordination in the lab was harder in both the LtFE and LtOE designs, significantly more subjects use simple adaptive decision rules. In contrast, subjects used more sophisticated decision rules in simple (i.e., low model parameter values) and highly complex environments (i.e., high model parameter values). This finding we observe in the LtF and also in the LtO experiment.

The paper is organized as follows. Section 2 presents the theoretical analysis of the OLG model and describes the design of the laboratory experiment and its implementation. Section 3 presents our experimental results, and Section 4 concludes.

2 The experimental designs

First, this section presents the underlying OLG model of the experiment along with theoretical predictions under various expectation assumptions (i.e., learning rules). Second, this section discusses the design of the LtFEs and the LtOEs, and the lab implementation.

[3]In the LtO experiment of Arifovic et al. (2019), subjects are jointly tasked with and rewarded with forecasting future returns.

2.1 The OLG model

We use the overlapping generation (OLG) model of Araujo and Maldonado (2000) because it conveniently yields production to evolve according to the widely studied quadratic map. In this model, the complexity of the dynamics increases with the value of a single parameter in the utility function. The offer curve is symmetric, which, we conjecture, may make coordination on two-cycles more difficult and coordination on higher-order cycles more likely—compared to the OLG model of Grandmont (1985) used in Arifovic et al. (2019), where the offer curve is asymmetric.

Individuals live for two periods in the model of Araujo and Maldonado (2000). They work in the first period when young and consume in the second period when old. For the experimental implementation, we assume that each generation consists of a finite number $N \geq 1$ of agents, indexed by i. Agent i derives utility from consumption when old, denoted by $c_{i,t+1}$, and suffers a disutility from labor when young, which linearly produces goods, denoted by $y_{i,t}$. These goods are sold from the young generation to the old generation at market price P_t. Households' two-period lifetime utility function is given by

$$U\left(c_{i,t}, y_{i,t}\right) = \lambda\ c_{i,t+1} - \frac{\lambda}{2} c_{i,t+1}^2 - y_{i,t}, \tag{2.1}$$

where $\lambda > 0$ is the parameter of interest—tuning the trade-off between consumption and leisure when young. Each young agent chooses how many hours to work when young to maximize consumption in the old age, subject to their budget constraint

$$P_{i,t+1}^e c_{i,t+1} \leq P_t y_{i,t}, \tag{2.2}$$

where the superscript e denotes (possibly boundedly rational) expectations. Market clearing yields

$$\sum_{i=1}^{N} y_{i,t} \equiv Y_t = \frac{M}{P_t}, \tag{2.3}$$

where $M > 0$ denotes the constant quantity of money available in the economy. Throughout the paper, capitalized letters denote aggregate variables.

As the detailed derivations in Appendix A show, at the symmetric perfect-foresight equilibrium, the first-order condition, expressed in terms of individual output, reads as

$$\lambda y_{t+1}(1 - y_{t+1}) = y_t, \tag{2.4}$$

with $y_{i,t} = y_t = \frac{Y_t}{N}$, $\forall i, t$. Equation (2.4) corresponds to the quadratic map where higher values of λ generate increasingly complex dynamics.

The feedback of the forecasts is complex and non-monotonic. The price depends positively on price expectations if the price expectation is larger than $2M$. If this condition does not hold, then the feedback is negative (see, again, Appendix A).

2.2 The experimental setting

As described in Section 2.1, individuals of the old generation take no decisions. Therefore, we implement a single-population design, where subjects act on behalf of an individual of the young generation in each period.[4] We also use a between-subject design, that is each subject is randomly assigned to a single treatment and may participate only once in the experiment. Online Appendix B and C present the instructions for the LtFE and the LtOE, respectively.

In each session, subjects are randomized into groups of seven. Each group represents one experimental economy (i.e., market). Participants are told that they act as consultants to an investment fund. The participants' decisions correspond to the decision of an individual of the young generation in the model presented in Section 2.1. Due to our focus on the long-run dynamics and convergence outcomes, the experiment lasts for 100 rounds, which is common knowledge.

The amount of fiat money in the OLG economy (M) is set to rule out the autarky solution where the asset price goes to infinity, and the young generation chooses not to work, see Appendix A. We use $M = 1.5$ and multiply the prices and forecasts displayed

[4]Arifovic et al. (2019) show that results on equilibrium selection are robust to more costly alternatives that require more participants. These OLG design alternatives include (i) participants alternating between the young and the old generation, and (ii) participants being randomly selected for entering the market (and being assigned to the young generation).

to the subjects by 50 to ensure intuitive values. For the same reasons, savings decisions are mapped into a $0-100$ range. In what follows, we present both experimental settings in parallel because they do not differ greatly.

Participants' role In the LtFE, their role is described in terms of professional forecasters, while the instructions speak about professional saving advisors in the LtOE design.[5] At the beginning of any period $t \in \{1, ..., 100\}$, the LtFE-participants are tasked with predicting the next period's asset *price*, while the LtOE-participants face an additional task. In the LtOE, participants first provide a forecast of the next period's asset *return*, and second, they need to make the saving decision. The realized return between period t and $t+1$ is calculated as follows:

$$R_t = \frac{P_t}{P_{t+1}}. \tag{2.5}$$

The return forecast helps subjects to pick a saving decision based on their savings payoff table (see Appendix Figure A2) and does not impact the aggregate outcome other than by impacting these saving decisions. Note that forecasting the return and not the next period's price is necessary to make savings decisions because the current market-clearing price may not be yet known before all savings decisions are submitted.

Sequence of events Once participants have completed their tasks, the aggregate savings Y_t is computed as follows. In the LtFE, conditional on the elicited individual forecasts $P^e_{i,t+1}$, we combine the first-order condition Eq. (A.3) and individual savings $y_{i,t}$ per Eq. (A.7), and calculate aggregate savings Y_t as per Eq. (A.6). In the LtOE, elicited individual saving decisions $y_{i,t}$ are directly summed up to $Y_t = \sum_{i=1}^{N} y_{i,t}$. The market-clearing price is given by $P_t = \frac{M}{Y_t}$. From period 2 on, period-t information together with the savings decisions and the market-clearing price in period $t-1$ gives the consumption and resulting utility level of each member of the old generation in period t and the realized

[5]We frame the decision $y_{i,t}$ in terms of savings that may more easily relate to asset returns than a production (output) decision. However, we clarify in the instructions that "savings" and "output" are used interchangeably.

return between period $t-1$ and t.

Payoff Forecasts, whether price or return forecasts, are rewarded based on their accuracy. Savings decisions are rewarded on the basis of their corresponding two-period utility level.

Subjects accumulate forecasting points in each period of the experiment, either for their price forecasts in the LtFE or for their return forecast in the LtOE. We use quadratic payoff functions, where higher forecast errors lead to lower earnings. Specifically, price forecasts yield the following amount of forecasting points:

$$\text{Price forecast payoff}_t = \max\left(0, 1300 - \frac{1300}{49}(P^e_{i,t+1} - P_{t+1})^2\right), \tag{2.6}$$

where forecast errors larger than 7 do not give rise to any point, and return forecasts are paid according to:

$$\text{Return forecast payoff}_t = \max\left(0, 1300 - \frac{1300}{4}(R^e_{i,t+1} - R_{t+1})^2\right), \tag{2.7}$$

where forecast errors higher than two yield zero point.[6]

Note that the forecasts are two-period-ahead. Hence, at the beginning of each period t, subjects have to forecast the price (or return) in the next period without yet knowing the price in the current period t. It follows that subjects discover their forecast error and the corresponding payoff for any forecast in period t at the end of period $t+1$. This time structure is made very clear in the instructions, and we check for participants' comprehension in the pre-experiment quiz (Appendix D).

Subjects are given the explicit formula (2.6) and (2.7) along with a payoff table that describes the relationship between forecast errors and possible earnings; see Online Appendix B and C.

As for the saving payoff in the LtOE, savings decisions map into utility points *via* a

[6]The exchange rate from the experimental points to euros is 0.25 euro for 1,300 points in the LtF sessions (0.3 euro for 1,300 points in the LtF pilot sessions) and 0.2 euro for 800 points in the LtO sessions.

monotonic transformation of the utility function u in Eq. (2.1):

$$U^* = 300 \times (\max(u, 0) + 3). \tag{2.8}$$

Such a transformation ensures that the saving payoff range is similar to the one in the LtFE.

Finally, in the LtOE, we use a monotonic transformation of the two-cycle payoff such that this equilibrium Pareto dominate the steady state as follows:

$$1300 * \left(\frac{U^*}{1300}\right)^{6.5}. \tag{2.9}$$

This feature is largely exploratory in light of the difficulty of coordination on a cycle as as discovered by Arifovic et al. (2019), where the two-cycles and the steady state yield the same payoff by design. We hypothesize that Pareto dominance may help favor coordination on a two-cycle.

Information set and graphical user interface (GUI) Participants know only the qualitative information about the experimental economy, not the exact equations, as standard in the related literature.

In period 1, participants enter their decisions without prior price information. The instructions only mention that prices in similar economies typically range between 10 and 100. We set the price of period 0 to 300. This price is unknown to the subjects but is used to compute the initial (period 0) return. This calibration ensures that the initial level of the return is not too close to 1—which could artificially lock-in the experiment towards the steady-state equilibrium.

In any subsequent period t, participants observe the past prices, their past own decisions, their past earnings and their cumulative payoff up until period $t - 1$. This information is shown on the GUI in table form and by a graph; see Figures 2.1 and 2.2 for the LtFE and the LtOE, respectively. It is important to note that subjects may not see the predictions and payoffs of other participants. However, they observe the aggregate

savings decision, which is equivalent, but more intuitive, as disclosing the quantity of money M in the economy.

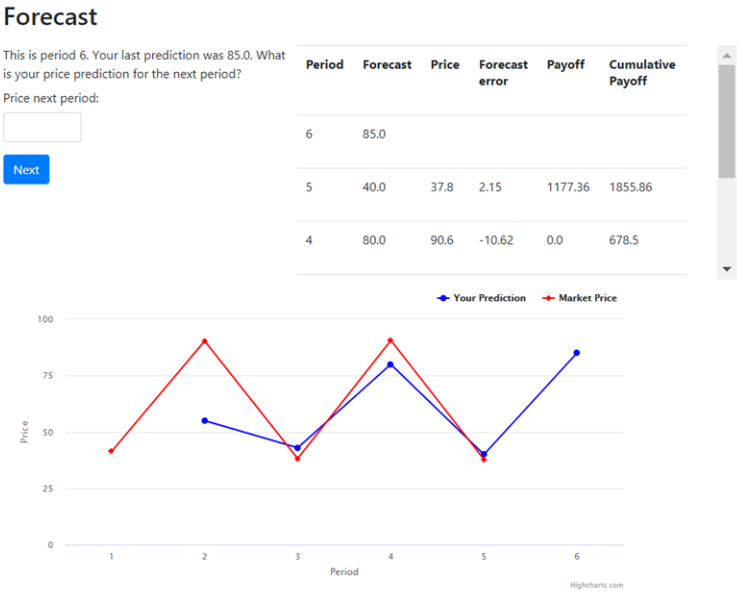


Figure 2.1: Decision Screen LtF Experiment

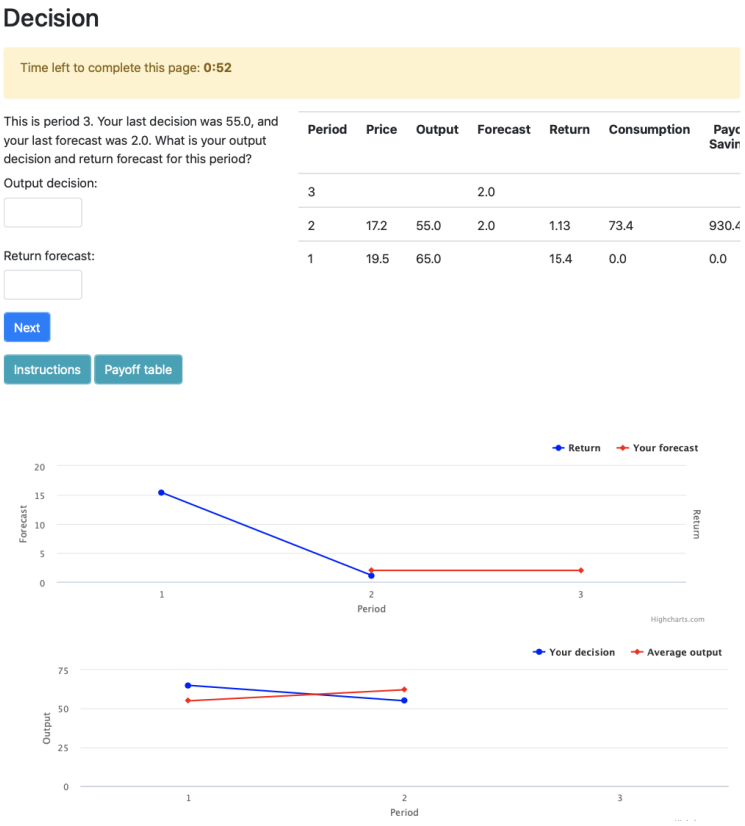


Figure 2.2: Decision Screen LtO Experiment

2.3 Experimental treatments

The model parameter λ in the utility function (2.1) is the treatment variable. Increasing λ give rise to increasingly complex economic dynamics in the model, as illustrated by the bifurcation diagram in Figure 2.3 under backward perfect foresight.

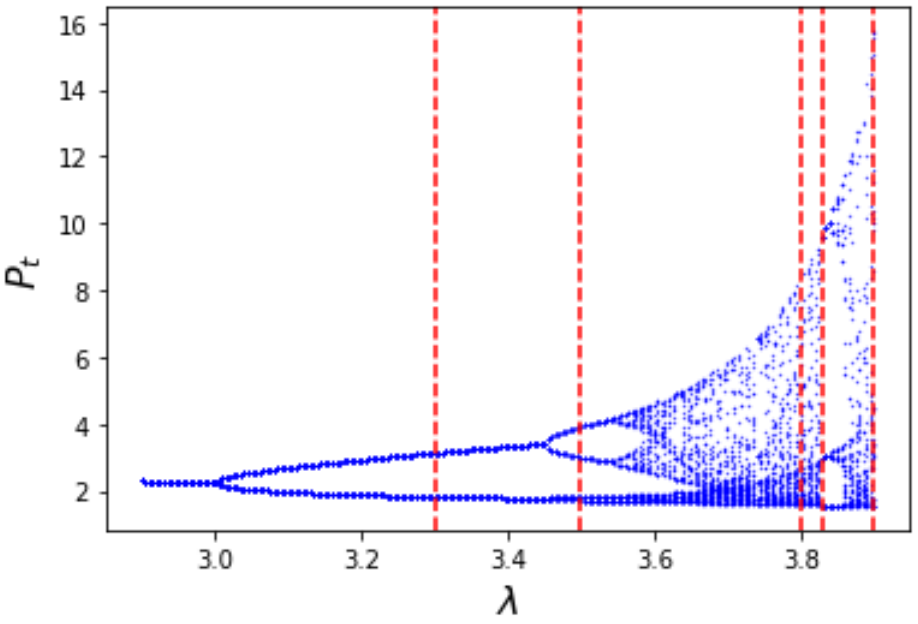


Figure 2.3: Bifurcation diagram of the price map

Notes: Bifurcation diagram under backward perfect-foresight dynamics with $P^e_{t+1} = P_{t-1}$, with $p_0 = 1.55, M = 1.5$. Dashed (red) vertical lines denote the numerical values of the treatment parameter λ.

We implement five treatments for the LtFE. We choose the treatments based on the bifurcation diagram of the widely studied quadratic map. Table 2.1 reports for each treatment the value of λ—corresponding to the red dotted vertical lines in Figure 2.3. We choose these treatments for the diversity of outcomes that they involve and the underlying complexity of the equilibria in the backward perfect foresight dynamics. In particular, we are interested in exploring parameter regions of high complexity beyond the dynamics studied in Arifovic et al. (2019). Because of this exploratory dimension, we do not bind ourselves to explicit hypotheses. Intuitively, we expect that achieving coordination is more challenging for higher values of λ, for which the dynamics become chaotic.

Specifically, we choose a treatment with $\lambda = 3.3$ which yields a stable two-cycle after a period-doubling bifurcation at $\lambda = 3$; one treatment with $\lambda = 3.5$ after another period-

Design	*LtFE*	*LtO*	
Dynamics under backward foresight	***# obs***		**λ-value**
Convergence to the two-cycle	4	-	3.3
Convergence to the 4-cycle	4	4	3.5
Chaotic behavior (without existence of the 3-cycle)	4	4	3.8
Convergence to the 3-cycle	4	4	3.83
Chaotic behavior (with existence of the 3-cycle)	4	4	3.9

Table 2.1: Experimental treatments and Model predictions

doubling bifurcation at $\lambda = 3.449$ that results in a stable four-cycle; one treatment with $\lambda = 3.8$ that involves chaotic behavior (for $3.56995 < \lambda < 3.8283$); one treatment with $\lambda = 3.83$ that gives rise to a stable three-cycle after a tangent bifurcation of the third iterate at $\lambda = 3.8283$. And finally, one treatment with $\lambda = 3.9$, after which the three-cycle has lost stability and chaotic behavior results.[7]

We implement four treatments for the LtOE, with λ values equal to 3.5, 3.8, 3.83, and 3.9. We focus on these four treatments because we observed interesting price dynamics in the corresponding intermediate λ-range in the LtFE.

Note that the above-discussed equilibria are obtained in the backward perfect-foresight dynamics, where $P^e_{t+1} = P_{t-1}$, which also corresponds to the criteria of strong E-stability (Evans and Honkapohja, 2001).[8] Figure 2.4 illustrates the convergence to the two-cycle for $\lambda = 3.3$ (Fig. 2.4a), to the four-cycle for $\lambda = 3.5$ (Fig. 2.4b), to chaotic behavior for $\lambda = 3.8$ (Fig. 2.4c) and $\lambda = 3.9$ (Fig. 2.4d) and convergence to the three-cycle for $\lambda = 3.83$ (Fig. 2.4e).

