

Learning to Believe in Simple Equilibria in a Complex OLG Economy - evidence from the lab*

Jasmina Arifovic[†] Cars Hommes^{‡§} Isabelle Salle^{‡§}

June 17, 2016

Abstract

We set up a laboratory experiment within the overlapping-generations model of Grandmont (1985). Under perfect foresight this model displays infinitely many equilibria: a steady state, periodic as well as chaotic equilibria. Moreover, there exists some learning theory predicting convergence to each of these equilibria. We use experimental evidence as an equilibrium selection device in this complex OLG economy, and investigate on which outcomes subjects most likely coordinate. We use two alternative experimental designs: learning-to-forecast, in which subjects predict the future price of the good, and learning-to-optimize, in which subjects make savings decision. We find that coordination on a steady state or 2-cycle are the only outcomes in this complex environment. In the learning-to-forecast design, coordination on a 2-cycle occurs frequently, even in the chaotic parameter range. Simulations of a behavioral heuristic switching model result in initial coordination on a simple AR(1) rule though sample autocorrelation learning, with subsequent coordination on a simple second-order adaptive rule once the up-and-down pattern of prices has been learned.

JEL codes: C90, D83, C62.

Keywords: laboratory experiments, learning, complex dynamics, equilibrium selection.

***Acknowledgments:** We would like to thank Shyam Sunder and Mike Woodford for stimulating discussions and feedback. We also thank the participants at the CREED seminar at the University of Amsterdam, October 28, 2014, the CEF conference in Taipei, June 20-22, 2015 and the workshop on Behavioral Macroeconomics, Tinbergen Institute Amsterdam, October 26, 2015 for helpful comments. This research has been financed by the EU FP7 projects *RAstaNEWS*, grant agreement No. 320278 and *MACFINROBODS*, grant agreement No. 612796. None of the above are responsible for errors in this paper.

[†]Simon Fraser University

[‡]CeNDEF, Amsterdam School of Economics, University of Amsterdam

[§]Tinbergen Institute

1 Introduction

A well-known concern among theorists is that (macro)economic rational expectation models are prone to indeterminacy. These models may possess multiple long-run equilibria, some of which may involve sub-optimal outcomes or volatility in real variables that are undesirable from the point of view of policy makers aiming at stabilizing aggregate fluctuations. Selecting among those equilibria is then a critical issue in the model and for policy design.

From a theoretical point of view, adaptive learning has been frequently advocated as an equilibrium selection device. The main idea of this literature is that agents are not endowed with rational expectations beforehand, but rational expectations may be viewed as the long run outcome of some adaptive learning process. Only equilibria that emerge as a long-run outcome of the adaptive learning process are then regarded as plausible (see Evans & Honkapohja (2001) for a comprehensive discussion).

A problem with the theory of adaptive learning is that "anything goes", that is any equilibrium may be stable under some suitable form of adaptive learning. For example, in an OLG economy with infinitely many periodic equilibria, any equilibrium cycle can be learned provided that the adaptive rule of agents is consistent with the periodicity of the cycle (Grandmont 1985, Guesnerie & Woodford 1991, Evans & Honkapohja 1995). In a similar set-up, Woodford (1990)'s *learning-to-believe in sunspots* shows that a suitable adaptive learning rule may lead to convergence to a sunspot equilibrium with probability one. The goal of our paper is to design laboratory experiments as an empirical test for equilibrium selection and stability under learning in a complex environment.

Laboratory experiments with human subjects constitute a promising way of empirically testing the learning predictions, and determining which rules agents are likely to adopt. Experimental evidence can then also serve as an equilibrium selection device in case of indeterminacy in economic models, as in the pioneering experimental work of Marimon et al. (1993), Marimon & Sunder (1993, 1994). Lucas (1986) already stressed the importance of an experimental approach in studying expectations and stability of equilibria under learning:

Recent theoretical work is making it increasingly clear that the multiplicity of equilibria [...] can arise in a wide variety of situations [...]. All but a few equilibria are, I believe, behaviorally uninteresting: They do not describe behavior that collections of adaptively behaving people would ever hit on. I think an appropriate stability theory can be useful in weeding out these uninteresting equilibria [...]. But to be

useful, stability theory must be more than simply a fancy way of saying that one does not want to think about certain equilibria. I prefer to view it as an experimentally testable hypothesis [...]. (Lucas 1986, pp. S424-S425)

This paper presents such an experimental study within the environment of a complex OLG economy *à la* Grandmont (1985). This environment is particularly interesting in regards to the question of equilibrium selection because it possesses infinitely many long-run equilibria, including a steady state, cycles of all periods, and even chaotic dynamics. Cycles are not caused by exogenous or policy shocks as the OLG model is deterministic. Complicated dynamics arise as an equilibrium outcome of the model as soon as there is a strong conflict between substitution and wealth effects of a change in the return on savings, which can be easily tuned by varying the risk aversion parameter of the utility function in the model. This model has been extensively studied in the learning literature. All these equilibria are stable under adaptive learning if agents use a suitable rule that is consistent with their periodicity (Grandmont 1985, Guesnerie & Woodford 1991). Interestingly, there are two adaptive learning theories that predict only simple outcomes in this complex OLG environment. Evolutionary genetic algorithms (GA) learning only selects forecasting rules that are consistent with the steady state or the 2-cycle (Bullard & Duffy 1998). Hommes et al. (2013) reach similar conclusions when agents forecast prices by sample autocorrelation learning, where agents learn a simple but optimal AR(1) rule with correct sample average and sample autocorrelation (cf. Hommes & Zhu (2014)).

We design several experimental treatments and formulate two hypotheses. The first hypothesis states that coordination on simple equilibria (for instance a steady state or a two-cycle) is more likely to emerge than coordination on more complicated equilibria. We therefore consider different treatments with increasingly complicated dynamics. The second hypothesis postulates that the ability to coordinate also depends on the experimental task. Previous experiments show that when agents have to directly submit quantity decisions, the experimental economies display more variability and coordination of subjects' strategies is more challenging than when subjects have to make forecasts, see e.g. Bao et al. (2013, 2016). We therefore consider two designs, a learning-to-forecast (LtF) design, where subjects only forecast the price and the optimal savings decision is computed based on their forecasts, and a learning-to-optimize (LtO) design, where subjects make directly savings decisions. The second hypothesis is that coordination may arise on simpler equilibria when subjects have to optimize than when they have to forecast.