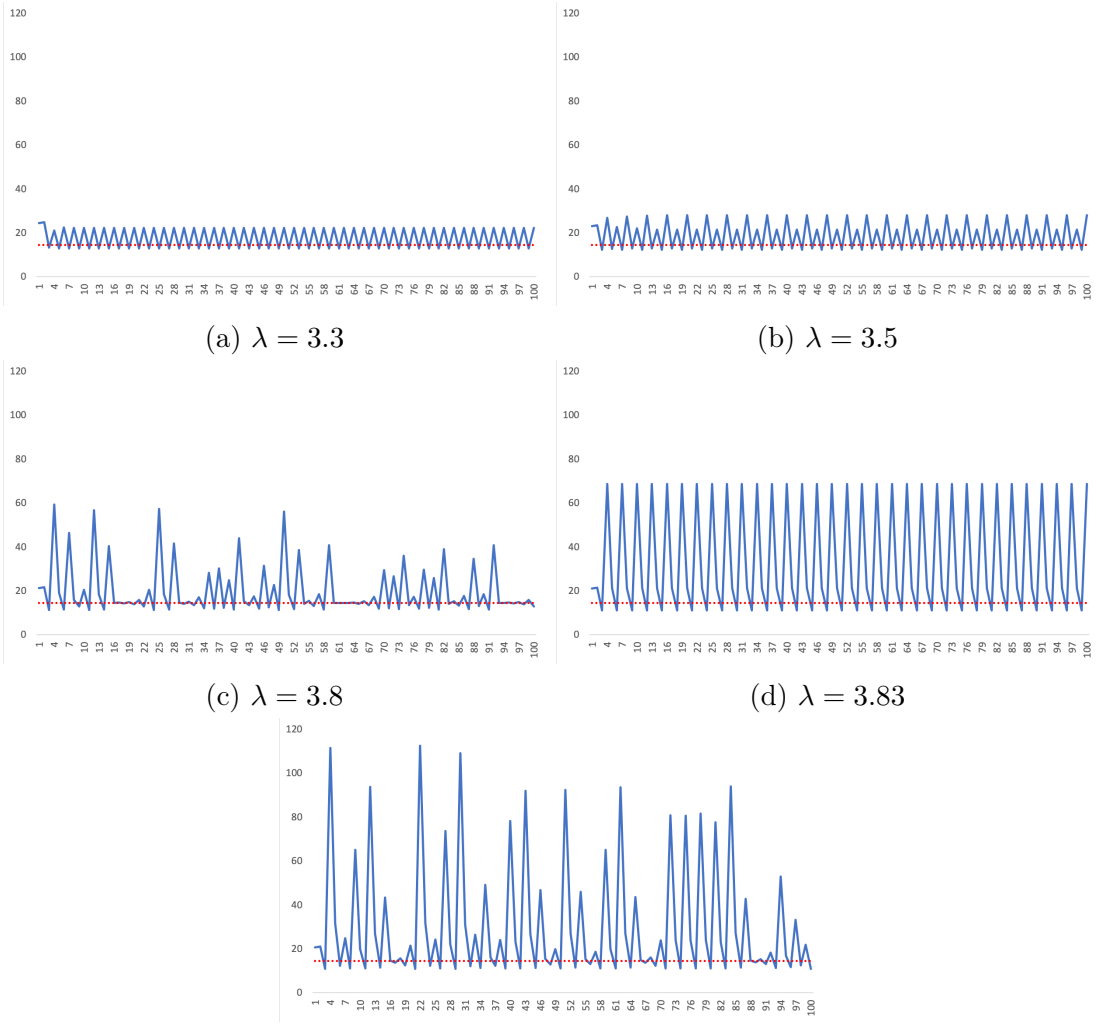
We defer the details of the stability (and convergence) results under various alternative expectation schemes to Appendix Table A1. To help the reader develop intuition on what type of convergence to expect in the experiment, we next present some illustrative

[7]Arifovic et al. (2019) did not conduct a treatment in this highly chaotic region in the context of Grandmont (1985) OLG model (which would correspond to a large value of the risk aversion of the old generation) because in this model, the price amplitude then becomes arbitrarily large, which makes it impractical for the purpose of lab implementation. By contrast, the model used in this paper makes the exploration of this parameter region in the lab more doable.

[8]Within the context of our deterministic model, the classical definition of rational expectations ($P^e_{t+1} = P_{t+1}$) correspond to forward perfect foresight. The equilibria that are stable in the backward-perfect foresight dynamics are unstable in the forward dynamics and the opposite is also true (Grandmont, 1985). In particular, this implies that for $\lambda = 3.83$, the steady state and all cycles except three-cycles are stable, while in the chaotic region with $\lambda = 3.9$, all are stable.

simulations—using different expectation schemes.

The less stringent criterion of weak E-stability is worth discussing here because it has been found to be a sufficient, but not a necessary, condition for predicting coordination in the experiment of Arifovic et al. (2019). Under this criterion, convergence towards a particular equilibrium occurs if and only if agents use a learning rule that involves the number of lags consistent with the periodicity of this equilibrium. Within the context of our model, the monetary steady state is weakly stable for all λ-values, and the two-cycle is stable for all treatments.



(a) $\lambda = 3.3$

(b) $\lambda = 3.5$

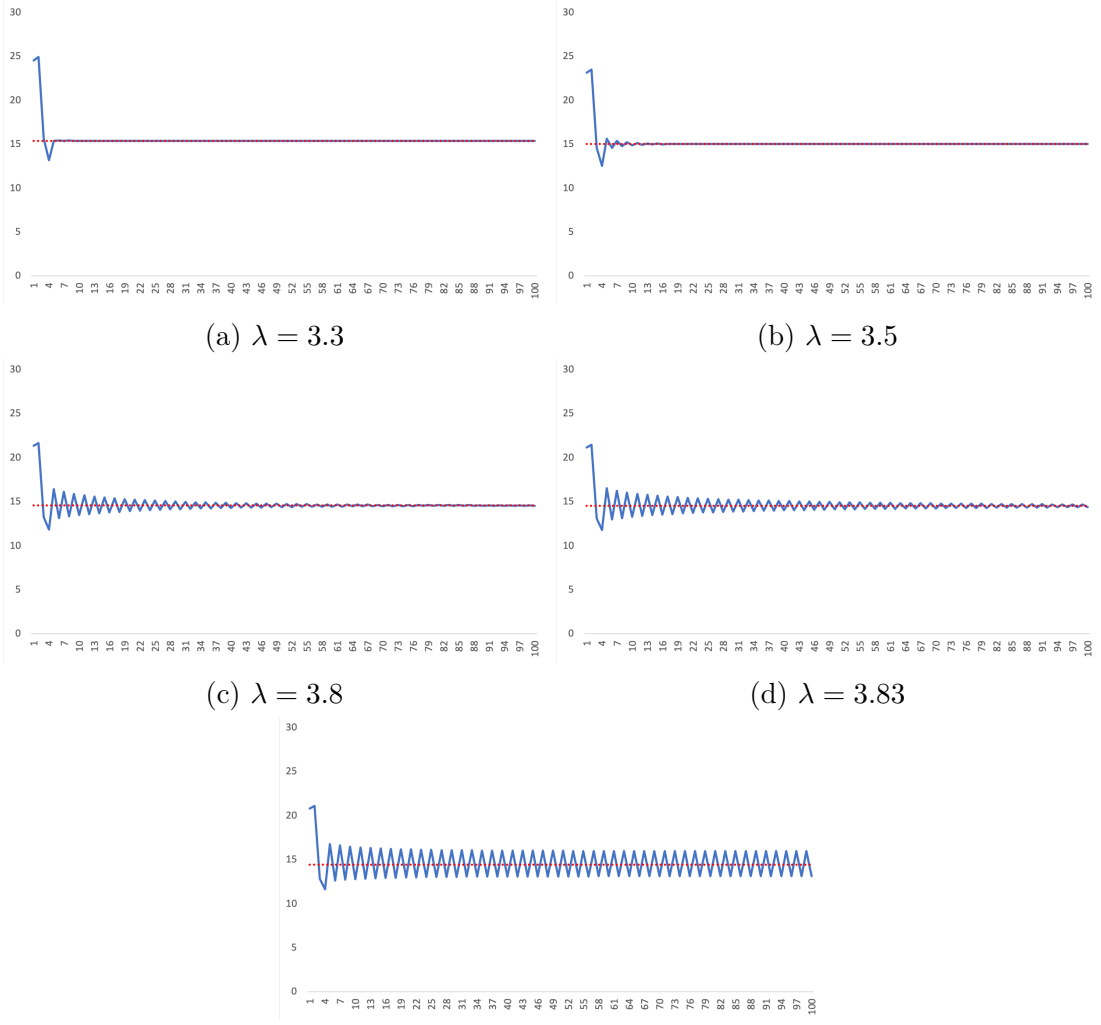
(c) $\lambda = 3.8$

(d) $\lambda = 3.83$

(e) $\lambda = 3.9$

Notes: Model simulations for each treatment (i.e., model parameter λ) under the assumption of backward perfect-foresight: $P^e_{t+1} = P_{t-1}$. The thick lines illustrates the simulated price (in blue) and the dashed red line indicates the steady state.

Figure 2.4: Dynamics under backward perfect-foresight: $P^e_{t+1} = P_{t-1}$



(a) $\lambda = 3.3$ (b) $\lambda = 3.5$

(c) $\lambda = 3.8$ (d) $\lambda = 3.83$

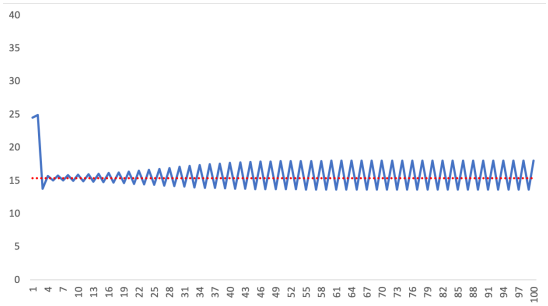
(e) $\lambda = 3.9$

Notes: Model simulations for each treatment (i.e., model parameter λ) under the assumption of adaptive expectations: $P^e_{t+1} = 0.3P^e_{t-1} + 0.7P_{t-1}$. The thick lines illustrates the simulated price (in blue) and the dashed red line indicates the steady state.

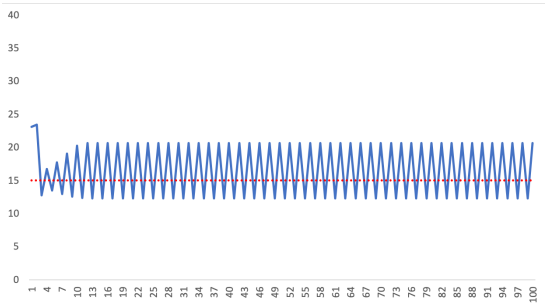
Figure 2.5: Dynamics under adaptive expectations: $P^e_{t+1} = 0.3P^e_{t-1} + 0.7P_{t-1}$

We focus on adaptive expectations of the form $P^e_{t+1} = (1-\beta)P^e_{t-1} + \beta P_{t-1}$. Adaptive expectations with a coefficient β smaller than 0.7 result in convergence to the steady state for every treatment (see Figure 2.5). Adaptive expectations with a coefficient β larger than 0.8 produce stable price dynamics with convergence to the steady state or to the two-cycle (see Figure 2.6). Higher-order cycles or complex behavior do not emerge. Adaptive expectations with a coefficient β equal to 0.5, leads to stable price dynamics as long as $\lambda < 3.8$, while chaotic dynamics occurs for larger λ-parameter values (Figure 2.7). It is worth noting that for high λ-values, chaotic dynamics imply sensitivity to initial conditions (namely, the value of initial forecasts) and, there exist initial states for which

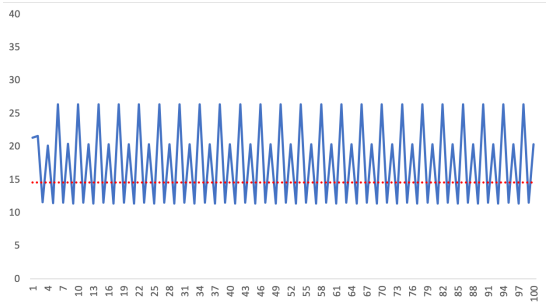
convergence towards stable dynamics, such as the steady state, occurs.



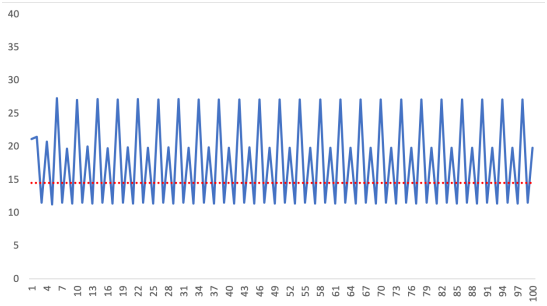
(a) $\lambda = 3.3$



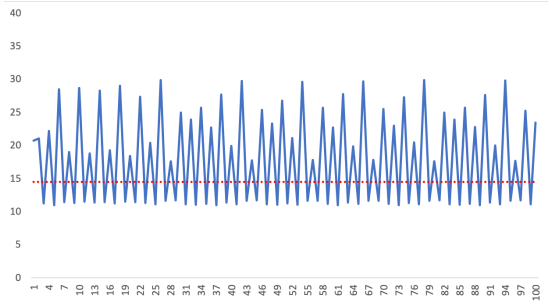
(b) $\lambda = 3.5$



(c) $\lambda = 3.8$



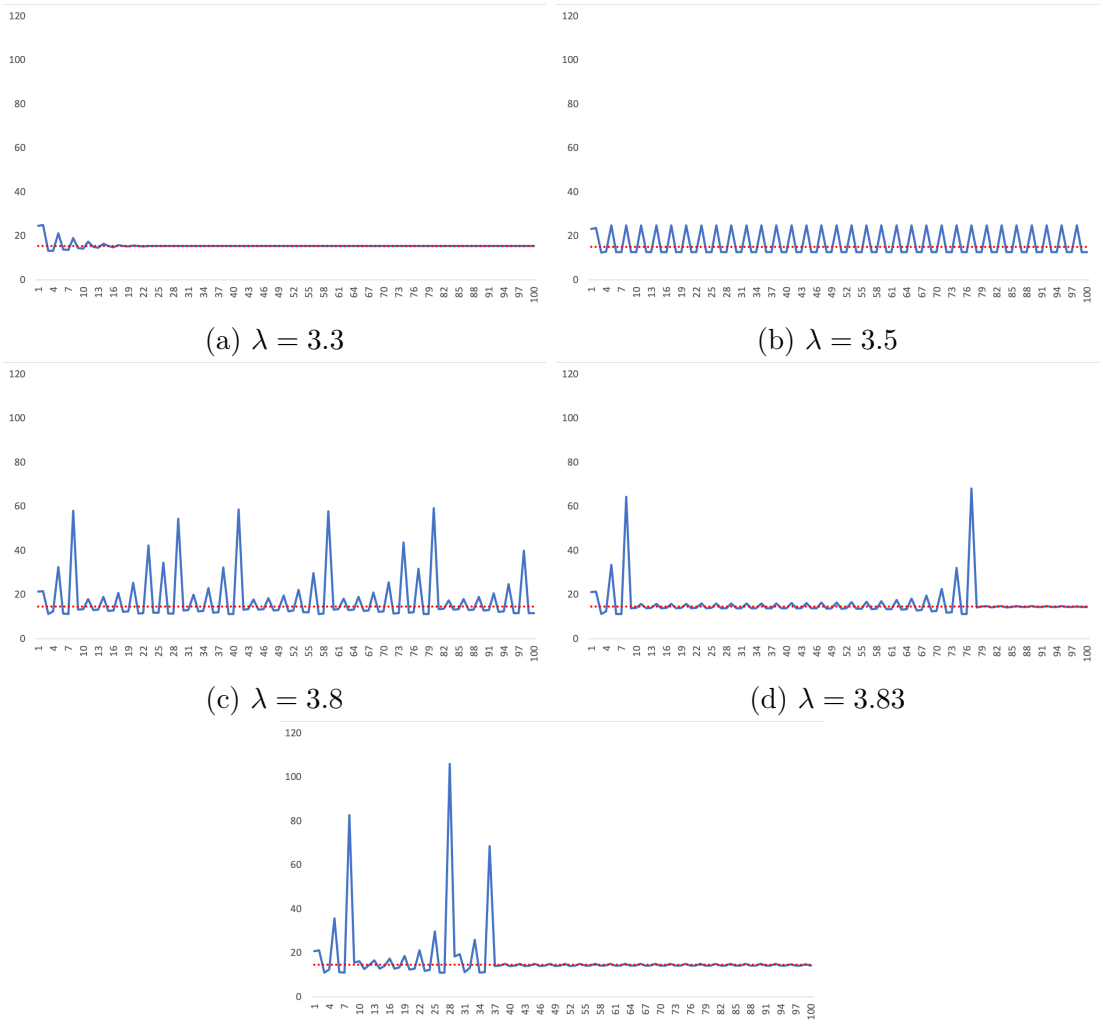
(d) $\lambda = 3.83$



(e) $\lambda = 3.9$

Notes: Model simulations for each treatment (i.e., model parameter λ) under the assumption of adaptive expectations: $P^e_{t+1} = 0.1P^e_{t-1} + 0.9P_{t-1}$. The thick lines illustrates the simulated price (in blue) and the dashed red line indicates the steady state.

Figure 2.6: Dynamics under adaptive expectations: $P^e_{t+1} = 0.1P^e_{t-1} + 0.9P_{t-1}$



(a) $\lambda = 3.3$

(b) $\lambda = 3.5$

(c) $\lambda = 3.8$

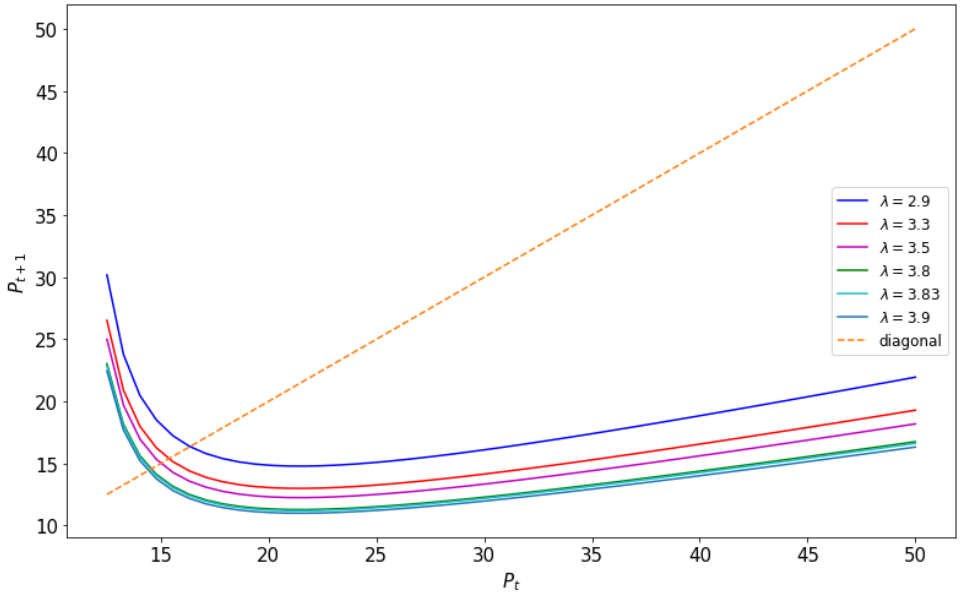
(d) $\lambda = 3.83$

(e) $\lambda = 3.9$

Notes: Model simulations for each treatment (i.e., model parameter λ) under the assumption of adaptive expectations: $P^e_{t+1} = 0.5P_{t-1} + 0.5P_{t-2}$. The thick lines illustrates the simulated price (in blue) and the dashed red line indicates the steady state.