Our results may be summarized as follows. First, we always observe coordination on either

the monetary steady state or the 2-cycle, but never on any higher-order cycle. Our experiment is the first in which coordination of expectations of a group of subjects on a 2-cycle equilibrium arises (see the discussion of related work by Marimon et al. (1993) below). Second, the more "unstable" the steady state, the more likely the coordination on the 2-cycle in the LtFE, even if the 2-cycle is unstable in the backward perfect foresight dynamics. This coordination arises spontaneously in the experiment. We obtain both aggregate convergence of price values and individual coordination of price forecasts on the steady state or the 2-cycle, possibly after a long transition. These outcomes are predicted by the weak E-stability criterion: only if the forecasting rule of the agents is *exactly* consistent with a steady state or a 2-cycle, these outcomes may be achieved under adaptive learning. However, this result is not robust to misspecification or overparametrisation of the forecasting rule. Hence, our result suggest that subjects make use of first or second order behavioral rules, but do not use higher order rules. By contrast, in the LtOE, subjects spontaneously coordinate on the monetary steady state, even if it is unstable under adaptive learning. They do so with more heterogeneity in the savings decisions, and aggregate price dynamics displays more volatility than under the LtFE. However, if subjects are "trained" at the beginning of the experiment with a fictitious group of players who are coordinated on the 2-cycle, we may observe convergence to a "noisy" or "attenuated" 2-cycle. This suggests that subjects tend to select simpler equilibria (a steady state rather than a 2-cycle) as a coordination device when the sophistication of their task increases.

Related literature A large number of experimental studies have explored the question of equilibrium selection in static or repeated games, see e.g. Camerer (2003) for a survey. We discuss here two contributions that are closely related to our experimental study, but with important differences. Van Huyck et al. (1994) investigate the question of equilibrium selection in an experiment within a coordination game with two efficient Nash equilibria. The myopic best response dynamics coincide with the chaotic quadratic map, while the interior equilibrium is stable under adaptive learning (if the gain is small enough). In all their experimental sessions, subjects coordinate on the interior solution, in line with the prediction of adaptive learning. An important difference with our experimental environment is that our set-up has infinitely many perfect foresight periodic cycles that arise as equilibrium outcomes of the model, without imposing a priori an expectation rule, and these cycles can be stable under adaptive learning. By contrast, in Van Huyck et al. (1994), the chaotic dynamics is not an equilibrium of the coordination game, but results from the assumption of myopic best response behavior. The

authors do not address the question whether or how subjects may coordinate on the best response, which seems especially difficult since it involves complicated dynamics. In our complex OLG environment, we are interested which of those many periodic equilibria, if any, subjects may coordinate on. The second closely related contribution is the work by Marimon et al. (1993), which have been the first (and, to the best of our knowledge, so far the only one) to observe, to some extent, coordination on two-cycle type of dynamics in a laboratory experiment. They use a design similar to our LtFE, but there are several major differences with respect to our framework and our results. First, they consider an OLG environment in which only a steady state, a two period cycle and two state sunspot equilibria exist, while our model involves infinitely many periodic and chaotic equilibria, making our equilibrium selection problem more complex. Second, they employ a three group design, in which each generation is renewed from a pool of subjects. We use a single group design, so that the resulting course of events in the experiment is the same as in the adaptive learning literature, especially the seminal contribution of Grandmont (1985). Most importantly, Marimon et al. impose real shocks to the OLG economy by cyclically varying the number of subjects in each generation between a high and a low number in phase with the color of a blinking square on subjects' computer screens. This generates temporary "attenuated" 2-cycle oscillations driven by these exogenous shocks. However, these oscillations dampen out once the exogenous shocks to the generation size were turned off.¹ Hence, Marimon et al. find no evidence of 2-cycles arising spontaneously. In our LtFE, subjects *spontaneously* select the 2-cycle as a coordination device in an OLG economy where many more complex equilibria exist.

The rest of the paper is organized as follows. Section 2 introduces the OLG model of the experiment, and discusses its properties and learning dynamics. Section 3 details and motivates the experimental design and the hypotheses, while Section 4 presents the experimental results. Section 5 discusses a behavioral model explaining our experimental data, and Section 6 concludes.

2 The model

2.1 The underlying OLG economy

The underlying model of the experiment is a deterministic OLG economy à la Grandmont (1985). This is an exchange economy with a single perishable consumption good and constant population. In each period t , a continuum (of measure 1) of identical agents is born and lives for two periods,

¹See the five economies in Marimon et al. (1993, Figure 3, p. 89).

so that, in each period, two generations coexist: the young and the old generations. Individuals receive an endowment $e_1 > 0$ of the consumption good when young, and $0 < e_2 < 1$ when old. Young individuals can save part of their first-period endowment by selling to the old individuals a quantity $s_t \in (0, e_1]$ of the good at the market-clearing price P_t , and holding the corresponding (non-negative) money balances, denoted by $m_t = p_t s_t$. The aggregate quantity of money supply in the economy is held constant, and endogenously set to $M > 0$. Once old, in the next period ($t + 1$), any individual purchases back goods from the young generation using all his savings at the market-clearing price, denoted P_{t+1} .

Formally, at any arbitrary period, a young individual chooses his current consumption, denoted by c_t (conversely his real money balances s_t), to maximize his two-period expected utility function, denoted by $U(c_t, c_{t+1}^e)$, subject to his current (when young) and expected (when old) budget constraints:

$$\begin{cases} c_t & \leq e_1 - s_t \\ c_{t+1}^e & \leq e_2 + \frac{P_t}{P_{t+1}^e} s_t, \end{cases} \quad (1)$$

where $R_{t+1}^e \equiv \frac{P_t}{P_{t+1}^e}$ corresponds to the expected gross return on savings.