Figure 2.7: Dynamics under adaptive expectations: $P^e_{t+1} = 0.5P_{t-1} + 0.5P_{t-2}$

Before turning to the lab implementation and results, we take a closer look at the model's dynamics—to develop intuition for the experimental results. Figure 2.8 shows the price map, that is the price in $t+1$ as a function of the price in t assuming backward-looking expectations, for each treatment (e.g., for each λ-value). The price in period t is illustrated on the x-axis and the price in period $t+1$ is on the y-axis. The dashed line is the diagonal line and illustrates the steady state. This Figure gives information about the so-called expectation feedback in the system. Figure 2.8 reveals that the price map is an "upside down" non-monotonic map. At first, the smaller the forecasts, the higher the realized prices, which corresponds to negative feedback, but for higher price values,



Notes: Model simulations of the price map for each treatment (i.e., model parameter λ) under the assumption of backward perfect-foresight: $P^e_{t+1} = P_{t-1}$. The dashed line is the diagonal line and illustrates the steady state.

Figure 2.8: Price map under backward dynamics

the relationship reverses and features positive feedback. Such a non-monotonic map illustrates the underlying complexity of the model.

2.4 Lab implementation

The experiment was programmed in oTree (Chen et al., 2016), and subjects were recruited from the pool of the CREED lab at the University of Amsterdam. Due to the COVID-19 social-distancing restrictions, we conducted all sessions online in June, July and September 2020, and in February and May 2021. Our experiment was one of the very first group experiments to be conducted fully remotely *via* the help of the zoom platform at CREED.[9] This novel setting for a group experiment unavoidably induced some challenges in 2020.

We conducted a total of 36 experimental groups (i.e., markets).[10] To be precise, we have data from 20 LtF experimental markets and from 16 LtO experimental markets, with a total of 259 participants. In general, one market consists of seven subjects. However,

[9] A complete description of the online procedures can be found in Appendix A.

[10] One additional session with two groups was conducted, but the experiment could not be finished—due to severe server problems. Subjects still received the participation fee of 5 euro and a payment for the number of rounds they participated in.

due to recruiting difficulties, four out of the 36 markets consist of six subjects only. The average duration was approximately two hours for the LtFE and three hours for the LtOE.

The subject pool consists of Bachelor and Masters students enrolled at the University of Amsterdam. The average age is 22.2, and 53.6% of the participants are women. The payment consisted of a 5 euro participation fee and performance-based payment. The average payoff in the LtFE was 27.4 euro, while the average payment in the LtOE sessions amounted to 33 euro. The balancing tables per treatment are presented in the Online Appendix F.

Online, the experimenter may have less control than in the physical lab, in particular regarding possible communication between subjects. However, in our experiment, subjects only have qualitative information about the experimental economy and the underlying equations are quite complicated, which limits the value of collusion. Another problem arises when participants do not pay attention to the task, even in the presence of monetary incentives, notably because of boredom or screen fatigue. Another related problem is the occurrence of dropouts, either due to inattention or poor internet connection. To address these issues, we implemented a timer of 90 seconds for every decision page. Whenever a decision page would time out, an additional 10-second timer would appear to check whether the subject was still active. In the event of a second time-out, the subject would be suspended, the total amount of money balances adjusted to the reduced number of players to keep the equilibrium price values constant and subsequent rounds would not incorporate their decisions any longer. This procedure ensures that the dropouts do not substantially slow down the experiment. A suspended subject could return to the experimental task in later rounds, in which case the initial configuration of the experiment would resume.[11] The pre-experiment quiz (Online Appendix D) was also adjusted to limit the interactions between the experimenter and the participants. Two successive wrong answers to a question triggers a multiple-choice version of the ques-

[11]To be specific, in our experiments, none of the subjects dropped out entirely. However, it occurred that subjects dropped out for some periods, but came back, and resumed the experiment. In total, subjects dropped out for 49 periods in LtF sessions (0.4% of the total number of decisions) and 136 periods in LtO (1.3% of the total number of decisions).

tion with four possible answer options. After a third wrong answer, the right answer is displayed in bold font.

We now turn to the experimental results. To ease the presentation, we first discuss the results of the LtFE, and then the robustness of the results using the LtOE.

3 Experimental results

Section 3.1 provides an overview of the main results of the LtF and the LtO experiments. Sections 3.2 and 3.3 provide further details on the LtF and LtO results; and are dedicated to highlight the treatment differences.

3.1 Overview of the experimental results

Figure 3.1 illustrates the price and price forecast dynamics in six experimental economies that are representative of the dynamics observed in the LtFE. And Figure 3.2 shows two experimental economies representative for the dynamics in the LtOE. The exhaustive collection of figures is deferred to Online Appendix G.1.

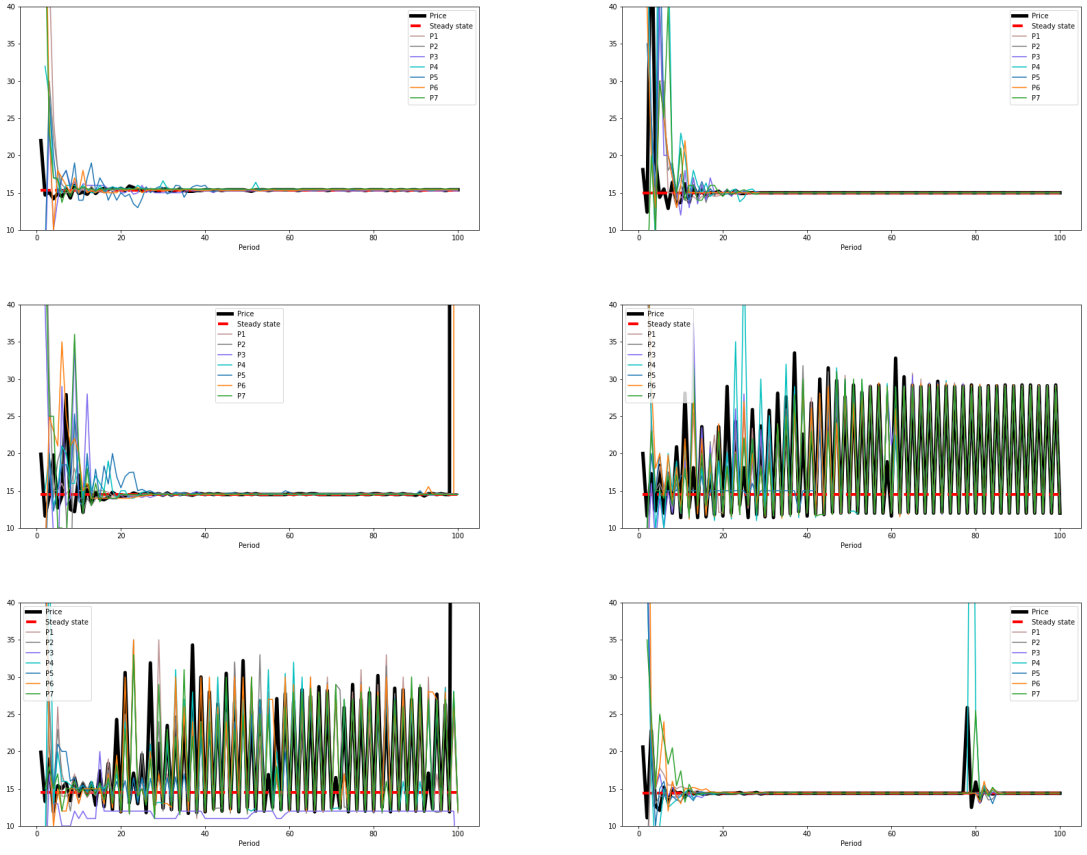
Let us first discuss the LtFEs. The price and the individual price forecasts strikingly converge towards the monetary steady state in the large majority of the cases (in 17 economies). In the remaining three economies, the price and the individual price forecasts converge towards the two-cycle. The convergence towards the two-cycle is only observed for the intermediary values of the λ-value, namely 3.8 and 3.83.

This result contrasts the observations from the experiment of Arifovic et al. (2019), where all economies converge towards the two-cycle once the underlying dynamics is chaotic in the backward perfect-foresight dynamics, and more strikingly so as the complexity parameter increases. In particular, once the three-cycle becomes stable, the experimental economies in Arifovic et al. (2019) invariably converge towards the two-cycle, while we report only a few instances of convergence towards the two-cycle. In our experiment, no instance of convergence towards any cycle emerges in the newly explored region with $\lambda = 3.9$, once the three-cycle has lost stability in the backward perfect-foresight dynam-

ics. Therefore, there is a non-monotonicity in the outcomes as complexity increases—we will focus on this discontinuity in Section 3.2.

In the LtOEs, the first important difference with respect to the dynamics in the LtFEs is the absence of convergence on the two-cycle, although this equilibrium Pareto-dominates the steady state. Second, in LtOEs, the distance of the price to the steady state remains larger than in the LtFE sessions. Nevertheless, by the end of the experiment the price in most LtO sessions is close to the equilibrium.

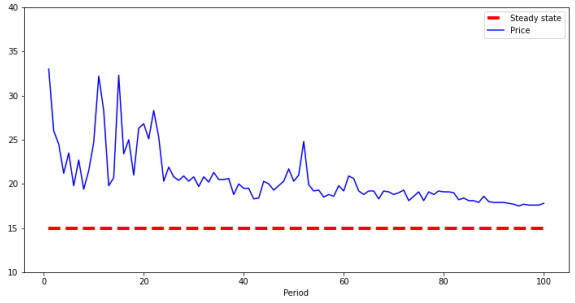
As for coordination of individual decisions, we can notice a rapid and high degree of price forecast coordination in the LtFEs. In the LtOE, the coordination between



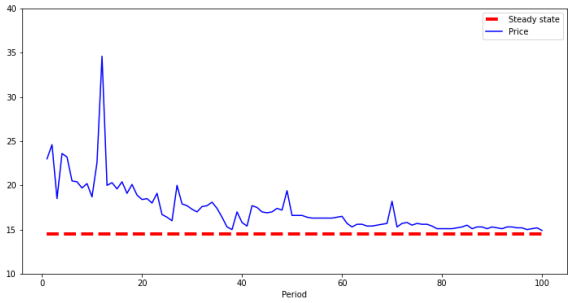
Notes: Each plot represents an experimental LtF economy where the thick lines report the price (in black) and its steady-state benchmark (in red) and each thinner line represents one subject' price forecasts. In most sessions, the price remains in the range $[0, 100]$. Price spikes (higher than 150) are almost always due to temporary dropouts which in turn push optimal output to 0 and the price to infinity. In such cases, the price is set to a fixed high number (750 in session 5 and 175 in all subsequent sessions). The parameter M was adjusted after session 14, and in the sessions with adjusted M dropouts do not have an effect on price.

Figure 3.1: Examples of price and forecast dynamics observed in the LtFEs

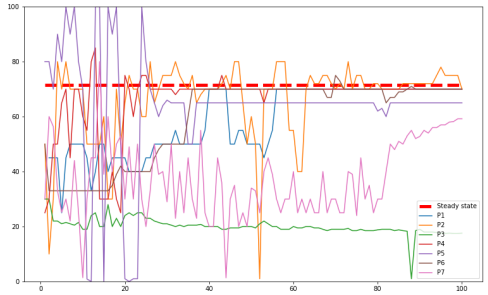
return forecasts appears higher than for savings decisions. Although the variety of savings decisions is especially high at the beginning, it remains substantial even at the end of the experiment. By contrast, return forecasts are less dispersed. Most of the spikes in return forecasts happen because subjects accidentally mix up the two tasks and enter savings decisions instead of return forecasts (on the basis of end-of-experiment questionnaire).



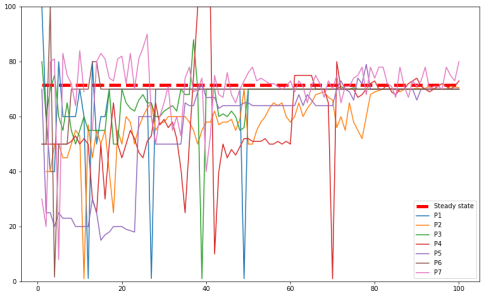
(a) Price in Group 1. $\lambda = 3.5$



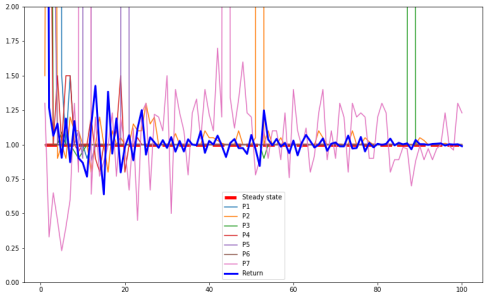
(b) Price in Group 6. $\lambda = 3.8$



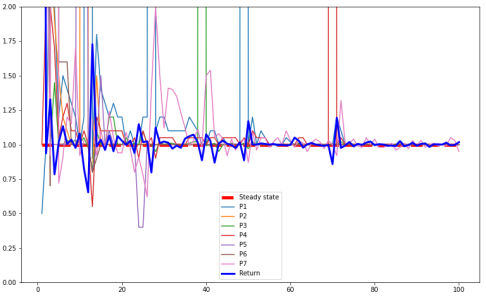
(c) Saving decisions in Group 1. $\lambda = 3.5$



(d) Saving decisions in Group 6. $\lambda = 3.8$



(e) Return forecasts in Group 1. $\lambda = 3.5$



(f) Return forecasts in Group 6. $\lambda = 3.8$

Figure 3.2: Examples of price, saving, and forecast dynamics in the LtOEs

These visual impressions are confirmed by Table 3.1 that provides descriptive statistics by treatment in both designs.[12] The first two rows refer to price aggregate convergence in the experimental economies. We use a definition of convergence as commonly employed in the literature, see e.g. Bao et al. (2013): we define an instance of ϵ-convergence in any period t if the price in this period lies within an ϵ-radius of its equilibrium value—whether it is a steady state or a cycle—and remains there until the end of the experiment. In what follows, we use $\epsilon = 5\%$ but results are robust to $\epsilon = 10\%$. The third and fourth rows of Table 3.1 show metrics for individual coordination. The third row reports a measure of coordination of price forecasts in the LtF design, and a measure for coordination of return forecasts in the LtO design. The fourth row reports such a metric for (implied) savings decisions in the LtF and elicited savings decisions in the LtO design. The smaller the reported numbers, the more similar are subjects' individual variables.

The distance of the price to equilibrium (first row) is smaller in sessions with $\lambda \leq 3.8$ than in sessions with $\lambda = 3.83$ and $\lambda = 3.9$. The distance to equilibrium is also several times smaller in the LtFE than in the LtOE and is largest in the LtO design with $\lambda = 3.83$. Treatments in the chaotic parameter regions take a particularly long time to converge (second row), especially in the LtO design. Again the treatment with $\lambda = 3.83$, whether implemented in the LtF or the LtO design, takes the longest time to converge.

We find the same pattern for individual coordination. It is higher for lower values of λ (in both LtFEs and LtOEs) and higher in LtFEs than in LtOEs (independent of looking at savings or forecasts).[13] In both designs, the treatment with $\lambda = 3.83$ stands out because it features the lowest degree of coordination between subjects.

Figures 3.3 and 3.4 illustrate coordination between individual decisions over time for the LtF and LtOEs, respectively. For readability, groups are split into two graphs. For the LtFE, we see that the relative standard deviation in all groups drops to almost zero by the end of the experiment. It decreases already substantially in the first 20 rounds and stays low until the final round 100. We also notice that the coordination of the forecasts

[12] We defer the statistical analysis of the cross-treatment comparisons to Sections 3.2 and 3.3 and do not discuss statistical significance in this section that aims to provide an overview.

[13] A similar conclusion is reached if we look at the time to coordinate, that is lower in the LtFEs than in the LtOEs and lower for lower λ-values.

Treatment	$\lambda = 3.3$	$\lambda = 3.5$		$\lambda = 3.8$		$\lambda = 3.83$		$\lambda = 3.9$	
Design	*LtFE*	*LtFE*	*LtO*	*LtFE*	*LtO*	*LtFE*	*LtO*	*LtFE*	*LtO*
Equilibrium	4 SS	4 SS	2 SS*	3 SS,1 2-c	2 SS*	2 SS, 2 2-c	1 SS*	4 SS	3 SS*
$ARDE$	0.3	0.1	26.4	0.6	27.2	5.6	38.3	1.6	12.2
TTC_{10}	4.5	36.3	84.8	38.3	92.5	67.5	98.3	51.0	97.0
RSD_f	0.4	0.9	28.3	0.6	18.3	7.8	56.7	2.8	22.8
RSD_s	0.24	0.7	25.4	0.4	24.0	2.6	30.7	2.6	25.4
EER_f	95.4	91.0	91.3	89.3	90.5	81.5	88.5	86.5	91.5
EER_s	-	-	85.5	-	86.8	-	77.8	-	83.8

Notes: all numbers are averages over all groups of a given treatment. * denotes approximate convergence, defined as when the average price stayed within 25% from the steady state in the last 25 rounds. Outlier price values due to subjects' drop-outs, typos or experimentation are excluded (0.95% of the total number of periods excluded). $ARDE$: average price is x% away from the equilibrium for the last 25 rounds. TTC: time to converge to an equilibrium and stay within 10% from it until the end of the experiment. In the case of no convergence, TTC is set to 100 periods. RSD_f: standard deviation of the forecasts divided by the average forecast over the last 25 rounds. RSD_s: standard deviation of the savings decisions divided by the average savings over the last 25 rounds. For the LtF we use savings derived from the first-order conditions of the model given price forecasts. EER_f – average payoff for the forecasting task relative to the maximum possible payoff. EER_s - average payoff for the savings task relative to the maximum possible payoff.