2.2 Definition of a perfect foresight equilibrium

We follow Grandmont (1985) and assume a separable utility function, i.e. $U(c_t, c_{t+1}) = V_1(c_t) + V_2(c_{t+1})$, with the functions $V_{1,2}(\cdot)$ being continuous, strictly increasing and concave on $[0, +\infty)$, twice continuously differentiable on $(0, +\infty)$, with $\lim_{c \rightarrow 0} V'(c) = +\infty$. These properties together with the compactness of the budget constraints ensure that the maximization problem of the young individuals has a unique solution. The first order condition can be expressed as:

$$V_1'(c_t)P_{t+1} = V_2'(c_{t+1})P_t \quad (2)$$

or equivalently, in terms of real money balances:

$$V_1'(e_1 - s_t)s_t = s_{t+1}V_2'(s_{t+1} + e_2). \quad (3)$$

Once the optimal savings and consumption decisions are determined, it is possible to define a perfect foresight equilibrium sequence of prices (or equivalently of real money balances) using

market clearing condition in the money and the good markets:

$$m_t = M, s_t = \frac{M}{P_t} \text{ and } s_{t+1} = \frac{P_t}{P_{t+1}} s_t \quad \forall t \quad (4)$$

At this stage, it is convenient to define functions $v_1(s) = sV_1'(e_1 - s)$, that maps $[0, e_1)$ onto $[0, +\infty)$, and $v_2(s) = sV_2'(e_2 + s)$ that maps $[0, +\infty)$ into itself. $v_1(\cdot)$ is strictly increasing, so that $v_1^{-1}(\cdot)$ exists. The dynamics of s under perfect foresight can then be described by the continuous map $\chi = v_1^{-1} \circ v_2$. The graph of χ is called the *offer curve*, defined as the locus of points representing optimal consumption when young and when old (c_t, c_{t+1}) when the return on savings R_{t+1} varies. Equivalently, the dynamics of the model can be expressed in terms of prices by a continuous map, denoted by G :

$$P_t = G(P_{t+1}^e) = \frac{M}{\chi(M/P_{t+1}^e)}. \quad (5)$$

The maps G and χ are topologically equivalent, i.e. describing the dynamics of the model in terms of prices or savings is equivalent.

A *perfect foresight (periodic) equilibrium* is a (periodic) sequence of prices that is a solution of (5) with $P_{t+1}^e = P_{t+1}$. A perfect foresight steady state is a fixed point \bar{P} of the map G , so that $\bar{P} = G(\bar{P})$. A periodic perfect foresight equilibrium of period k (k being the smallest integer greater than one satisfying this) is a sequence of k prices $\{P_1, P_2, \dots, P_k\}$ (or orbit), such that $P_j = G^k(P_j)$, $j = 1, \dots, k$, where G^k denotes the k^{th} iterate of the map G .

2.3 Existence of perfect foresight equilibria

There are at most two steady states in the model (Gale 1973): one monetary steady state where real money balances are strictly positive, and the sequence of the returns on savings equals unity, and one non-monetary steady state where aggregate savings is zero and individuals consume their endowment every period. In his seminal paper, Grandmont (1985) shows that this economy may possess infinitely many perfect foresight equilibria when the income effect of a change in the return on savings R is sufficiently strong, as an increase in R has an ambivalent effect on consumption when young. These equilibria include periodic equilibria of any period and infinitely many chaotic equilibria.²

²Azariadis & Guesnerie (1986) show that if the model has an equilibrium cycle, it also has many sunspot equilibria, on which expectations may coordinate.

In our lab experiments, we use CRRA utility functions, as often in the related literature:

$$V(c_1) = \frac{c_1^{1-\rho_1}}{1-\rho_1}, V(c_2) = \frac{c_2^{1-\rho_2}}{1-\rho_2}, \quad (6)$$

where we further assume $0 < \rho_1 < 1$ and $\rho_2 > 0$. Parameters ρ_1 and ρ_2 measure the degree of relative risk aversion of the young and the old individuals, and play a critical role in the long-run dynamics of the economy. Given the characteristics of the model and the further assumption that $e_1 + e_2 > 1/e_2$, Grandmont (1985, Corollary to Proposition 4.4, p. 1023) shows that complex dynamics arise as a long-run outcome of the model as soon as ρ_2 is high enough. When $\rho_2 \leq 1$, substitution effects dominate, the offer curve is monotonic for all consumption values, and the dynamics always converge to the unique monetary steady state. When $\rho_2 > 1$, the offer curve becomes non-monotonic. As ρ_2 increases, the offer curve becomes more and more humpy, the long-run price dynamics becomes increasingly complicated, after an infinite cascade of period-doubling bifurcations. For values of ρ_2 sufficiently high, the map has infinitely many periodic as well as chaotic perfect foresight equilibria, together with the monetary steady state.

2.4 Stability of perfect foresight equilibria under learning

Grandmont (1985) distinguishes between the *forward perfect foresight dynamics* (when the map G in (5) is defined with $P_{t+1}^e = P_{t+1}$) and the *backward perfect foresight dynamics* (when $P_{t+1}^e = P_{t-1}$ in (5)).³ Equilibria that are (locally) unstable in the forward perfect foresight dynamics are (locally) stable in the backward perfect foresight dynamics. However, those two dynamics are largely theoretical outcomes, as forward perfect foresight dynamics can be regarded as the long-run outcome of some learning process, and backward dynamics is fictitious (time goes backward).⁴ Grandmont (1985) advocates an expectation formation process that is based on past prices (akin to econometric learning), together with mild assumptions on the expectation function, and proves that an equilibrium which is stable in the backward perfect foresight dynamics is stable in the *forward dynamics with learning*. Provided that the memory of past prices in the expectation function is consistent with the periodicity of such a cycle, any cycle can be learned under adaptive learning.

For later use, we summarize the stability conditions obtained in the literature under different

³The forward dynamics may not be globally well-defined. Gardini et al. (2009) characterize the forward perfect foresight equilibria as iterated function systems with fractal attractors.