Table 3.1: Summary statistics of the LtFEs by treatment

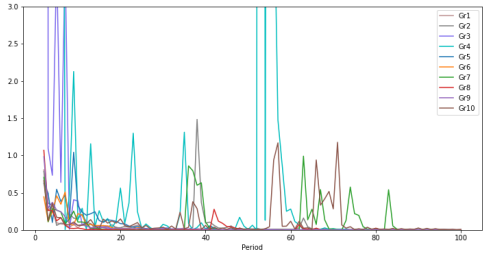
in the treatment with $\lambda = 3.83$ is highest. In the LtOE, coordination is lower in general (Figure 3.4). The coordination of the savings decisions is higher, but, unlike in the LtFE, the index of coordination does not drop to 0 by the end of the experiment.[14]

The last two rows of Table 3.1 look at payoff and earnings efficiency in the experiment. Efficiency is measured by comparing the earnings of subjects to the maximum possible amount of points for the corresponding task. Fast convergence to the equilibrium and a high degree of coordination naturally result in low forecast errors, high utility level and a high level of efficiency in almost all groups. Therefore, we remark that, again, earnings are higher for lower λ-values and lower in the LtFEs compared to the LtOEs. In the LtOEs, efficiency is higher for the return forecasting task than for the savings task. Again, the treatment with $\lambda = 3.83$ stands out with the lowest earnings in both designs.

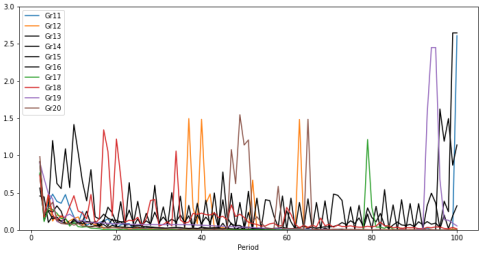
In the LtFEs, this treatment is where we observe several instances of convergence towards the two-cycle. Along a two-cycle, a slower convergence and more fluctuations than when the economy converges towards the steady state are associated with higher forecast errors and lower efficiency in these economies; see Figure 3.5 for a striking illustration.

Average forecast errors are reported in Figure 3.5. They drop to zero in all sessions

[14]In Figures 3.3 and 3.4 we observe occasional spikes. For the LtFE, the occasional spikes are mostly caused by temporary dropouts and subjects' experimentation with price forecasts. For the LtOE, many spikes in return forecasts are caused by mistakenly filling in savings decision instead of the return forecast.



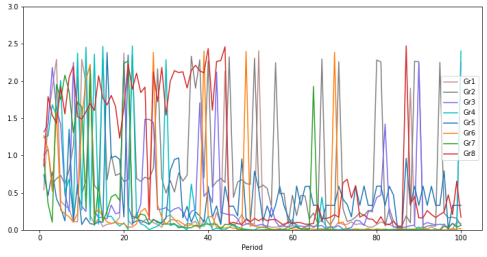
(a) Groups 1-10



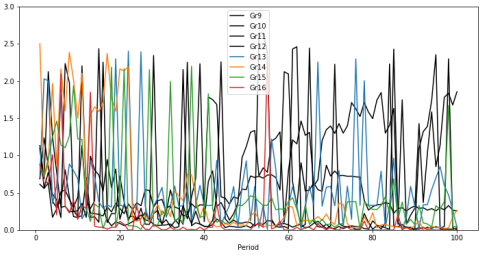
(b) Groups 11-20

Notes: Panel (a) and (b) illustrate the standard deviation of price forecasts divided by the mean forecast by group and over time.

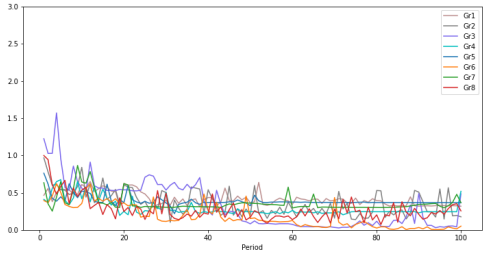
Figure 3.3: Standard deviation of price forecasts (RSTD) in the LtFE



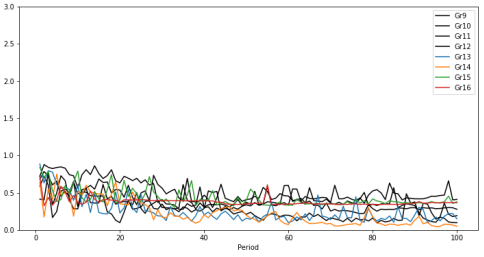
(a) Return forecasts (groups 1-8)



(b) Return forecasts (groups 9-16)



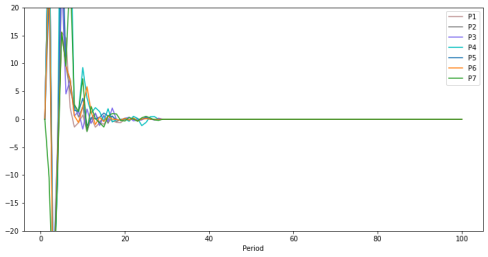
(c) Savings decisions (groups 1-8)



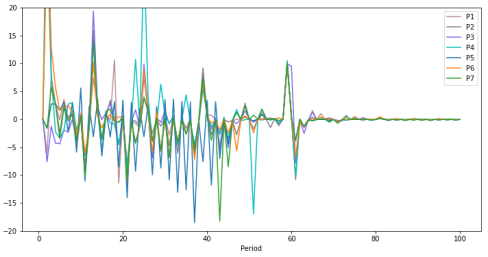
(d) Savings decisions (groups 9-16)

Notes: Panel (a) and (b) illustrate the standard deviation of return forecasts divided by the mean forecast by group and over time. Panel (c) and (d) illustrate the standard deviation of the savings decisions divided by the mean savings by group and over time.

Figure 3.4: Relative standard deviation of return forecasts and savings decisions (RSTD) in the LtOE

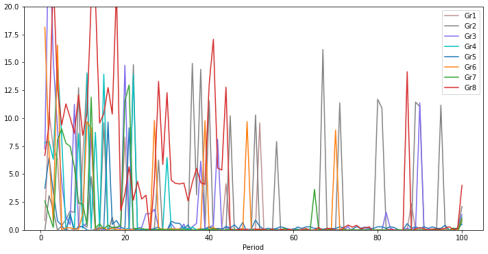


(a) Session that converged to the SS

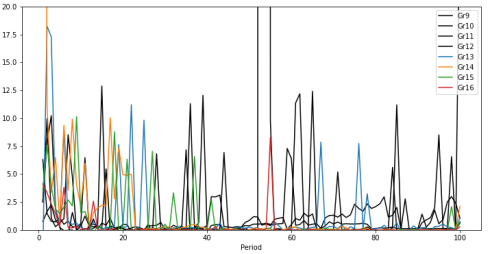


(b) Session that converged to the two-cycle

Figure 3.5: Average price forecast errors in the LtFEs

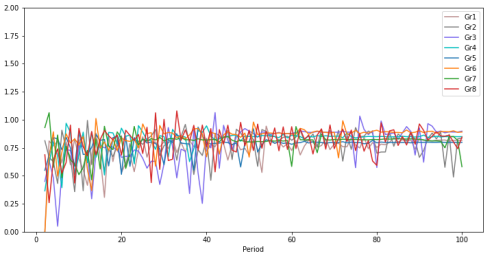


(a) Groups 1-8

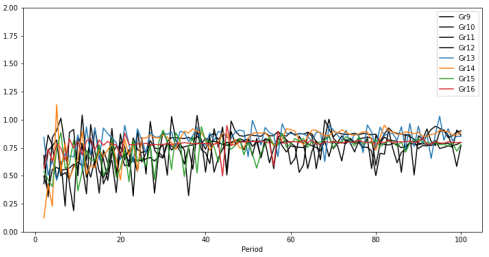


(b) Groups 9-16

Figure 3.6: Average return forecast errors in LtOEs



(a) Groups 1-8



(b) Groups 9-16

Figure 3.7: Average utility of agents in LtOEs

by the end of the experiment in the LtFE and in most of these sessions, zero average forecast error is achieved even before round 50. In LtOE, return forecasts are on average accurate (Figure 3.6). The majority of the spikes are caused by occasional mistakes of the subjects. The accuracy of the return forecasts is also reflected by high payoff efficiency for the forecasting task (Table 3.1). Finally, we can notice that the groups with $\lambda = 3.83$ have higher forecast errors and consequently lower average utility (Figure 3.7).

From this overview, we conclude that there exists a non-monotonicity in the experimental results as the complexity parameter λ increases. We now turn to a detailed

analysis of this pattern, first for the LtFEs, then for the LtOEs.

3.2 Non-monotonic dynamics in the LtFE

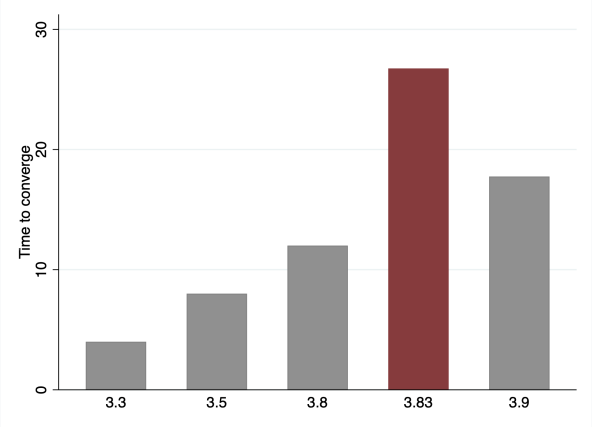
Figure 3.8 illustrates this non-monotonicity result for all indicators considered in Section 3.1. The Figure shows the average value of each indicator by treatment. The treatment $\lambda = 3.83$ is strikingly different from all other treatments—in a non-monotonic way.

In addition, we investigate treatment differences for two indicators measuring the cognitive load of the task. The first indicator measures the average decision time of the subjects in each round (Figure 3.8e), where a longer time indicates a higher cognitive load. The second indicator is an uncertainty index (Figure 3.8f). The uncertainty index is computed following the idea of Binder (2017): participants who round their forecasts are considered uncertain. Non-rounded forecasts correspond to an index equal to zero, forecasts rounded to 0.5 correspond to an index equal to one and the index for integer forecasts is equal to two. Hence, subjects are less sure about their forecasts, the larger the uncertainty index. For both indicators, the treatment with $\lambda = 3.83$ stands out (Figures 3.8e-3.8f). Again, we observe non-monotonic dynamics: In the treatment with $\lambda = 3.83$, subjects take more time to submit their forecasts and reveal a higher level of uncertainty—than in treatments with simpler dynamics (lower λ values) or in treatments with highly complex dynamics (with $\lambda = 3.9$).

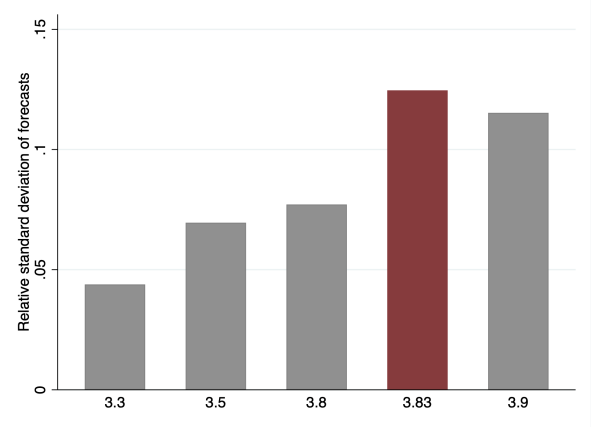
To test for the statistical significance of this non-monotonicity, we first pool all LtFE groups together and regress these aggregate statistics (i.e., the indicators) on a dummy variable that takes the value one for the treatment $\lambda = 3.83$ and zero otherwise. We estimate a linear probability model using the following specification:

$$Y_j = \beta_0 + \beta_1' Tr_{3.83} + \epsilon_j, \tag{3.1}$$

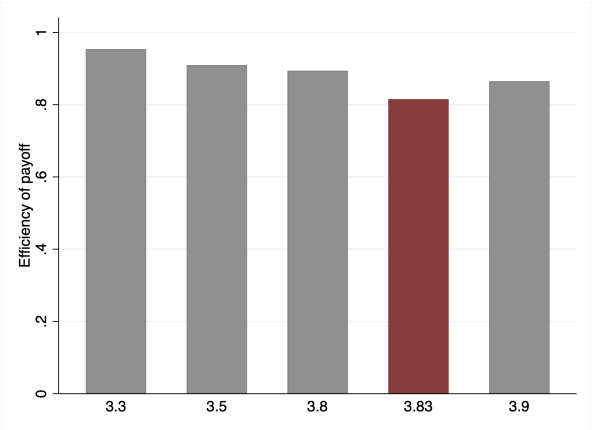
where Y_j denotes the indicator of group j (average over periods 1–100). $Tr_{3.83}$ denotes the dummy variable that equals one for the treatment $\lambda = 3.83$ and zero otherwise. The robust standard errors are denoted by ϵ_j.



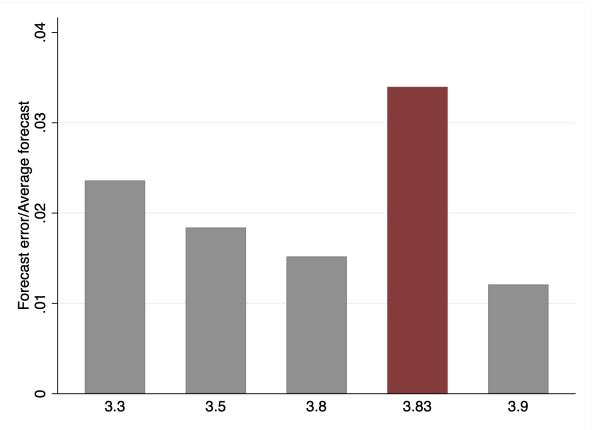
(a) Time to converge to equilibrium



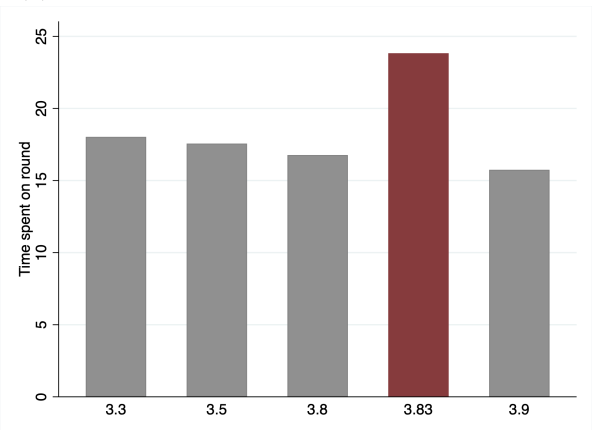
(b) Average relative SD of price forecasts



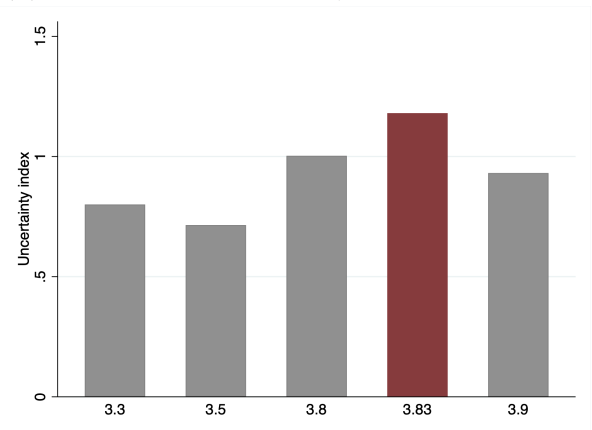
(c) Average forecast earnings efficiency ratio



(d) Average forecast error / the mean forecast



(e) Average forecasting decision time



(f) Uncertainty index

Notes: Panel (a) illustrates the time to converge to equilibrium, with the condition of staying within 10% from equilibrium. Panel (b) shows the average relative standard deviation of price forecasts. Panel (d) illustrates the average forecast error divided by the mean forecast. Panel (e) shows the average forecasting decision time, measured in seconds. Panel (f) shows the uncertainty index based on the rounding of forecasts.

Figure 3.8: Summary indicators of LtFEs by treatment

For some indicators, it is possible to compute the indicator also at the individual level. For these indicators, we use the following second specification:

$$Y_{i,j} = \beta_0 + \beta_1' Tr_{3.83} + F_j + \epsilon_{i,j}, \tag{3.2}$$

where $Y_{i,j}$ is the indicator for individual i, belonging to group j. $Tr_{3.83}$ denotes the treatment dummy that is equal to one for the treatment $\lambda = 3.83$ and zero otherwise. In addition, we control for group fixed effects, denoted by F_j. The standard errors are clustered at the group level and denoted by $\epsilon_{i,j}$.

	EER		TTC	RMSE	Uncertainty		Time on round
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\lambda = 3.83$	-0.0903*	-0.2191***	16.31*	37.31***	0.318**	0.720***	11.16*
	(0.0316)	(0.0352)	(6.623)	(8.697)	(0.114)	(0.182)	(3.452)
constant	0.906***	0.968***	10.44**	29.39***	0.862***	0.749***	14.39***
	(0.0169)	(0.0022)	(2.873)	(5.279)	(0.0657)	(0.121)	(1.905)
Group FE	-	+	-	+	-	+	+
N	20	139	20	140	20	140	70
R^2	0.255	0.742	0.261	0.657	0.225	0.324	0.264

Notes: Clustered standard errors in parentheses (at treatment level). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. EER - average payoff relative to the maximum possible payoff; TTC - time to converge to equilibrium and stay within 5% from it for at least 10 rounds; RMSE - square root of the mean squared forecast error; Uncertainty - uncertainty index based on rounding of forecasts; Time on round - average time spent on experimental round.

Table 3.2: Testing non-monotonicity in LtFE

Table 3.2 reports the most interesting regression results.[15] The treatment difference between the treatment $\lambda = 3.83$ and all other treatments, in other words, the non-monotonicity around the parameter value of 3.83 is statistically significant for all indicators of earnings (Cols. 1, 2 and 4), aggregate convergence (Col. 3) and cognitive load (Cols. 5, 6 and 7). Subjects make significantly higher forecast errors, take significantly more time to make a decision, and are significantly more uncertain about their forecasts in the treatment with $\lambda = 3.83$ than in treatments with any other λ-value considered. The sessions with $\lambda = 3.83$ also take significantly longer to converge to an equilibrium.

Now that we have established the statistical significance of the non-monotonicity result, we dig into the individual forecast time series to shed further light on this result.

[15] Table 3.2 reports all statistically significant results. Appendix G.2 provides a comprehensive overview of the regression results for all indicators tested.

To do so, we estimate the following common forecasting rules for each subject (see, e.g., Heemeijer et al. (2009a)).