⁴Another interpretation of the backward perfect foresight dynamics is that agents have naive expectations.

learning schemes consistent with periodic equilibria, and how they relate to each other. Any period k -cycle $\{P_1^*, \dots, P_k^*\}$ ($k \geq 1$) of the map G is stable under *backward perfect foresight* if and only if $|DG^k(P_k^*)| = |\prod_{i=1}^k DG(P_i^*)| < 1$, where DG^k is the derivative of the k -th iterate of G . This condition corresponds to the determinacy condition of any cycle under perfect foresight (see e.g. Guesnerie & Woodford (1991)). This condition is also equivalent to the *strong E-stability condition* under recursive learning of Evans & Honkapohja (1995, Proposition 3, p. 197), defined as a stability criterion which is robust to over-parametrization of the forecasting rules of agents. If an equilibrium is stable only when the forecasting rule of the agents is exactly consistent with its periodicity, it is said to be *weakly E-stable*. The weak E-stability criterion is less stringent, and requires $DG^k(P_k^*) < 1$ for $k \leq 2$, and $-\cos(\pi/k)^{-k} < DG^k(P_k^*) < 1$ for $k > 2$. When $k \rightarrow +\infty$, this condition is equivalent to strong E-stability.

Guesnerie & Woodford (1991) consider a period- k adaptive expectation scheme with a constant parameter $0 < w < 1$:

$$p_{t+1}^e = wp_{t+1-k} + (1-w)p_{t+1-k}^e \quad (7)$$

which is consistent with a period $k \geq 1$ equilibrium, and reduces to the backward perfect foresight dynamics when $w = 1$ and $k = 2$. The strong E-stability is a sufficient condition for stability under adaptive expectations. The necessary and sufficient condition for stability is a complicated function of k and w which cannot be solved in closed form for $k > 2$, but reduces to weak E-stability when $w \rightarrow 0$, and to strong E-stability when $w \rightarrow 1$. Moreover, when $k = 1$, a steady state P^* is stable under adaptive expectations if and only if $DG(P^*) < 1$ or $DG(P^*) > \frac{2-w}{w}$. When $k = 2$, the stability condition of a 2-cycle becomes $-\frac{(2-w)^2}{w} < DG^2(P_{1,2}^*) < 1$. If we note $d \equiv DG^2(P_{1,2}^*)$, a 2-cycle is stable under the rule (7) if and only if i) $w \in (0, \underline{w})$, where $\underline{w} = \frac{4-d-\sqrt{d(d-8)}}{2}$, if $d \leq -1$, ii) is always stable if $d \in (-1, 1)$.

Figure 1 reproduces the bifurcation diagram of the backward perfect foresight dynamics in Grandmont (1985, p. 1030) using the parametrization: $e_1 = 2$, $e_2 = 0.5$, $\rho_1 = 0.5$, that we also use in the experiments. The long run outcomes of the model are displayed in terms of real money balances (y-axis) for any value of $\rho_2 > 2$ (x-axis). Under this calibration, Grandmont (1985, Lemma 4.6, p.1026) shows that there is at most one periodic equilibrium, say of period k , that is stable in the backward perfect foresight dynamics for each ρ_2 value. As long as the forecasting rules of the agents are consistent with the k -periodicity of this equilibrium, it is equivalently stable in the forward dynamics under learning, and it follows that it is strongly E-

stable under recursive learning. All other equilibria that may co-exist are either weakly E-stable or E-unstable, but (locally) stable in the forward perfect foresight dynamics. In particular, when there is a cycle of period 3, it is well-known that cycles of any periodicity co-exist with the period 3 cycle as equilibrium solutions of the system. This happens e.g. for values of ρ_2 beyond 13.




Figure 1: Bifurcation diagram under backward perfect foresights dynamics in Grandmont (1985, p. 1030). The red vertical lines indicate the five different ρ_2 values run in the experiments (see Section 3).

At least two other learning mechanisms have been applied to this specific OLG economy, and predict different outcomes from the ones under adaptive learning. Bullard & Duffy (1998) use an heterogeneous agent version of this model, with two populations, in which agents forecast using evolutionary learning. A genetic algorithm (GA) encodes the lag in the past price series to be chosen to form the next period's price forecast, so that their algorithm can in principle learn higher order cycles. They conduct numerical simulations of the OLG economy under the same range of ρ_2 values as in Figure 1. The model displays infinitely many equilibria, including many cycles of different periods. However, they only observe two outcomes under GA learning: convergence to the monetary steady state or to the two-cycle when it exists, i.e. for values of ρ_2 roughly higher than 4.⁵ This result suggests that learning agents tend to coordinate on simple equilibria.

Hommes et al. (2013) apply the so-called Sample-AutoCorrelation (SAC) learning (Hommes & Zhu 2014) to this OLG economy with $\rho_2 = 12$, i.e. when the dynamics in the backward

⁵They also find one 4-period cycle for one simulation at a specific value of ρ_2 , as well as two cases of non-convergence.

perfect foresight is chaotic. SAC learning assumes that boundedly rational agents make use of a parsimonious linear AR(1) forecasting rule, and update the two parameter values using the observed sample average and first order autocorrelation of past prices. SAC-learning also only selects simple equilibria – steady state or a noisy two-cycle – in this complex environment.⁶

We design an experiment to test which theories of learning outcomes are empirically relevant in this complex OLG environment.

3 Experimental design

The experiment is a *single-group design* and a within-session randomization. At the beginning of every experimental session, participants are divided into groups of $N = 6$ subjects, and each group represents an experimental economy. Each experimental economy is governed by the OLG model described in Section 2. In the single group design, each participant repeatedly plays the role of a professional advisor bureau to one young individual for T periods (see Heemeijer et al. (2012) for a similar OLG design). As the role of the old individuals in the OLG framework is essentially passive (they just consume the amount of goods that their savings can buy), they do not make any strategic decision, and subjects do not need to advise them. This single group design is in line with the learning literature in this framework, e.g. as considered in Grandmont (1985) and the references to the adaptive learning literature previously cited, and thus forms a natural empirical test for learning in an OLG framework.

We consider two different experimental designs, learning-to-forecast (LtF) and learning-to-optimize (LtO), where subjects have to perform different tasks, either to submit price forecasts or savings decisions. These different designs have been introduced by Marimon et al. (1993).

3.1 The Learning-to-Forecast Experiment (LtFE)

In the LtF design, subjects provide a forecast of the future price of the consumption good to the members of the young generation, who then use the forecast to optimally decide upon their savings. Subjects have to form two-period ahead forecasts: at the beginning of every period/generation t , each subject $i = 1, \dots, N$ has to submit a forecast of the price $P_{i,t}^e(t+1)$ in the next period $t+1$. We assume that every member of the young generation then makes

⁶Along a 2-cycle (p_1, p_2) , SAC-learning parameters (α_t, β_t) converge to $\alpha^* = \frac{p_1+p_2}{2}$ and $\beta^* = -1$. Hommes et al. (2013) consider a stochastic model and refer to a "noisy two-cycle" when the first order autocorrelation of prices converges to a value close to -1 .

the optimal savings decision, conditional on the forecast he receives from his advisor. Formally, using the CRRA utility functions (6), and combining the first order condition (2) with the budget constraint (1), the optimal consumption of any young individual i , denoted by $c_{i,t}$ (or conversely his real money balance $e_1 - c_{i,t}$) is implicitly defined by:

$$c_{i,t} + c_{i,t}^{(\rho_1/\rho_2)} \frac{P_{i,t}^e(t+1)^{[(\rho_2-1)/\rho_2]}}{P_t} = e_1 + e_2 \frac{P_{i,t}^e(t+1)}{P_t} \quad (8)$$

where the market clearing price at time t is given by $P_t = \frac{M}{\sum_{i=1}^N s_{i,t}} = \frac{M}{Ne_1 - \sum_{i=1}^N c_{i,t}}$.

Sequence of events In each generation t , once the N price forecasts have been submitted by the subjects, the corresponding level of consumption and savings of each individual, together with the market clearing price P_t , are numerically solved for (as condition (8) does not allow for a closed-form solution of the optimal individual consumption levels), and displayed to the subjects. The experimental economy then goes to the next generation $t+1$, etc. Note that for the first period, subjects have to submit two price forecasts, for the current period 1 and the next period 2, before the first market clearing price P_1 can be computed.

Payoff Subjects earn points as a function of their forecast errors. The lower their forecast error, the higher their payoff. We use the quadratic payoff function, as in Bao et al. (2016) :

$$\max \left(1300 - \frac{1300}{49} (P_{i,t}^e(t+1) - P_{t+1})^2, 0 \right) \quad (9)$$

in which the payoff is maximal and equal to 1300 points in case of perfect prediction, and equals zero if the prediction error is higher than 7. Note that the timing of the payoff is two-period ahead, as subjects only observe the realized price and their forecast error at the end of the next period. A payoff table that shows the pay-off value for a grid of forecast errors was provided in the instructions.

This LtF design has been first introduced in the OLG experiments of Marimon et al. (1993). It allows us to focus entirely on the expectation formation process of subjects, and to assess which expectation model may describe the forecasts data observed in the lab, and explain the resulting price dynamics.

3.2 The Learning-to-Optimize Experiment (LtOE)

In the second design – a so-called Learning-to-Optimize experiment (LtOE), see e.g. Bao et al. (2013, 2016), we drop the assumption of optimal conditional savings decisions, and ask the subjects to directly submit the savings quantity of a young individual. This savings decision may be based on his forecast of the return on savings P_t/P_{t+1}^e , but we do not explicitly elicit return forecasts from subjects. This is essentially because this design focuses on quantity decisions of subjects, and we did not want to introduce a more demanding cognitive load, by combining a forecasting and an optimizing task (see Bao et al. (2016) for more discussion). However, subjects are instructed (see Appendix E) that they face a two-stage decision process, and they first have to forecast the return on savings, and then choose the corresponding optimal value of savings using their two-dimensional payoff table (see below). Additionally, visual information (e.g. question marks in the table on their screen, see Figure 27 in Appendix F) indicates the two-period ahead nature of the forecast of the return on savings. Quantities that are displayed to the subjects are also scaled by a factor of 100, so that they make decisions in the interval between 0 and 200, and not between 0 and 2. This allows an easier interpretation of the savings task.

Sequence of events At any generation t , once every subject $i = 1, \dots, N$ has submitted a savings decision for a member of the young generation, denoted by $s_{i,t}$, the market clearing price for the consumption good is given by $P_t = \frac{M}{\sum_{i=1}^N s_{i,t}}$, and the experimental economy goes to the next generation $t + 1$, etc.⁷

Payoff Subjects are rewarded by the realized utility of the young individual over his two-period life. As in the LtFE, the timing of the payoff of any savings decision is then *two-period ahead*: a savings decision made for any member of the young generation in period t is rewarded at the end of period $t + 1$. In order to implement this payoff scheme in the lab, we use two transformations of the utility function U (with separable utility functions given by (6)). First, in order to rule

⁷We also initialize the price P_0 so that subjects can observe a value of the return on savings right from period 1, after having submitted their first savings decision, i.e. $\frac{P_0}{P_1}$. This is to avoid that subjects have to submit two savings decisions in a row without seeing the first realization of the return on savings. In this case, pilot sessions indicate that they would have no reason to change their decisions, and the first two realizations of aggregate savings would be similar, and so would the first two realizations of the price. The first realization of the return on savings would then be close to one, artificially driving the experimental economies towards the steady state. We chose the initial values P_0 (specifically 50 for $\rho_2 = 3,5$ and 10 for $\rho_2 = 8$, see below) i) to be consistent with the initial price ranges given in the LtFE, and ii) in order for the first return to be sufficiently different from unity, but not too extreme, so that the plots on the subjects' screen remain readable.

out negative payoff, we apply the following transformation of the utility:

$$\tilde{u} = \max(K \times (U(c_{i,t}, c_{i,t+1}) + C), 0) \quad (10)$$

where the parameters $K, C > 0$ are chosen to keep the values of the payoff function under LtOE in the same order of magnitude as the ones under LtFE, and to ensure that any equilibrium real money balances gives rise to a non-zero payoff.⁸ We use the payoff function (10) for $\rho_2 = 3$, when the monetary steady state is the only equilibrium solution of the model (see Subsection 3.3).