- Naive expectations: agents set forecasts equal to the latest available price.

$$p^e_{t+1} = p_{t-1}$$

- Trend-following expectations: expectations are a function of the latest available price and the latest period-to-period change in price:

$$p^e_{t+1} = \beta p_{t-1} + \delta(p_{t-1} - p_{t-2}).$$

- Adaptive expectations: expectations are a weighted average of their own past forecast and the latest observable price (or, equivalently, expectations are adjusted towards the latest observable forecast error):

$$p^e_{t+1} = w p_{t-1} + (1-w)p^e_{t-1}, \quad \text{where} \quad 0 < w \leq 1.$$

- Sample average expectations: expectations are equal to the average of the last t past prices, where t is often restricted to a low value, such as 2:

$$p^e_{t+1} = \frac{1}{t-1}\sum_{j=1}^{t-1} p_j.$$

- Anchoring-and-adjustment heuristic (Tversky and Kahneman, 1974): subjects choose a weighted sum of a constant, the latest price and their last forecast as an anchor, and adjust this sum based on the latest price change:

$$P^e_{t+1} = \beta_1 P_{t-1} + \beta_2 P^e_t + \alpha + \gamma(P_{t-1} - P_{t-2}). \tag{3.3}$$

After the estimation of the rules, we choose the best-fitting rule for each subject based on the nested-model approach: we progressively add more variables into a regression and

choose the best model with a Likelihood-Ratio test.

Figure 3.9 shows the average composition of the different price forecasting rules by treatment. Trend-following, adaptive and anchoring and adjustment expectations are the most widely used strategies. In the majority of cases, the intercept in the trend-following rule is significant—which resembles an anchoring-and-adjustment behavior: a weighted average of some constant, possibly a perceived steady state, and the last price serves as an anchor for many subjects. Such a strategy is the most frequent strategy in the LtFEs.

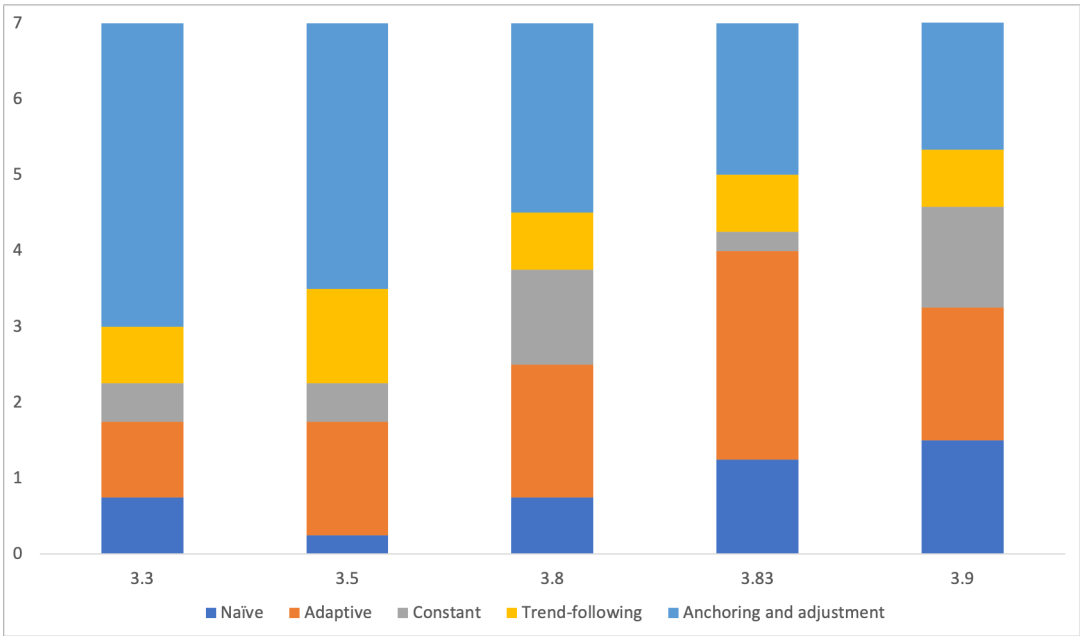


Figure 3.9: Composition of forecasting rules by treatment in the LtFEs (144 subjects)

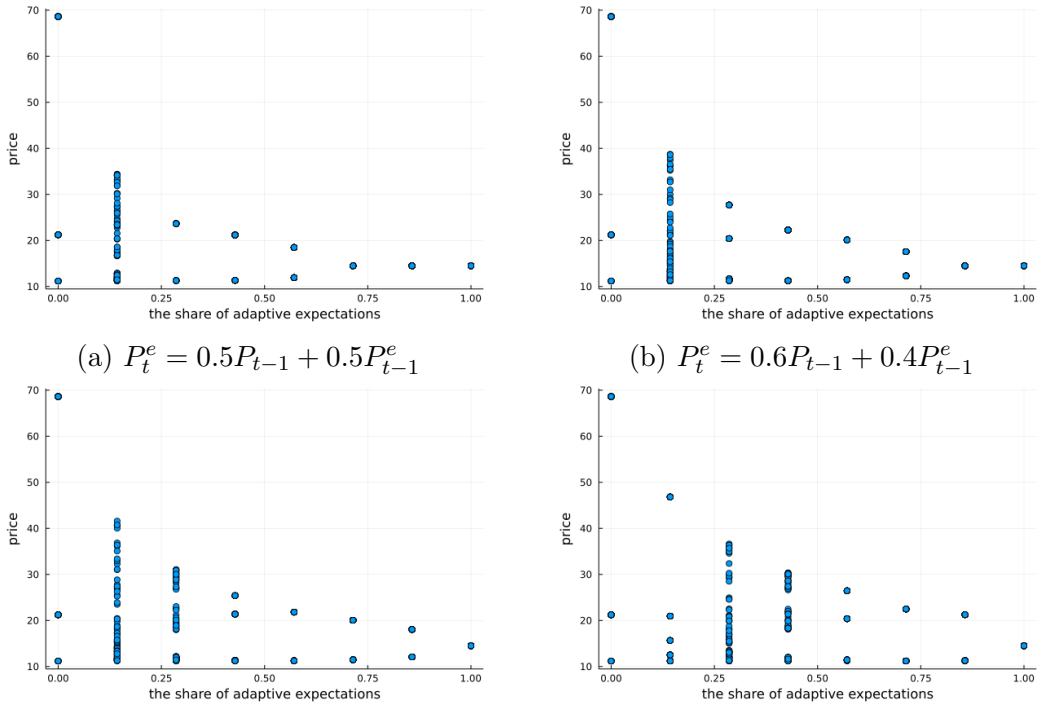
Most interestingly, the treatment with the parameter value $\lambda = 3.83$ clearly stands out. Recall that in this region, we observe the majority of the two-cycles and the composition of the forecasting rules for this parameter value is different: we observe the majority of adaptive expectations in the treatment with $\lambda = 3.83$. We tested if the share of subjects who use adaptive, trend-following and anchoring and adjustment heuristics is identical in the treatment with $\lambda = 3.83$ and in all other treatments using the Chi-squared test. The differences in composition are statistically different (p-value$= 0.0002$)[16]

This difference in forecasting heuristics, and in particular a higher share of adaptive

[16]To be precise, we present the largest p-value from the pairwise comparisons using Chi-squared test.

expectations in the treatment $\lambda = 3.83$ compared to all other treatments clearly relates to the observed non-monotonicity of the dynamics as complexity increases. To see how, we conduct simulations by varying the share of agents who use adaptive expectations in the group with $\lambda = 3.83$. The share of adaptive expectations observed in the $\lambda = 3.83$-treatment (about 40%, see Figure 3.9) leads to coordination on the two-cycle for the majority of the parameter values. Figure 3.10 shows the summary of the simulations in the form of bifurcation diagrams.

After establishing the non-monotonicity result in the LtFE, both at the aggregate and at the individual levels, we show in the next section that this result is robust in the LtOE.



(a) $P_t^e = 0.5P_{t-1} + 0.5P_{t-1}^e$ (b) $P_t^e = 0.6P_{t-1} + 0.4P_{t-1}^e$

(c) $P_t^e = 0.7P_{t-1} + 0.3P_{t-1}^e$ (d) $P_t^e = 0.8P_{t-1} + 0.2P_{t-1}^e$

Notes: Model simulations. All agents who do not use adaptive expectations are assumed to follow naive forecasting rule.

Figure 3.10: Bifurcation diagram with varying share of adaptive expectations and varying coefficients in adaptive expectations for $\lambda = 3.83$

3.3 Robustness of the non-monotonic result in the LtOEs

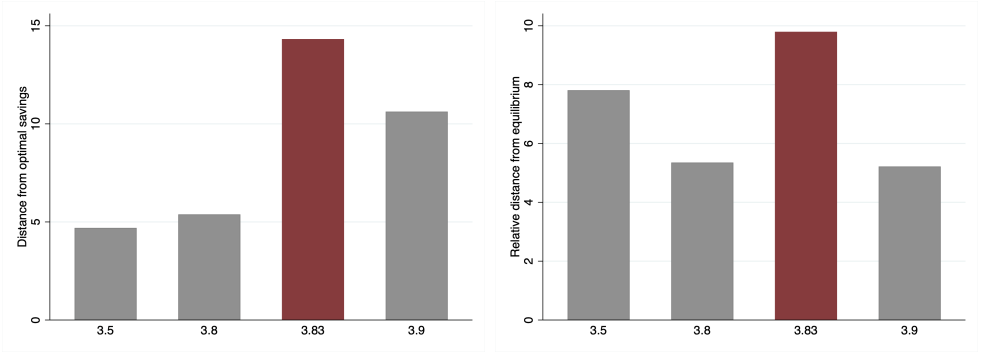
The non-monotonicity of behavior depending on the complexity is also present in the LtOE sessions. Figure 3.11 highlights this non-monotonicity for a series of aggregate and individual indicators. For the LtOE sessions and in contrast to the LtFE, we can calculate the optimal savings decision conditional on the return forecast. Recall that we elicit return forecasts, and hence, we may evaluate for each individual whether the savings decision is optimal—that the saving decision results in the maximal payoff, conditional on the return forecast submitted by the subject. Figure 3.11h shows the average difference between optimal and actual savings for each treatment. While savings decisions are on average downward-biased (i.e., subjects save too little), this downward bias is largest for the treatment with $\lambda = 3.83$.

Next, we investigate whether the treatment differences with $\lambda = 3.83$ are statistically significant for the following indicators: coordination of savings and return forecasts, price relative distance to steady state, efficiency, and relative forecast errors. Table 3.3 reports the regression results on the statistical significance of the differences observed with $\lambda = 3.83$. We find that the treatment with $\lambda = 3.83$ is significantly different compared to the other treatments. The price remains further away from its steady-state value, subjects make large forecast errors, earn less utility points when making savings decisions, and are less coordinated with $\lambda = 3.83$ than in simpler or highly complex parameter regions.

	D_o	RD	EER_s	FE_r	RSD_s
	(1)	(2)	(3)	(4)	(5)
$\lambda = 3.83$	7.423*	3.670***	-0.0758***	-0.7309***	0.0892*
(dummy)	(4.200)	(1.158)	(0.0243)	(0.2753)	(0.127)
Group FE	-	-	-	+	-
N	16	16	16	108	16
R^2	0.195	0.207	0.434	0.406	0.031

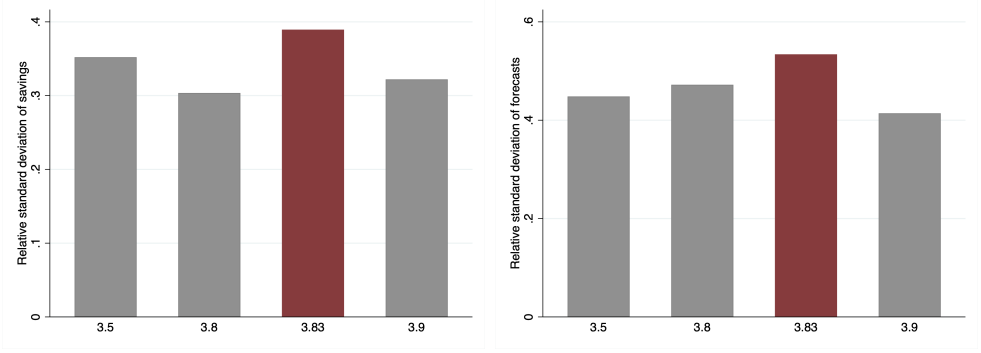
Notes: Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. D_o: distance of actual savings decisions from the optimal savings given subjects' return forecasts (negative number refers to savings too much and positive number to savings too little); EER_s: average payoff for the savings task relative to the maximum possible payoff; RSD_s - relative standard deviation of the savings decisions; RD: average relative distance to the equilibrium; FE_r: forecast error divided by the mean forecast.

Table 3.3: Testing non-monotonicity in the LtOE



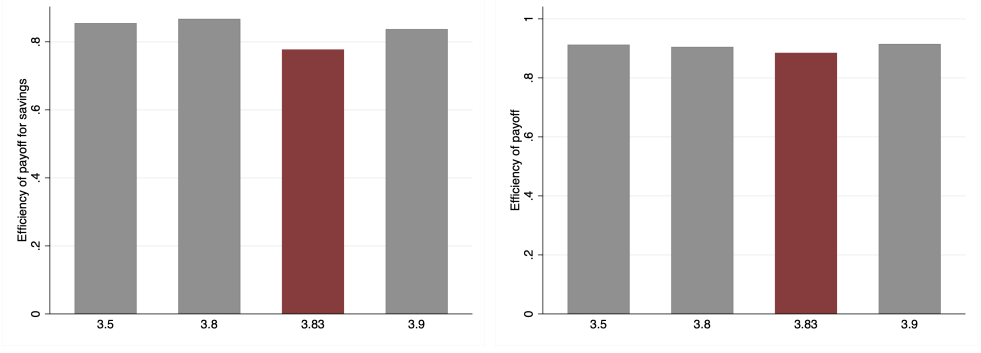
(a) Distance from optimal savings

(b) Relative distance of price from SS



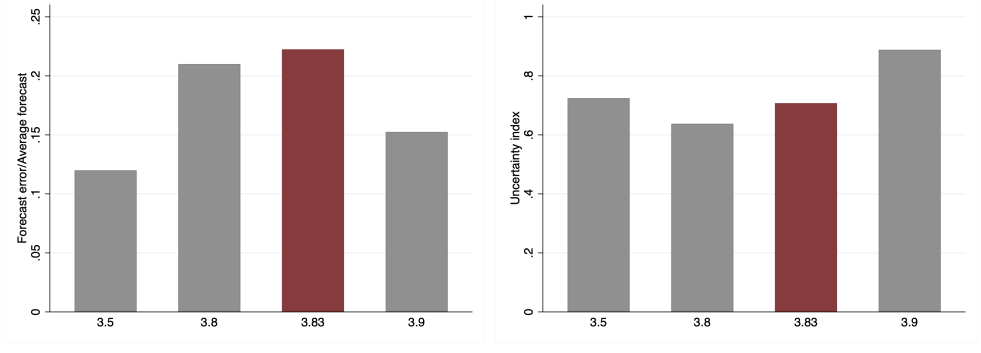
(c) Average relative SD for savings task

(d) Average relative SD for forecasts



(e) Saving earning efficiency ratio

(f) Forecast earning efficiency ratio



(g) Forecast error / mean forecast

(h) Uncertainty index

Notes: Panel (a) illustrates the average distance of savings from optimal savings. Panel (b) shows the relative distance of the price from its steady state value. Panel (c) illustrates the average relative standard deviation of the saving task. Panel (d) shows the average relative standard deviation of the forecasts. Panel (e) shows the earning efficiency ratio for savings decisions. Panel (f) shows the return forecasting earning efficiency ratio. Panel (g) illustrates the average return forecast error divided by the mean forecast. Panel (h) shows the uncertainty index based on the rounding of savings.

Figure 3.11: Summary indicators of LtOs by treatment

Finally, we estimate several benchmark heuristics using the time series on individual savings:

- Naive decisions: agents make savings decision equal to the past decision

$$y_{t+1} = y_t.$$

- Trend-following decisions: savings are a combination of the past savings and the latest change in savings:

$$y_{t+1} = \beta y_t + \delta(y_t - y_{t-1}).$$

- Adaptive decision-making: savings are equal to the weighted average of their own past savings decision and the average savings decision in the last period:

$$y_{t+1} = w y_{t-1} + (1-w)\bar{y}_{t-1}, \quad \text{where} \quad 0 < w \leq 1$$

- Average savings: savings are equal to the last average savings of the group:

$$y_{t+1} = \bar{y}_t.$$

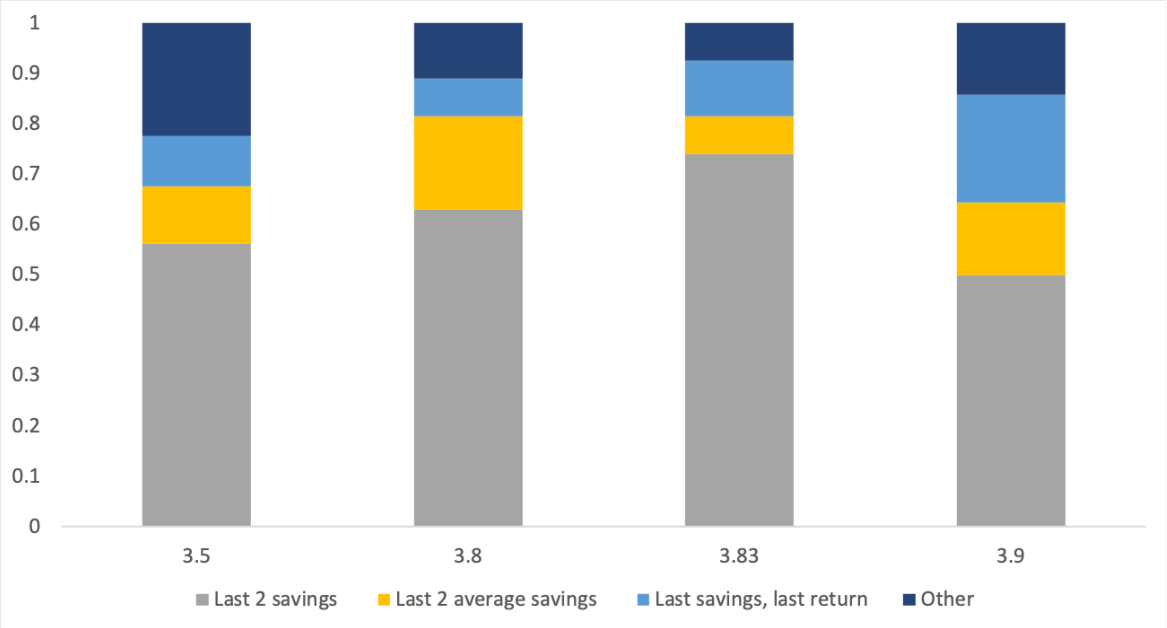
- Average trend-following: savings are equal to the weighted average of the last two average savings of the group:

$$y_{t+1} = \frac{1}{2}\sum_{j=1}^{2} \bar{y}_{t-j+1}.$$

- Sophisticated decision: savings are a function of their own last savings and the last observed return:

$$y_{t+1} = \beta_1 y_t + \beta_2 r_{t-1}.$$

We choose for each subject the best-fitting rule based on the Akaike information criterion. The results by treatment are presented in Figure 3.12. The most popular rule in all treatments is the average of the last two savings decisions, but even more so for the treatment with $\lambda = 3.83$ than for lower and higher λ-values. In this treatment, the fewest subjects rely on the group information, namely the average of the last two savings. It is worth noting that constant output decisions are observed only in the lowest and highest lambda range. The differences between the treatment with lambda equal to 3.83 and all other treatment are statistically significant (Chi-squared test: p-value = 0.000).



Notes: The fitting decision rules are chosen based on Akaike information criterion. Unlike in LtFE, the models in LtO are not nested, hence the LR test cannot be used.

Figure 3.12: Composition of decision rules by treatment (108 subjects)

4 Conclusion

This paper investigates group coordination and price dynamics in complex environments. We implement an overlapping generations model in the lab and conduct learning-to-forecast and learning-to-optimize experiments. We have chosen the model of Araujo and Maldonado (2000) because the output dynamics are given by the well-known chaotic quadratic map. Depending on the value of the utility function parameter, infinitely many perfect foresight equilibria can arise. In contrast to the related literature, we are particularly interested in the parameter region that allows for complex dynamics. We conduct five treatments of the learning-to-forecast experiment with different values of this parameter which correspond to the following theoretical predictions of price dynamics under the backward perfect foresight: convergence to the two-cycle, convergence to the 4-cycle, convergence to the 3-cycle, and chaotic behavior. In addition, we conduct four treatments using a learning-to-optimize design with parameter values corresponding to the following theoretical predictions: convergence to the 4-cycle, convergence to the 3-cycle, and chaotic behavior. This paper contributes mainly to two strands of literature. First, it contributes to the literature that studies equilibrium selection empirically using laboratory experiments. And second, this paper contributes to the experimental research on chaotic dynamics in general equilibrium models.

In all markets of the learning-to-forecast experiment, the price converges to a perfect foresight equilibrium. It is striking that convergence occurs on the simplest equilibria. In 17 out of 20 markets, the price converges to the monetary steady state, while in the remaining three sessions, the price converges to the two-cycle. In the learning-to-optimize experiment, the price approximately converges to the steady state in most sessions. These findings confirm the results of Arifovic et al. (2019) that the convergence occurs on the simplest equilibria.

Our paper documents an interesting and novel finding that the relationship between the complexity of chosen equilibria and model parameters is non-monotonic. The coordination on the two-cycles occurs only in the intermediate range of parameter values, while the coordination on the steady state happens for both low and high lambda values. This

non-monotonicity is also observed for convergence, forecasting errors, and subjects' uncertainty. The treatment with the parameter $\lambda = 3.83$ clearly stands out and differs from all other parameter values (i.e., treatments). The non-monotonicity in this parameter region is also present in the learning-to-optimize experiments.

One potential reason for the observed non-monotonic behavior could be differences in subjects' used forecasting strategies. In particular, we find that significantly more subjects use adaptive expectations in treatments of the intermediary parameter region—which corresponds to the treatments where most two-cycles are observed. Such behavior is not observed in Arifovic et al. (2019). We confirm our finding by conducting simulations where we vary the share of subjects using adaptive expectations in an experimental economy. Using the average share of adaptive expectations observed in our experiment with $\lambda = 3.83$ leads to coordination on the two-cycle in the simulations. In general, market convergence to the two-cycle is more likely if 3-5 (out of 7) subjects exhibit adaptive expectations.

Finally, our paper finds that the convergence in the learning-to-optimize experiment is much more challenging to achieve. In this experiment, we observe many suboptimal savings decisions and less efficient behavior. Also, the two-cycle is never observed, although it Pareto dominates the steady state in terms of payoff.

Our results show that people coordinate on the simplest possible equilibria—even in a highly complex and non-linear environment. This finding has important policy implications. In many cases, multiple equilibria co-exist. Policy-makers aiming to improve social welfare are advised to make the welfare-improving equilibria as salient and simple as possible to ease coordination.

Bibliography

Aliprantis, C. D. and C. R. Plott (1992): "Competitive equilibria in overlapping generations experiments," *Economic Theory*, 2, 389–426.

Araujo, A. P. and W. L. Maldonado (2000): "Ergodic chaos, learning and sunspot equilibrium," *Economic Theory*, 15, 163–184.

Arifovic, J. (1994): "Genetic algorithm learning and the cobweb model," *Journal of Economic dynamics and Control*, 18, 3–28.

——— (1995): "Genetic algorithms and inflationary economies," *Journal of Monetary Economics*, 36, 219–243.

——— (1998): "Stability of equilibria under genetic-algorithm adaptation: An analysis," *Macroeconomic Dynamics*, 2, 1–21.

Arifovic, J., C. Hommes, and I. Salle (2019): "Learning to believe in simple equilibria in a complex OLG economy-evidence from the lab," *Journal of Economic Theory*, 183, 106–182.

Arifovic, J. and J. Ledyard (2018): "Learning to alternate," *Experimental Economics*, 21, 692–721.

Bao, T., J. Duffy, and C. Hommes (2013): "Learning, forecasting and optimizing: An experimental study," *European Economic Review*, 61, 186–204.

Binder, C. C. (2017): "Measuring uncertainty based on rounding: New method and application to inflation expectations," *Journal of Monetary Economics*, 90, 1–12.

Brock, W. A. and C. H. Hommes (1997): "A rational route to randomness," *Econometrica: Journal of the Econometric Society*, 1059–1095.

Bullard, J. and J. Duffy (1995): "On learning and the stability of cycles," *Federal Reserve Bank of St. Louis Working Paper Series*.

——— (1998): "A model of learning and emulation with artificial adaptive agents," *Journal of Economic Dynamics and Control*, 22, 179–207.

Chen, D. L., M. Schonger, and C. Wickens (2016): "oTree—An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 9, 88–97.

DeCanio, S. J. (1979): "Rational expectations and learning from experience," *The quarterly journal of economics*, 93, 47–57.

Duffy, J. (2016): "Macroeconomics: a survey of laboratory research," *Handbook of experimental economics*, 2, 1–90.

Evans, G. (1985): "Expectational stability and the multiple equilibria problem in linear rational expectations models," *The Quarterly Journal of Economics*, 100, 1217–1233.

Evans, G. W. and S. Honkapohja (2001): "Learning and Expectations in Macroeconomics," .

——— (2009): *Learning and macroeconomics*, vol. 1, Annual Reviews.

Grandmont, J.-M. (1985): "On endogenous competitive business cycles," *Econometrica: Journal of the Econometric Society*, 995–1045.

Guesnerie, R. and M. Woodford (1991): "Stability of cycles with adaptive learning rules," in *Equilibrium Theory and Applications: Proceedings of the Sixth International Symposium in Economic Theory and Econometrics*, Cambridge University Press Cambridge, 111–134.

Heemeijer, P., C. Hommes, J. Sonnemans, and J. Tuinstra (2009a): "Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation," *Journal of Economic dynamics and control*, 33, 1052–1072.

Heemeijer, P., C. H. Hommes, J. Sonnemans, and J. Tuinstra (2009b): "Inflation expectations and stability in an overlapping generations experiment with money creation," .

Lim, S. S., E. C. Prescott, and S. Sunder (1994): "Stationary solution to the overlapping generations model of fiat money: experimental evidence," in *Experimental Economics*, Springer, 75–97.

Lucas Jr, R. E. (1978): "Asset prices in an exchange economy," *Econometrica: journal of the Econometric Society*, 1429–1445.

——— (1986): "Adaptive behavior and economic theory," *journal of Business*, S401–S426.

Marimon, R., S. E. Spear, and S. Sunder (1993): "Expectationally driven market volatility: an experimental study," *Journal of Economic Theory*, 61, 74–103.

Marimon, R. and S. Sunder (1993): "Indeterminacy of equilibria in a hyperinflationary world: experimental evidence," *Econometrica: Journal of the Econometric Society*, 1073–1107.

——— (1994): "Expectations and learning under alternative monetary regimes: an experimental approach," *Economic Theory*, 4, 131–162.

——— (1995): "Does a constant money growth rule help stabilize inflation?: experimental evidence," in *Carnegie-Rochester Conference Series on Public Policy*, Elsevier, vol. 43, 111–156.

Sargent, T. J. et al. (1993): "Bounded rationality in macroeconomics: The Arne Ryde memorial lectures," *OUP Catalogue*.

Tversky, A. and D. Kahneman (1974): "Judgment under uncertainty: Heuristics and biases," *science*, 185, 1124–1131.

A Appendix: Derivations of the OLG model

A.1 Finite number of agents

$P^e_{i,t+1}$ denotes the individual i's expectations in period t about the price level in $t+1$. Using the individuals' price expectations $P^e_{i,t+1}$, the compactness of the budget set (2.2), and the concavity of the utility function (2.1), the first-order condition of an individual i's maximization program reads as

$$\lambda \frac{P_t}{P^e_{i,t+1}} - \lambda \left(\frac{P_t}{P^e_{i,t+1}} \right)^2 y_{i,t} - 1 = 0. \tag{A.1}$$

Market clearing as described by (2.3) implies

$$\lambda \frac{M / \sum_{i=1}^N y_{i,t}}{P^e_{i,t+1}} - \lambda \left(\frac{M / \sum_{i=1}^N y_{i,t}}{P^e_{i,t+1}} \right)^2 y_{i,t} - 1 = 0, \tag{A.2}$$

rewriting yields that

$$\lambda M P^e_{i,t+1} \sum_{i=1}^N y_{i,t} - \lambda M^2 y_{i,t} - (P^e_{i,t+1} \sum_{i=1}^N y_{i,t})^2 = 0 \tag{A.3}$$

must hold for each individual $i \in \{1, ..., N\}$.

For both experiments (LtF and LtO), we impose that an individual price forecast is strictly positive; i.e., $P^e_{i,t+1} > 0 \quad \forall i, t$. Summing up all individual first-order conditions (A.3) gives

$$\lambda M \sum_{i=1}^N y_{i,t} (\sum_{i=1}^N P^e_{i,t+1}) - \lambda M^2 (\sum_{i=1}^N y_{i,t}) - (\sum_{i=1}^N (P^e_{i,t+1})^2)(\sum_{i=1}^N y_{i,t})^2 = 0. \tag{A.4}$$

Solving for aggregate output, denoted by $Y_t \equiv \sum_{i=1}^N y_{i,t}$, yields the following equation:

$$Y_t \left(\lambda M \sum_{i=1}^J P^e_{i,t+1} - \lambda M^2 - \sum_{i=1}^J (P^e_{i,t+1})^2) Y_t \right) = 0, \tag{A.5}$$

Equation A.5 admits two solutions for the temporary equilibrium of the model. The first one is the autartic equilibirum where production is always zero, i.e. $y_{i,t} = 0 \quad \forall i, t$. In addition, a monetary equilibrium exists as soon as at least one individual output is strictly positive. Note that we restrict output decisions to be strictly positive in the implementation of the LtO experiment. In any non-autartic temporary equilibrium, aggregate output is then equal to

$$Y_t = \frac{\lambda M (\sum_{i=1}^N P^e_{i,t+1}) - \lambda M^2}{(\sum_{i=1}^N (P^e_{i,t+1})^2)}. \tag{A.6}$$

Note that aggregate output in (A.6) is non-negative as soon as $\sum_{i=1}^N P^e_{i,t+1} > M$. The value of the parameter M is chosen to be equal to 1.5 in the calibration so that the condition is likely to hold for reasonable values of individual outputs. In the rare cases, when

aggregate output turns out negative, we set it to 0.

Plugging (A.6) in the individual first-order condition of the form (A.3) gives the individual output $y_{i,t}$ for each young individual i—expressed as a function of price forecasts only

$$y_{i,t} = P^e_{i,t+1} \frac{\frac{\sum_{i=1}^N P^e_{i,t+1}}{M} - 1}{\sum_{i=1}^N (P^e_{i,t+1})^2} \left(\lambda M - P^e_{i,t+1} \frac{\lambda M \sum_{i=1}^N P^e_{i,t+1} - \lambda M^2}{\sum_{i=1}^N (P^e_{i,t+1})^2} \right). \tag{A.7}$$

The aggregate price P_t is derived using the market-clearing condition and is given by

$$P_t = \frac{M}{\sum_{i=1}^N y_{i,t}} = \frac{M}{Y_t}. \tag{A.8}$$

A perfect-foresight equilibrium assumes that $P^e_{i,t+1} = P_{t+1}$, $\forall i, t$. Hence, all agents work and consume the same quantity such that $y_{i,t} = y_t$ $\forall i, t$. It follows

$$\lambda y_{t+1}(1 - y_{t+1}) = y_t, \tag{A.9}$$

which corresponds to the quadratic map.

The feedback from price expectations to realized prices is a complex non-monotonic function. To derive the conditions for a positive feedback, we differentiate the price function with respect to price expectations:

$$P_t = \frac{(P^e_{t+1})^2}{\lambda(P^e_{t+1} - M)} \tag{A.10}$$

$$\frac{\partial P_t}{\partial P^e_{t+1}} = \frac{P^e_{t+1}}{\lambda} \frac{(P^e_{t+1} - 2M)}{(P^e_{t+1} - M)^2} \tag{A.11}$$

where $P^e_{t+1} = \sum_{i=1}^N P^e_{i,t+1}$.

A.2 Stability analysis under various expectation schemes

Table A1 summarizes the results of the stability analysis and shows the equilibria—depending on the learning regime.

2* Treatment\ Stability concept	Forward perfect foresight	Backward perfect foresight	Strong E-stability	Weak E-stability	Adaptive Expectations
	Grandmont (1985)[a]		Evans and Honkapohja (2001)[b]		Guesnerie and Woodford (1991)[c]
$\lambda = 3.3$	SS	2-cycle	2-cycle	SS 2-cycle	SS 2-cycle
$\lambda = 3.5$	SS 2-cycle	4-cycle	4-cycle	SS 2-cycle 4-cycle	SS 2-cycle 4-cycle
$\lambda = 3.8$	SS All cycles except period 3	none	none	SS 2-cycle	SS 2-cycle All cycles except period 3 (if w is low enough)
$\lambda = 3.83$	SS All cycles except period 3	3-cycle	3-cycle	SS 2-cycle 3-cycle	SS 2-cycle 3-cycle All cycles (if w is low enough)
$\lambda = 3.9$	SS All cycles	none	none	SS 2-cycle	SS 2-cycle All cycles (if w is low enough)

Notes: SS denotes the monetary Steady State. w is the weight on the previous price realization in the adaptive expectations rule. The stability of any cycle under adaptive expectations is conditional on agents using an adaptive rule consistent with the cycle's periodicity.

Table A1: Stability analysis under different learning criteria

[a] Grandmont (1985)
[b] Evans and Honkapohja (2001)
[c] Guesnerie and Woodford (1991)

Online Supplementary Material

Learning in a Complex World: Insights from an OLG lab experiment

Cars Hommes, Stefanie J. Huber, Daria Minina, Isabelle Salle

Section A describes the implementation of the online experiment. Section B provides the instructions of the LtF experiment, and Section C the instructions for the LtO experiment. Section D shows how we make sure that the subjects understood the instructions. Section E reports the Questionnaire that we ask subjects to complete after the experiment ended. Section F reports demographics statistics across treatments. And Section G reports additional results.

A Online procedure

We use the CREED laboratory at the University of Amsterdam and send the invitation emails to 50–100 randomly selected students from the CREED pool. The email briefly explains the online procedure (including the use of the Zoom Software), the requirement to fill in an IBAN bank account number for receiving the payment, the expected duration of the experiment, and importantly that participants need to finish the experiment to receive the payment. The session is open for 5–7 extra participants to insure against no show-up.

One day before the experiment, we send a reminder email to the registered participants. Subjects receive the link to the Zoom meeting 15 minutes before the beginning of the experiment. After they open the link, they are assigened to the Zoom waiting room. All participants in the waiting room see the following message:

"Please wait, the host will let you in the virtual lab within the limit of the required number of participants (first-come first-served basis). If you can't enter this time, you will be able to register for another session soon!"

For the registration, we let subjects in—one by one. After letting them into the main Zoom room but before checking the ID, we rename subjects into "participant 1", "participant 2", etc. Then, we ask each subject to turn the video on, check their IDs, and send them back to the waiting room with the following message:

"Thank you! You will be participating in the experiment but please, you have to wait for a few more minutes until all participants have been registered so I'm putting you back now into the waiting room. I'll let you in once I start the experiment."

After the required number of participants is registered, these subjects are moved back to the main Zoom room, and the remaining participants in the waiting room are sent

home with the show-up fee of 7 euro and the following message:

"I'm sorry: the groups are full and there aren't enough participants for an additional one! I have to send you away but feel free to register for another upcoming session! Thank you for your participation!"

In the main room, the video is turned off, and the subjects are muted. Communication between participants is disabled. The room is locked so that no other participant can join. All the participants see the message:

"Welcome! I will now send you a link to the experiment in this chat box. Each link is private and anonymous. You may open it in any browser but Google Chrome is preferred. Please open the link by clicking on it and start reading the instructions at your own pace. After the instructions, there is a quiz. Once everybody has answered the questions correctly, the experiment will start. Good luck!"

Participants can ask questions through the "Raise your hand" Zoom option or in the private chat. The link to the experiment are sent via the private chat to each subject separately.

After the main experimental task is over, subjects receive the following message in the Zoom meeting:

"The experiment is now over. Once you have filled in the end-questionnaire and provided us with your International Bank Account Number (IBAN), you can leave the meeting. Thank you once again for your participation!"

After filling in the questionnaire, subjects leave the Zoom meeting. The payment is made to the subjects' accounts by the financial administration of the University of Amsterdam. The anonymity is preserved as experimenters know only the IBAN number but not the participant's name.

B Instructions: Learning-to-forecast Experiment

Welcome

Welcome! The experiment is anonymous, the data from your choices and information about your payment will only be participant number, not to your name. If you follow these instructions carefully, you can earn a considerable amount of money. Your earnings will be transferred to your account right after the experiment. We will ask you to fill in your IBAN number before the experiment begins in case of technical difficulties on our side Before the payment, you will also be asked to fill out a short questionnaire. During the experiment, you can use scratch paper and a calculator if you feel the need to do so. Before starting the experiment, you have to answer 5 questions at the end of the instructions section to make sure that you understand your role in the experiment. Each participant has the same role, and the rules are the same for all participants. From now until the end of the experiment, you are not allowed to communicate with other participants. Please address your questions to us by using the zoom interface.