Additionally, all periodic equilibria in the OLG are Pareto-optimal, but differ in terms of intergenerational equity (Grandmont 1985). This means that utility values along cycles, and especially along the 2-cycle and at the monetary steady state may differ. In order to be consistent with the LtF design where subjects' payoff is maximized and the same along perfect foresight cycles and at the steady state, and to give an equal chance to coordination on the 2-cycle or the steady state, we apply the following transformation of the payoff function:

$$\hat{u} = 1300 \times \left(\frac{\tilde{u}}{1300} \right)^\alpha \quad (11)$$

where the scale parameter 1300 is chosen in consistency with the payoff function under LtFE, and α is adjusted so that the average payoff along the two cycle is in the same order of magnitude as the payoff at the steady state. In the case of $\rho = 5$ and 8 (see Subsection 3.3), the payoff of any given savings decision s_t is given by (11).

The instructions given to the subjects include a two-dimensional payoff table that report the expected payoff of a discrete grid of savings decisions as a function of a range of expected values of the return on savings (see Appendix E). The optimal savings decisions conditional on each expected return on savings then correspond to the consumers' offer curve, and the shape of the offer curve is unaffected by the transformations of the utility functions that we have considered.

3.3 Experimental treatments and hypotheses

We adopt the calibration used in Section 2, i.e. $e_1 = 2$, $e_2 = 0.5$, $\rho_1 = 0.5$. We then vary the parameter ρ_2 to define different treatments with increasingly more complicated equilibrium outcomes⁹ (see Section 2 and Figure 1).

⁸See Marimon et al. (1993) for a similar transformation of the utility function.

⁹The implementation of the OLG model in the lab rules out the possibility of chaotic dynamics, as price values are rounded to two digits on subjects' screen, and it becomes impossible to construct a

In the LtFE, we consider five treatments: $\rho_2 = 3$, $\rho_2 = 5$, $\rho_2 = 8$, $\rho_2 = 12$ and $\rho_2 = 13.5$ (see Figure 1). Table 1 summarizes the stable outcomes predicted by the learning literature in each treatment. For $\rho_2 = 3$, the monetary steady state is strongly E-stable, and the map G does not have any other equilibrium.¹⁰ For $\rho_2 = 5$, the map has both a monetary steady state and a period 2 cycle, that is the only stable outcome under backward perfect foresight and the only strongly E-stable outcome, while the steady state is weakly stable. For $\rho_2 = 8$, the map G has no strongly E-stable cycle, the monetary steady state, the 2-cycle and the 4-cycle are the only weakly E-stable equilibria, while all cycles¹¹ are (locally) stable in the forward perfect foresight dynamics. For $\rho_2 = 12$ and $\rho_2 = 13.5$, the map G has chaotic dynamics, and infinitely many periodic cycles (for all periodicities but three for $\rho_2 = 12$, and including three for $\rho_2 = 13.5$), and chaotic orbits (together with the monetary steady state). None of the cyclic equilibria are strongly E-stable or stable in the backward perfect foresight dynamics with $\rho_2 = 12$, while they are all locally stable under the forward perfect foresight dynamics. With $\rho_2 = 13.5$, the only strongly E-stable outcome is a period 3 cycle. The stability of the 2-cycle under the adaptive rule (7) in Treatments $\rho_2 = 5$, $\rho_2 = 8$, $\rho_2 = 12$ and $\rho_2 = 13.5$ depends on the weight w on past prices. In this economy, cycles of period 2^k are created by period doubling bifurcations, so that the derivative of the second iterate of the map G , $DG^2(P_{1,2}^*)$ at the 2-cycle is negative. In case of $\rho_2 = 8, 12$ or 13.5 , we have $DG^2(P_{1,2}^*) < -1$, while $DG^2(P_{1,2}^*) \in (-1, 0)$ when $\rho_2 = 5$ as the 2-cycle is strongly E-stable. For our calibration, the stability threshold \underline{w} for adaptive expectations in (7) are: for $\rho_2 = 8$, $\underline{w} \simeq 0.8$, for $\rho_2 = 12$, $\underline{w} \simeq 0.61$ and for $\rho_2 = 13.5$, $\underline{w} \simeq 0.57$.

We use a non-linear, *topologically equivalent* transformation of the map G in treatments with $\rho_2 = 12$ and $\rho_2 = 13.5$. This is because, for sufficiently large ρ_2 values, real money balances become close to the bounds 0 and 2, which produce price values that are too wide to be readable on the graph and the table of the subjects' screen (see Figure 27 in Appendix F). Therefore, for $\rho_2 = 12$ and $\rho_2 = 13.5$, we map the subjects' price forecasts into the actual price values as follows:

$$p = H(\tilde{p}) = 3.5^{\frac{\tilde{p}}{8}} - 1 \text{ or equivalently the inverse } \tilde{p} = H^{-1}(p) = 8 \times \frac{\ln(p+1)}{\ln(3.5)} \quad (12)$$

bounded path that never repeats any past value. However, it still leaves room for high order cycles.

¹⁰In the lab, price forecasts are bounded, and the autarkic steady state in which agents only consume their endowment and do not save at all is not feasible. Similarly, in the LtOE, subjects are instructed to submit a strictly positive savings decision, so that the price level is always defined, and the monetary steady state is the only feasible steady state.

¹¹These cycles have periodicity that is a multiple of 2^k after the cascade of period doubling bifurcation.