General information about the experimental market

You participate in a market, in which individuals trade chips at a given price in each period. You are a Professional Forecaster, and you have to predict the price of the chips in the next period. In every period, two generations of individuals – the young and the old – trade chips with each other. Imagine that a period in this economy represents a generation: in each period, the young generation from the previous period becomes old, and a new young generation enters.

The young generation produces the chips and earns income by selling them to the old generation who consumes them. This income is saved till the next period, when the young generation becomes old, and spent fully to buy all the chips produced by the new young generation. Hence, young agents have to decide how many chips to produce. This decision depends on the price of the chips that will prevail **in the next period**, when they will be buying chips with their earnings. They need your forecasts of the chip price in the next period to make their production decision. The old generation does not need your advice as they simply spend all their savings at the prevailing price then. The savings of a **young** individual in money then equals:
savings in money = number of chips earned when young × current price of the chips

To sum up, in each period:
*- earnings of young individuals in money = number of chips produced *current price of the chips*
- amount of chips consumed by the old generation = earnings when young / current price of the chips (which means that whatever the price they spend all their earnings)

How the price of chips is determined

The price of chips is always determined in such a way that the chips produced by the young individuals can be exactly bought by the money of the old individuals. **As a professional forecaster, at the beginning of each period, you have to predict the price of the chips in the next period, and your prediction is then used by a young individual for making a savings decision in the current period.** In each period, there are 7 young individuals, each of them is advised by a forecaster. Each forecaster is played by a participant like you.

The price predictions of participants for the next period determines the amount of chips produced by the young generation. This means that your price prediction for the next period only influences the price in the current period, not the price in the next period. The price is a complicated function of your forecast and the forecasts of other financial advisors. In economies similar to this one, the price of chips has historically been between 1 and 100.

Information about your prediction task
The experiment lasts for 100 periods or generations. At the beginning of each period, you have to submit a forecast of the price of the chips in the next period. This means that you will observe the realized value of the price that you predicted in a given period only at the end of the next period. Your payoff in each period depends on your forecast error, that is the difference between your price forecast for a given period and the realized value of the price (we explain below how your payoff is exactly computed). You will then observe your forecast error and your corresponding payoff for a forecast made at the beginning of any period at the end of the next period.

The experiment starts at period 1. You are asked to submit your price forecast for the next period (period 2). Once all participants have submitted their price forecasts, all young individuals decide how many chips to produce and thus how many chips to save and sell to the old in period 1, and this determines the price of the chips in period 1. You are then entering period 2, you have to submit your price forecast for period 3. After all participants have submitted their price forecasts, young individuals decide how many chips to produce and sell in period 2, and the price of chips in period 2 is disclosed. You then observe your forecast error based on the forecast that you made in period 1 for period 2, and your corresponding score (payoff) for period 2. You are then entering period 3. This sequence of events repeats in each of the 100 periods of the experiment.

Computer interface
The computer interface is mainly self-explanatory. When making your forecast in any period, the following information will be displayed in the table (right panel of the computer screen) and the graph (bottom panel):

- The price level from the beginning of the experiment (period 1) up to the previous period;
- Your price forecasts from the beginning of the experiment up to the current period;
- Your forecast errors from the beginning of the experiment up to the current period;
- Your payoffs from the beginning of the experiment up to the previous period.

All these elements can be relevant for your forecasts, but it is up to you to determine how to use this information in order to make accurate forecasts. You have to enter your price predictions in the top left part of the screen (Figure A.1). When submitting your prediction, use a decimal point if necessary (not a comma). For example, if you want to submit a prediction of 2.5, type 2.5. Computer interface will be telling you when you can enter your prediction and when you have to wait for other participants.

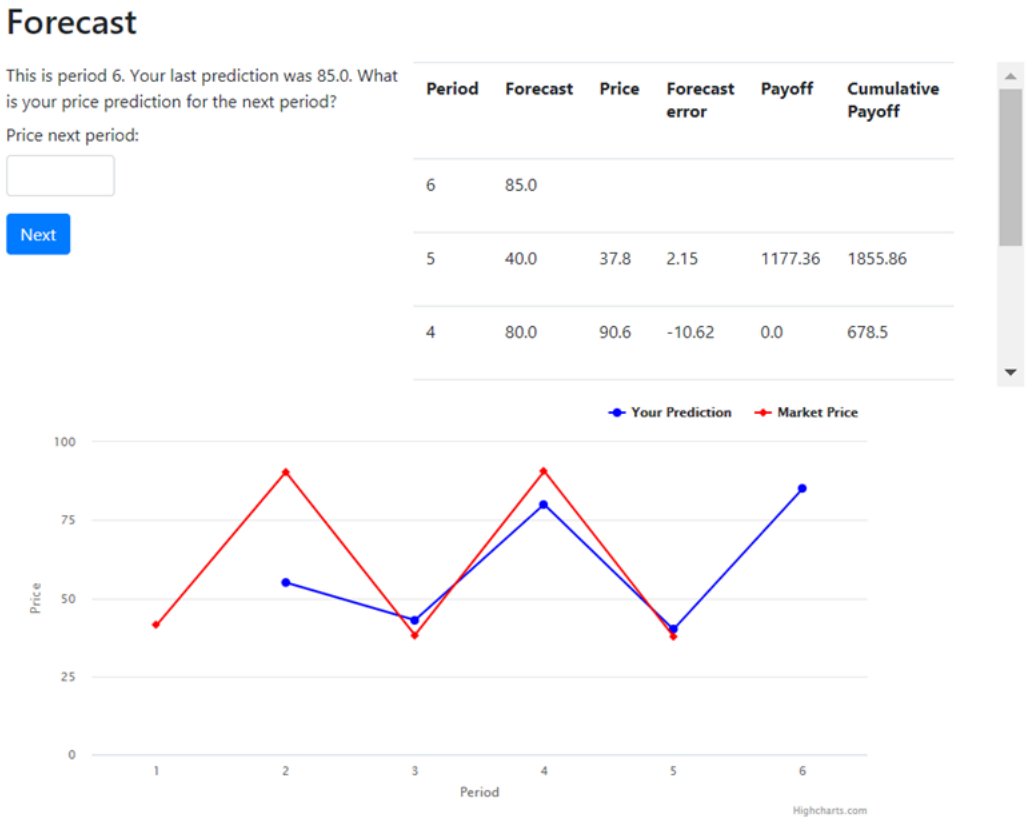


Figure A1: Computer screen

How payoff is computed

In each period, your payoff depends on the accuracy of your price forecast. The accuracy of your forecast is measured by the squared difference between your price forecasts and the realized values of price. Your payoff will be displayed on the computer screen in terms of points, and is computed as follows:

$$\text{Payoff}= max\{0, 1300 - 1300/49(\text{your forecast} - \text{actual price})^2\}$$

The payoff depending on your forecast error is also displayed in Figure A.2. Figure A.3 shows your payoff for different values of forecast errors.

If you forecast the price perfectly, your squared error is zero and you get 1300 points. This is the highest payoff that you can get in any period. The more accurate your forecast, the lower your squared forecast error, and the higher your payoff. If your forecast error is higher than 7, you get 0 point, and this is the minimum payoff you can get in any period.

Example: If your price forecast was 6 and the realized price is 5.7, your squared error is $(6 - 5.7)^2 = 0.09$, and your payoff is: $1300 - 1300/49 * 0.09 = 1297.6$ points.

If your prediction of the price was 32 and the realized price is 42, your squared error is $(32 - 42)^2 = 100$, and your payoff is: 0. So, you do not earn any points.

The sum of your payoffs over the different periods is shown in the top right of the screen. At the end of the experiment, your cumulative payoff over all 100 periods is computed, and converted into euro. For each 1300 points you make, you earn 0.35 euros. You will also receive a show-up fee of 5 euro on top of it. If you drop out before the end

of the experiment, the show-up fee will not be paid to you.

Please fill out the questionnaire on the next page. We will make sure that every participant has filled out the questionnaire with the correct answers for each of the five questions before starting the experiment.

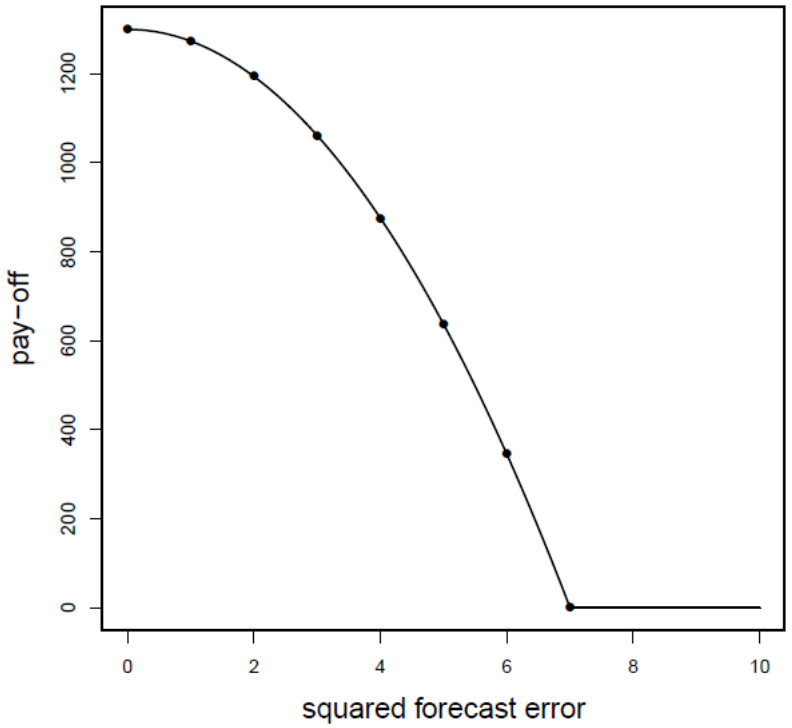


Figure A2: Payoff as function of forecast error

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

Figure A3: Payoff table

C Instructions: Learning-to-optimize Experiment

Welcome

Welcome! The experiment is anonymous, the data from your choices and information about your payment will only be linked to your participant number, not to your name. If you follow these instructions carefully, you can earn a considerable amount of money. Your earnings will be transferred to your account right after the experiment. We will ask you to fill in your IBAN number before the experiment begins in case of technical difficulties on our side. Before the payment, you will also be asked to fill out a short questionnaire. During the experiment, you can use scratch paper and a calculator if you feel the need to do so. Before starting the experiment, you have to answer 6 questions at the end of the instructions section to make sure that you understand your role in the experiment. Each participant has the same role, and the rules are the same for all participants. From now until the end of the experiment, you are not allowed to communicate with other participants. Please address your questions to us by using the zoom interface.

General information about the experimental market

You participate in a market, in which individuals trade chips at a given price in each period. In every period, two generations of individuals – the young and the old – trade chips with each other. Imagine that a period in this economy represents a generation: in each period, the young generation from the previous period becomes old, and a new young generation enters. The young generation consists of individuals of working age who work and produce chips. The old generation does not work anymore, and therefore consumes the income saved while being young.

The young generation produces the chips and earns income by selling them to the old generation who consumes them. This income is saved till the next period, when the young generation becomes old, and spent fully to buy all the chips produced by the new young generation. **Hence, young agents have to decide how many chips to produce which is equivalent to how much chips they want to save.**

You are a Professional Advisor working for the Professional Saving Advisor Bureau, and you have to decide in each period on the quantity of chips a young individual will save. Also, you have to forecast the return on savings in the next period. The calculation of the return is explained later on. In each period, there are seven young individuals, each of them follows the savings decision of a professional advisors. Each advisor is played by a participant like you.

To carry out the saved chips to the next period, the young individual converts these chips into money, by selling them to the old individuals. The quantity of money in the economy remains constant. The savings of a **young** individual in money then equals:

savings in money = number of chips earned (when young) × current price of the chips

The maximum amount of chips that can be produced by one young individual is 100.

The current price of chips is always determined in such a way that the chips produced by the young individuals can be exactly bought by the money of the old individuals. **The more chips the young individuals save, the lower the realized price of chips,**

and the more chips the old individuals can purchase with their savings and consume. As old individuals just consume the number of chips their savings can buy from the new young individuals, they do not need your savings advice. The consumption of chips of an old individual then equals:

consumption of chips when old = savings in money/ price of the chips when old

Your savings decision influences what the individual consumes when old in the next period. The **price of the chips in the current period** determines how much in money the young individual saves. The **price of chips in the next period** will determine how many chips the individual will be able to buy with his savings when old. Therefore, the consumption of chips when old also depends on the return on savings between the current period and the next period, defined as:

return on savings = current price (when young) /future price (when old)

The return on savings tells you how many chips the individual will be able to buy when old with one chip you choose to save for him when young.

You do not know yet the prices of the current and the next periods, so you do not know yet the return on savings when making your savings/production decision. Instead, **you should make a forecast of the return on savings. This forecast may also guide your savings decision in the current period.**

Information about your task as an advisor

The savings advisor bureau exists for 100 periods or generations. Each individual lives for two periods, produces and saves when young, and consumes when old. At the beginning of each period, you have to submit your savings/production decision and the forecast of the return on savings for a young individual for this period. **Your payoff for the savings task depends on the consumption of chips of this individual when old** (we explain below how your payoff is exactly computed). This means that **you will observe the quantity of chips this individual has consumed over his two-period life, and the corresponding payoff of your savings decision, only at the end of the next period, when he becomes old. Your payoff for the forecasting task depends on the accuracy of your return forecast.** At the time of the forecast, you do not know the current price and the price in the next period, so you will observe the forecast error and the payoff for the forecasting task 2 periods after the forecast is made.

The experiment starts at period 1. From period 1 to the end of the experiment (period 100), you have to make a savings/production decision and forecast the return on savings. Once all participants have entered their decisions and forecasts in period 1, all young people produce and save chips according to the advisors' decisions, all old individuals trade the money they earned in the young age against the saved chips of the new young and consume them. This determines the price of chips for period 1. Based on the initial price level, that usually ranges from 1 to 100, you observe the first return on savings. You are then entering period 2. After all participants have submitted their savings/production advice and return forecast for period 2, young individuals produce and save chips, old individuals buy and consume chips, and the realized price of chips for period 2 is disclosed, which determines the return on savings between period 1 and 2. You then observe the consumption of the young person you advised in period 1 in period 2 (when old), and

therefore the corresponding payoff of your savings decision made in period 1. You also observe the payoff you get for forecasting period 1 return. You are then entering period 3. This sequence of events takes place in each of the 100 periods of the experiment.

Computer interface
The computer interface is mainly self-explanatory. When making your savings/production decision and forecasting return in any period, the following information will be displayed in the table (right panel of the computer screen) and the graphs (bottom panel):

- The price level from the beginning of the experiment (period 1) up to the previous period;
- The return on savings from period 1 up to the previous period;
- The average savings decisions among the 7 advisers from the beginning of the experiment (period 1) up to the previous period;
- Your savings/production decisions from the beginning of the experiment (period 1) up to the previous period;
- Your return forecasts from the beginning of the experiment (period 1) up to the previous period;
- The consumption of chips when old of the individual you advised when young from period 2 up to the previous period;
- Your payoff from period 2 up to the previous period;
- Your payoff from forecasting return from period 2 up to the previous period.

The two plots (bottom panel) indicate your savings decisions together with the average decisions and the returns on savings with your return forecasts.

All these elements can be relevant for your savings/production decisions and return forecasts, but it is up to you to determine how to use this information in order to make optimal decisions and predictions.

You have to enter your savings/production decisions and return forecasts in the top left part of the screen (Figure below). When submitting your decisions and predictions, use a decimal point if necessary (not a comma). For example, if you want to submit a savings decision of 15.5 chips, type 15.5. Computer interface will be telling you when you can enter your decision and when you have to wait for other participants.

How the payoff is computed
In each period, **your savings payoff depends on the quality of your savings/ production decisions. The higher utility the individual you are advising gets from his/her consumption when old, the higher the quality of your savings/production decisions, and the higher your payoff. While consumption is positively related to the utility, the amount of work done while young is negatively related to the utility.** You do not need to calculate his/her utility, and hence your payoff yourself. **There is a payoff table in the instructions** (Figure below). According to your forecast of the return on savings (vertical axis), it shows the number of points that you can earn for a given savings decision. **You can use this**

payoff table to make your savings decision in the current period (columns) according to your forecast of the return on savings in the next period (rows). Note that the payoff table displays only some possible savings decisions and forecasts of the return on savings, but you can choose other ones. For instance, you do not need to choose between 90 or 100, but you may submit 91.2. Equally, you do not have to choose between 0.7 and 0.8 for your forecast of the return on savings, you may choose 0.72.

Your payoff for the forecasting task depends on your forecasting accuracy. The accuracy of your forecast is measured by the squared difference between your return forecasts and the realized values of return. Your payoff will be displayed on the computer screen in terms of points, and is computed as follows:

$$Payoff = max\{1300 - \frac{1300}{4}(your \quad forecast \quad error)^2, 0\}$$

If you forecast the return perfectly, your squared error is zero and you get 1300 points. This is the highest payoff that you can get in any period. The more accurate your forecast, the lower your squared forecast error, and the higher your payoff. If your forecast error is higher than 2, you get 0 point, and this is the minimum payoff you can get in any period. There is a payoff table with the instructions. It shows your payoff for different values of forecast errors.

Example 1 If you have advised a young person to save 90 chips, and the current price turns out to be 10 and the next period's price 20, the return on savings is $10/20 = 0.5$, this person consumes $0.5 \times 90 = 45$ when old, and your payoff is 230 points. For the same savings decision and current price, if the next period's price turns out to be 5, the return on savings is $10/5 = 2$ and this person consumes $2 \times 90 = 180$ when old, and your payoff is 65 points.

Example 2 If your return forecast was 6 and the realized price is 5.7, your squared error is $(6 - 5.7)^2 = 0.3^3 = 0.09$, and your payoff is

$$Your \quad earning = max\{1300 - \frac{1300}{4}0.09, 0\} = 1270.8 \quad points.$$

If your prediction of the return was 32 and the realized return is 42, your squared error is $(42 - 32)^2 = 1-^2 = 100$, and your payoff is

$$Your \quad earnings = max\{1300 - \frac{1300}{4}100, 0\} = 0,$$

and you do not earn any point.

At the end of the experiment, your cumulative payoffs for both tasks over all 100 periods are computed, and converted into euro. Each 500 points you make in savings task are converted into 0.2 euro, and each 800 points you make in forecasting task are converted into 0.2 euro. You will also receive a show-up fee of 5 euro on top of it. If you drop out before the end of the experiment, the show-up fee will not be paid to you.

You are going to be paid only for one task, either savings decision or fore-

casting. This will be determined randomly at the end of the experiment.