ρ_2	backward p.f.	forward p.f.	strong E-stability	weak E-stability	adaptive expectations	SAC learning	GA
	Grandmont (1985)		Evans & Honkapohja (1995)		Guesnerie & Woodford (1991)	Hommes & et al. (2013)	Bullard & Duffy (1998)
3	SS	none	SS	SS	SS	SS iff $\beta = 0$	SS
5	2-cycle	SS	2-cycle	SS 2-cycle	SS 2-cycle $\forall w$	SS iff $\beta = 0$ 2-cycle	SS 2-cycle
8	none	2^k -cycles	none	SS 2-cycle 4-cycle	SS 2-cycle (if $w < 0.8$) 2^k -cycles (if w low enough)	SS iff $\beta = 0$ noisy 2-cycle ($\beta \simeq -1$)	SS 2-cycle
12	none	all cycles (period $\neq 3$)	none	SS 2-cycle	SS 2-cycle (if $w < 0.61$) any cycle (period $\neq 3$, if w low enough)	SS iff $\beta = 0$ noisy 2-cycle ($\beta \simeq -1$)	SS 2-cycle
13.5	3-cycle	all cycles (period $\neq 3$)	3-cycle	SS 2-cycle 3-cycle	SS 2-cycle (if $w < 0.57$) any cycle (if w low enough)	SS iff $\beta = 0$ noisy 2-cycle ($\beta \simeq -1$)	SS 2-cycle

Table 1: Summary of the theoretical learning predictions of stable outcomes in the five treatments. SS stands for the monetary steady state. Stability of any cycle under adaptive expectations is conditional on agents using an adaptive rule consistent with its periodicity. Results under GA learning are obtained through numerical simulations. In bold, we highlight the results in line with our experimental evidence, see Section 4.

with $p, \tilde{p} \in]0, +\infty[$ the actual value of the price.¹²

As the map H is one-to-one, G and H are topologically equivalent, and the non-linear transformation does not affect the dynamical properties of the system (in particular, the number of equilibria and their stability, the slope of any iterate of G and H being the same at any of their fixed points).

By running these five treatments, we test the following hypotheses:

Hypothesis 1. (LtFE): Coordination of forecasts from heterogeneous players, if any, on high-order cycles or complicated dynamics is less likely than on simple equilibria, such as a steady state or a cycle of low periodicity.

This may be because of cognitive and memory limitations, that makes the use of high-order adaptive rules unlikely, or because of the systematic forecasting errors that would result from the use of too simple heuristics, such as naive expectations. In this case, simple outcomes are more likely to serve as coordination devices of heterogeneous beliefs (Bullard & Duffy (1998), Hommes (2011)). We design these five treatments in the LtFE in order to test Hypothesis 1.

We then use the LtOE to assess to what extent the resulting outcomes depend on the nature

¹²For example, for $\rho_2 = 12$, the transformation maps the two-cycle $\{4.67, 13451.29\}$ in $\{11.08, 60.71\}$.

and the difficulty of the experimental task performed by the subjects. The design of the payoff in the LtOE requires a two-dimensional table where subjects need to forecast the return on savings.¹³

Hypothesis 2. (LtO vs. LtFE): Coordination of savings decisions from heterogeneous agents on high-order cycles is less likely under LtOE than under LtFE. In other words, aggregate outcomes are likely to be simpler under LtO than under LtFE.

Hypothesis 2 may be justified by the increasing difficulty of the experimental task in the LtOE compared to the LtFE, where subjects have to cope with a two-stage decision process, and an implicit forecasting task. This may favor simpler coordination device under LtOE than under LtFE. In this context, the monetary steady state may serve as a focal point. The nature of the tasks are also different. Previous experimental evidence suggest that optimizing is a more complicated task than forecasting, and generates noisier aggregate outcomes (see the trial sessions discussed in Marimon et al. (1993) in an OLG model, see Bao et al. (2013, 2016) in, respectively, a cobweb and an asset-pricing model). Additionally, in cognitive psychology, the sequence learning literature concludes that humans are good at learning patterns of up to few prior observations (see Spiliopoulos (2012) and the references herein). Forecasting is more akin to sequence prediction, where a period k -cycle is a pattern of length k , than making savings decisions. Subjects may be more likely to coordinate their forecasts on cycles but of low-order (Hypothesis 1) than their savings decisions.

After observing from pilot sessions and the experimental sessions presented in this paper that only coordination on the steady state arises in the LtOE, we decided to increase the likelihood of coordination on the 2-cycle to test for the robustness of the selection of the monetary steady state. We did so by training the subjects in the first periods of the experiment. Such a training phase is current practice in laboratory experiments, see e.g. Marimon et al. (1993), Duffy & Fisher (2005), Arifovic et al. (2014). Here, we follow the training design of Arifovic et al. (2014). During the first 10 periods of the experiment with $\rho_2 = 5$ and $\rho_2 = 8$, each subject interacts with $N - 1$ automatized players who make the savings decisions consistent with the period 2-cycle. Subjects do not see the individual decisions of the other advisers but can track the average savings decision on a graph (see Figure 27 in Appendix F). We call this training treatment **T**, and form the following hypothesis, by contrast to the treatment without training that we denote

¹³The transformation of prices (12) does not leave the return on savings invariant. Without, the values of the prices, and hence the returns on savings, along cyclical equilibria are too extreme. therefore, we only run the three different treatments in the LtOE $\rho_2 = 3, 5$ and 8 .

O.

Hypothesis 2’. (LtOE: **O** versus **T**): Subjects who have experienced an initial training phase are more likely to coordinate their savings decisions on the 2-cycle than subjects who have not.

Finally, after observing that coordination on the 2-cycle never arises in Treatment **T** with $\rho_2 = 8$, we have elaborated an additional treatment with a non-linear transformation of the savings values, in order to test whether the selection of the monetary steady state is not related to a specific feature of the model. Indeed, the equilibrium values of the real money balances along the 2-cycle are quite extreme when $\rho_2 = 8$ (roughly 1 and 146 in the payoff grid of the subjects where savings lie in $(0, 200]$). The difficulty of coordination on a periodic equilibrium may then be well explained by a so-called framing effect: when it comes to make savings decisions, following extreme variations in the decisions may appear a less natural strategy than when forecasting a given time series pattern. Even with a training phase on a 2-cycle, inertia in savings decisions or conservative strategies may be more likely to arise than a strategy that follows a 2-cycle pattern with wide oscillations. We use the following transformation for every quantity in the experimental economy:

$$\tilde{s} = 100 \times \left(\frac{s}{2}\right)^{0.16} \in]0, 100] \quad (13)$$

where \tilde{s} is the savings decision submitted by the subjects, and $s \in]0, 2]$. The resulting savings values along the 2-cycle are $\{42, 95\}$ and the steady state equals $s^* = 81$.