Please fill out the questionnaire on the next page. We will make sure that every participant has filled out the questionnaire with the correct answers for each of the six questions before starting the experiment.

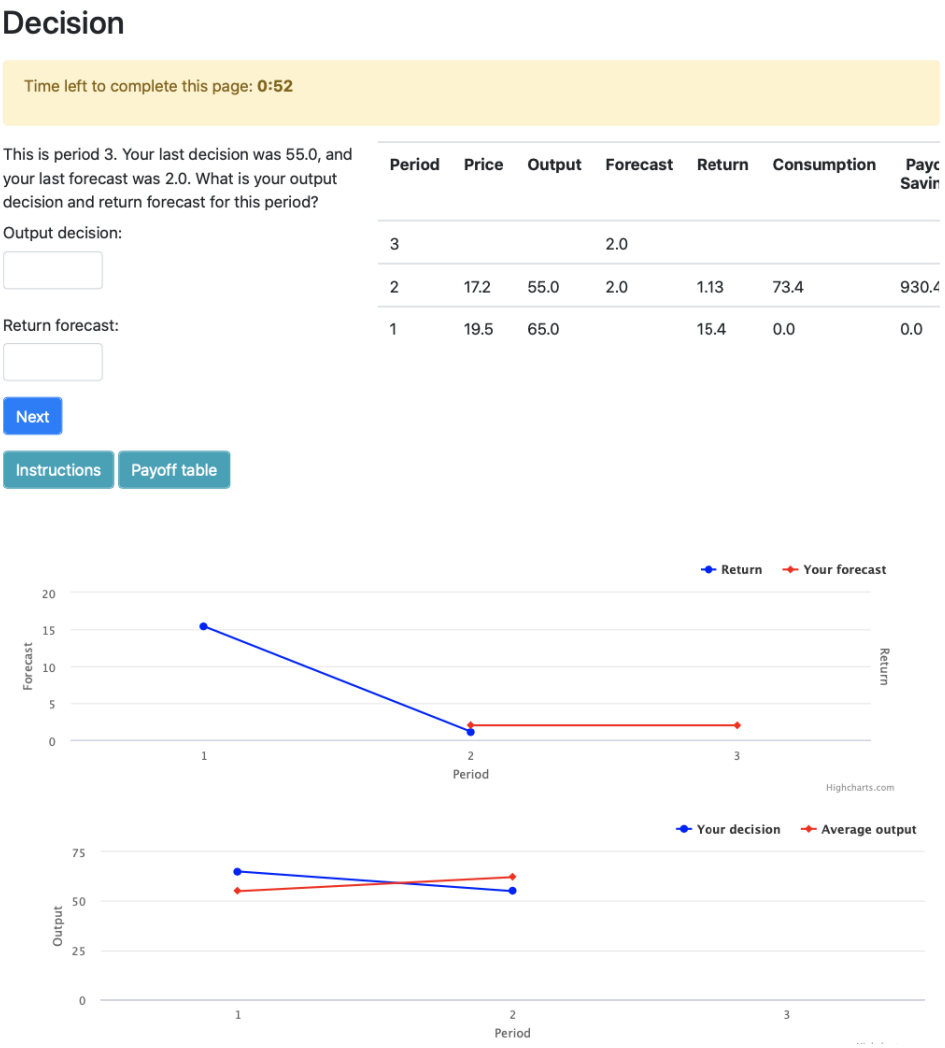


Figure A1: Computer screen

Your savings decision

Your return forecast	1	5	10	20	30	40	50	60	70	80	90	100
0.05	117	109	99	82	68	56	45	37	29	24	19	15
0.075	117	110	101	86	72	61	50	42	34	28	23	18
0.1	117	111	103	89	77	66	56	47	40	34	28	23
0.2	118	115	111	104	96	89	82	75	69	62	56	51
0.3	119	120	120	120	119	118	115	112	109	104	99	94
0.4	120	124	129	138	146	152	157	160	161	160	158	154
0.5	121	129	139	158	176	192	207	218	226	230	230	227
0.6	122	134	149	179	210	239	265	286	302	311	312	307
0.7	123	138	159	203	248	291	331	363	386	397	397	384
0.8	124	143	170	228	289	350	404	447	475	485	476	449
0.9	125	149	181	255	335	414	483	535	564	566	541	493
1	126	154	193	284	384	482	566	623	647	634	585	508
1.1	127	159	206	316	437	555	651	709	721	683	603	493
1.2	128	165	219	349	493	630	736	789	779	709	591	449
1.3	128	170	233	384	552	708	819	858	819	709	553	384
1.4	129	176	247	421	614	787	896	915	837	683	492	307
1.5	130	182	261	459	678	865	967	955	832	634	415	227
1.6	131	188	277	500	744	941	1029	977	804	566	331	154
1.7	132	194	292	542	811	1014	1079	980	756	485	248	94
1.8	133	201	309	586	878	1082	1116	964	691	397	173	51
1.9	134	207	325	631	946	1144	1138	930	612	311	111	23
2	135	214	343	678	1013	1199	1146	879	526	230	65	9
2.5	140	249	438	925	1314	1333	967	476	137	16	0	0
3	145	287	547	1181	1506	1199	566	127	7	0	0	0
4	156	376	798	1614	1394	482	36	0	0	0	0	0
5	167	478	1082	1785	788	47	0	0	0	0	0	0
6	179	595	1373	1614	235	0	0	0	0	0	0	0
7	191	725	1644	1181	23	0	0	0	0	0	0	0
8	204	865	1863	678	0	0	0	0	0	0	0	0
9	217	1014	2006	284	0	0	0	0	0	0	0	0
10	231	1168	2056	76	0	0	0	0	0	0	0	0
15	310	1891	1082	0	0	0	0	0	0	0	0	0
20	404	2204	96	0	0	0	0	0	0	0	0	0

Figure A2: Payoff table. Savings task.

$Your\ payoff = max[0, 1300 - \frac{1300}{4}(your\ forecast\ error)^2]$			
800 points = 0.2 euro			
error	points	error	points
0	1300	1.05	942
0.05	1299	1.1	907
0.1	1297	1.15	870
0.15	1293	1.2	832
0.2	1287	1.25	792
0.25	1280	1.3	751
0.3	1271	1.35	708
0.35	1260	1.4	663
0.4	1248	1.45	617
0.45	1234	1.5	569
0.5	1219	1.55	519
0.55	1202	1.6	468
0.6	1183	1.65	415
0.65	1163	1.7	361
0.7	1141	1.75	305
0.75	1117	1.8	247
0.8	1092	1.85	188
0.85	1065	1.9	127
0.9	1037	1.95	64
0.95	1007	≥ 2	0
1	975		

Figure A3: Payoff table. Forecasting task

D Quiz: Comprehension Checks

Learning-to-forecast design

1. If you enter period 6, for which period are you asked to submit a price forecast?
   Answer: 7

2. If you enter a price prediction for period 10, which period's price will be influenced by your prediction?
   Answer: 9

3. Suppose that in a period, your prediction for the market price was 40, and the market price turns out to be 41, how many points do you earn in this period? (use the payoff table)
   Answer: 1273

4. Suppose that in a period, your prediction for the price was 10, and the price turns out to be 25, how many points do you earn in this period? (use the payoff table)
   Answer: 0

5. Suppose the total amount of chips sold by the young generation in period 2 is 5, and the total amount of chips sold in period 3 is 20. In which period will the price be the highest?
   Answer: period 2

Learning-to-optimize design

1. If you enter period 6, for which period are you asked to submit production/savings decision and return forecast?
   Answer: 6

2. If you enter savings/production decision for period 10, which period's price will be influenced by your decision?
   Answer: 10

3. If the total amount of chips saved by the young generation is 150, how many chips will the old generation consume?
   Answer: 150

4. Suppose that in a period 9, you advised to save 4 chips, the price of the chips was 30 in this period, and 10 in the next period (period 10). What is the return on savings between period 9 and period 10?
   Answer: 3

5. Suppose you forecast that the return on savings will be 9, how many chips should you advise to save? (use your payoff table)
   Answer: 10

6. Suppose the total amount of savings of the young generation in period 2 is 100, and the total amount of savings in period 3 is 200. In which period will the price be the highest?
   Answer: period 2

E End-of-experiment Questionnaire

Learning-to-forecast design

1. What is your age?
2. What is your gender?
   Options: male, female, other
3. What is your nationality?
4. What is your study field?
5. How clear did you find the instructions?
   Options: very clear, clear, understandable, fairly confusing, confusing, unclear
6. Have you participated in a similar experiment before?
   Options: yes, no, I don't know
7. What strategy did you use for forecasting asset price?
8. Do you feel your forecasts influenced price? If yes, in which way?
9. If you have other comments, write them here.

Learning-to-optimize design

1. What is your age?
2. What is your gender?
   Options: male, female, other
3. What is your nationality?
4. What is your study field?
5. How clear did you find the instructions?
   Options: very clear, clear, understandable, fairly confusing, confusing, unclear
6. Have you participated in a similar experiment before?
   Options: yes, no, I don't know
7. What strategy did you use in making savings decisions?
8. What strategy did you use in forecasting return?
9. Do you feel your decisions influenced price? If yes, in which way?
10. If you have other comments, write them here.

F Balancing tables across treatments

The Chi-squared test shows significant treatment differences in age for both LtO and LtFE treatments. The treatments are balanced in all other characteristics: gender, EU nationality, participation in similar type of experiments.

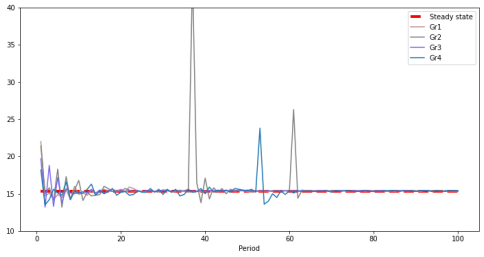
Treatment	$\lambda = 3.3$	$\lambda = 3.5$		$\lambda = 3.8$		$\lambda = 3.83$		$\lambda = 3.9$	
Design	LtF	LtF	LtO	LtF	LtO	LtF	LtO	LtF	LtO
The number of participants	28	28	26	28	27	28	27	28	28
Age (average)	22.0 (0.02)	21.0 (0.07)	25.0 (0.03)	21.8 (0.02)	21.8 (0.03)	21.8 (0.04)	21.9 (0.08)	22.0 (0.10)	23.0 (0.17)
Share of women	0.45 (0.11)	0.55 (0.28)	0.53 (0.33)	0.58 (0.11)	0.63 (0.09)	0.50 (0.28)	0.40 (0.08)	0.55 (0.42)	0.63 (0.08)
Share with EU nationality	0.35 (0.09)	0.65 (0.09)	0.68 (0.34)	0.63 (0.09)	0.80 (0.34)	0.55 (0.09)	0.70 (0.21)	0.45 (0.25)	0.85 (0.21)
Share of experienced participants	0.63 (0.18)	0.60 (0.18)	0.50 (0.12)	0.60 (0.06)	0.45 (0.41)	0.73 (0.06)	0.63 (0.17)	0.68 (0.11)	0.45 (0.12)
Share of participants who find instructions at least understandable	0.89 (0.28)	0.93 (0.28)	0.73 (0.68)	0.89 (0.55)	0.74 (0.46)	0.93 (0.55)	0.78 (0.46)	0.89 (0.30)	0.79 (0.56)

Notes: All the characteristics displayed in the table are computed as the average for all groups in the treatment. The largest p-value of the Chi-squared test for the pairwise comparisons between treatments is displayed in the brackets. The p-values larger than 0.05 mean no significant differences between treatments at the 95% confidence level.

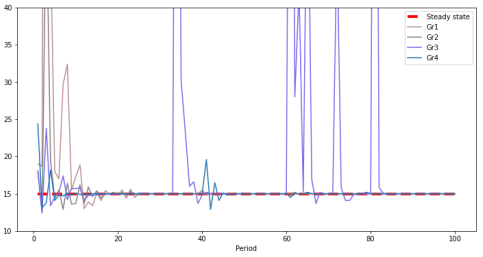
Table A1: Balancing table of LtF and LtO by treatment

G Additional results

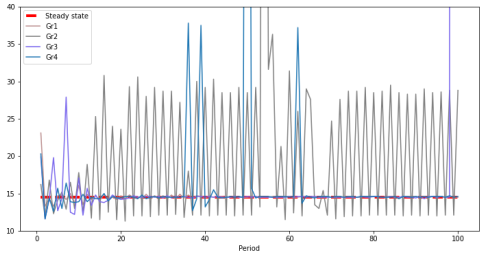
G.1 Price dynamics in each experimental economy



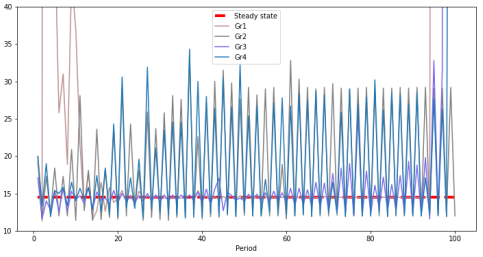
(a) $\lambda = 3.3$



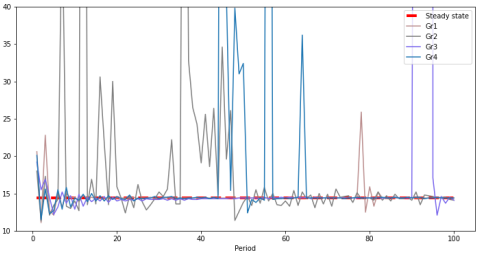
(b) $\lambda = 3.5$



(c) $\lambda = 3.8$

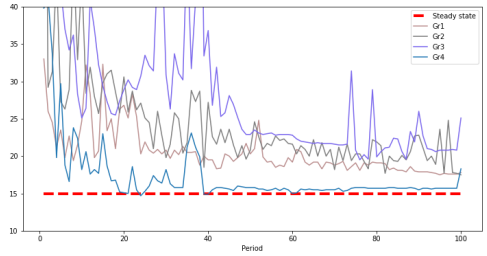


(d) $\lambda = 3.83$

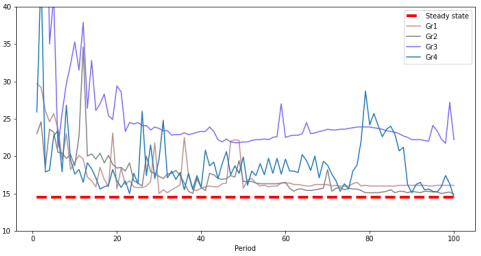


(e) $\lambda = 3.9$

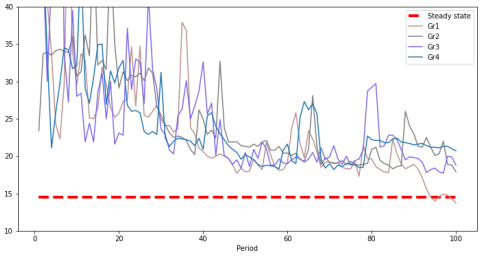
Figure A1: The price dynamics in LtFE sessions per treatment



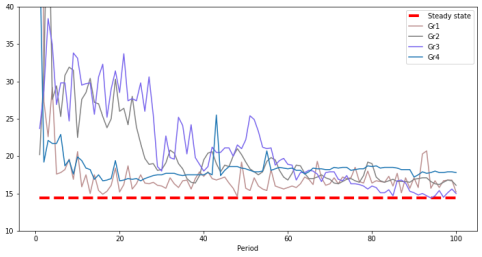
(a) $\lambda = 3.5$



(b) $\lambda = 3.8$



(c) $\lambda = 3.83$



(d) $\lambda = 3.9$

Figure A2: The price dynamics in LtO sessions per treatment

G.2 Regression tables

	EER		TTC	RSD	RMSE		ARDE	Uncertainty		Time on round
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$\lambda = 3.83$	-0.0903*	-0.2191 ***	16.31*	0.0659	16.93	37.31***	0.0740	0.318**	0.720***	11.16***
	(0.0316)	(0.0352)	(6.623)	(0.0427)	(18.25)	(8.697)	(0.0663)	(0.114)	(0.182)	(3.452)
constant	0.906***	0.968***	10.44**	0.012***	10.46***	29.39***	0.0141	0.862***	0.749***	14.39***
	(0.0169)	(0.0022)	(2.873)	(0.0040)	(2.195)	(5.279)	(0.00748)	(0.0657)	(0.121)	(1.905)
Group FE	-	+	-	-	-	+	-	-	+	+
N	20	139	20	20	20	140	20	20	140	70
R^2	0.255	0.742	0.261	0.319	0.136	0.657	0.189	0.225	0.324	0.264

Notes: Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. EER - average payoff relative to the maximum possible payoff; TTC - time to converge to equilibrium and stay within 5% from it for at least 10 rounds; RSD - relative standard deviation of the forecasts; ARDE - average relative distance to the equilibrium; Uncertainty - uncertainty index based on rounding of forecasts; Time on round - average time spent on experimental round. The data on the time spent on round spans 5 experimental sessions with 2 groups in each.

Table A1: Testing non-monotonicity in the LtFE

	D_o	RD	EER_s		EER_f		FE_r		RSD_s	RSD_f	$Uncertainty$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$\lambda = 3.83$	0.743*	3.670***	-0.0758***	-0.0685	-0.0258	-0.0799	0.0617	-0.7309***	0.0892*	0.0635	-0.0427	-0.0357
	(4.200)	(1.158)	(0.0243)	(0.0846)	(0.0176)	(0.0577)	(0.0527)	(0.2753)	(0.127)	(0.0345)	(0.0972)	(0.1864)
constant	6.903***	6.127***	0.853***	0.824***	0.911***	0.944***	0.161***	0.445***	0.826***	0.326***	0.750***	0.694***
	(1.970)	(1.101)	(0.0113)	(0.0715)	(0.0108)	(0.0176)	(0.0276)	(0.0685)	(0.2739)	(0.0232)	(0.0860)	(0.1247)
Group FE	-	-	-	+	-	+	-	+	-	-	-	+
N	16	16	16	108	16	108	16	16	108	16	16	108
R^2	0.195	0.207	0.434	0.104	0.103	0.135	0.084	0.406	0.031	0.136	0.006	0.232

Notes: Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. D_o - difference between the optimal savings and the actual savings decisions. EER_s - average payoff for the savings task relative to the maximum possible payoff; EER_f - average payoff for the forecasting task relative to the maximum possible payoff; RSD_s - relative standard deviation of the forecasts; RSD_s - relative standard deviation of the savings decisions; RD - average relative distance to the equilibrium; Uncertainty - uncertainty index based on rounding of forecasts;FE_r - forecast error divided by the mean forecast.

Table A2: Testing non-monotonicity in the LtOE