We call this treatment **S**, and form the following conjecture, by contrast to the training Treatment **T** with $\rho_2 = 8$:

Hypothesis 2’’. (LtOE, $\rho_2 = 8$: **T and **S**):** If variations of the savings strategies are of mild amplitude along a cycle, coordination of heterogeneous individual savings decisions on this cyclical pattern is more likely than when the amplitude of the cycle is wider.

3.4 Implementation

The experiment was programmed in Java using the PET software¹⁴ and was conducted at the CREED laboratory at the University of Amsterdam over the period November-December 2014 and February – May 2015. A total of 288 subjects were recruited from the CREED subject pool (composed of students from any field, both undergraduate and graduate) to participate in 48

¹⁴The PET software was developed by AITIA, Budapest, and is available at <http://pet.aitia.ai>, within the FP 7 European project *CRISIS*, Grant Agreement No. 288501.

experimental economies of $N = 6$ subjects each. In the LtF design, we ran 4 economies per treatment, for a total of 20 economies, and 120 subjects. For the LtO design, we ran 4 economies for each of the values $\rho_2 = \{3, 5, 8\}$ and Treatment **O**, 5 and 4 economies respectively for $\rho_2 = 5$ and 8 in treatment **T**, and 7 economies in Treatment **S**, for a total of 28 economies and 168 subjects. We run each experimental economy with $\rho_2 = 3$ for $T = 50$ generations/periods, as pilot observations indicate a very quick stabilization in the LtFE, and we run all the other treatments for $T = 100$.¹⁵ The computer interfaces of the LtF and the LtOE are reported in Appendix F, and the instructions, together with the payoff tables and the questionnaire for each design in Appendices D and E. Subjects receive a detailed description of the OLG environment underlying the experiment, their experimental task and their payoff. The consumption good is referred as "chips", following Marimon et al. (1993). The participants were given the opportunity to read the instructions at their own pace, and then were asked to fill in a quiz on paper. The instructors then checked that each subject, one by one, was able to correctly answer each question before starting the experiment. In case of a wrong answer, the experimentalist privately explained to the participant the correct answer. Only when all participants had answered all questions correctly was the experiment started. This procedure allows us to be sure that every subject has understood the economic environment underlying the experiment and his experimental task, in particular the use of the two-dimensional payoff table in the LtOE, before entering the experimental economy. Participants' payoff was expressed in points, that were converted into euros at the end of the experiment at an exchange rate that what announced in the instructions, and participants earned on average 23.6 euros. Each experimental session lasted on average for about 2 hours, including an average of 40 minutes for the instructions and questionnaire, with strong disparities across treatments (see Section 5 for details).

Table 8 in Appendix A summarizes the features of the different treatments and designs presented in this section and the parameter values, and reports the equilibrium values of prices and real money balances along the monetary steady state and the 2-cycle for the different ρ_2 values considered. We now discuss to the experimental results.

4 Experimental results

This section first provides an overview of the results in terms of aggregate convergence and coordination and then analyze individual behavior.

¹⁵Economy 1 in the LtOE with Treatment **T** and $\rho_2 = 5$ ended at period 58 due to a server crash.

4.1 Overview of experimental results

Tables 2 - 5 summarize the outcomes of the 48 experimental economies along five dimensions: the type of long-run attractor of the realized aggregate real money balances, the distance to its values along this attractor, its first-order autocorrelation throughout the experiment, the coordination between subjects' decisions (forecasts or savings) and the efficiency of the economies measured by the payoff of the subjects. For a visual representation, the dynamics of the individual price forecasts or savings decisions together with the aggregate price or savings level are reported for every single experimental economy in the figures in Appendix B.

We obtain three main outcomes from the experiments. First, *all 48 experimental economies either converge to the monetary steady state or to a 2-cycle*. We do not observe any other type of long-run dynamics in the experiments. It is worth stressing that our experiment is the first to observe spontaneous coordination of a group of subjects on a 2-cycle. This clearly shows that, *in this complex environment, only simple aggregate outcomes arise as a coordination device* in the lab, and simple equilibria may thus be viewed as most empirically relevant. These simple outcomes are all *weakly E-stable* equilibria, the steady state when $\rho_2 = 3$ and the 2-cycle when $\rho_2 = 5$ are also strongly E-stable. From the experimental evidence, it seems that neither backward perfect foresight nor forward perfect foresight dynamics provides an accurate description of the dynamics in the experimental economies. Moreover, the criterion of weak E-stability is a necessary condition for an equilibrium to be selected as a coordination device in our experiments, but not all weakly E-stable cycles are observed in our experiments.

Second, *the coordination on the 2-cycle is easier in the LtFE than in the LtOE*. In the LtFE, when the steady state is stable in the backward perfect foresight dynamics¹⁶, i.e. for $\rho_2 = 3$, all economies very quickly converge to the steady state. When the steady state is unstable (in 16 economies out of 20), we observe 14 convergences towards the 2-cycle out of 16 economies, possibly after a long transient, with first order autocorrelation of aggregate time series close to -1, and only two convergences to the steady state. A transient of at least 50 periods to the 2-cycle is observed in 6 economies (group 2 with $\rho_2 = 8$, groups 1 and 4 with $\rho_2 = 12$ and groups 1, 2 and 3 with $\rho_2 = 13.5$). The first case of convergence to the unstable steady state is observed when $\rho_2 = 5$, and occurs very quickly. However, in the second case ($\rho_2 = 8$, Group 4), we observe a "noisy" convergence towards the unstable steady state, and the aggregate price oscillates in a

¹⁶In the sequel, when we refer to a stable equilibrium, we refer to stability in the backward perfect foresight dynamics as depicted in Figure 1, which also corresponds to stable equilibrium under naive expectations, and to strong E-stability, see Section 2.

small neighborhood of the steady state, with negative autocorrelation. In the strongly unstable treatments $\rho_2 = 12$ and 13.5 , the only long run outcome is the 2-cycle. It is worth stressing that coordination on the 2-cycle arises spontaneously in the LtFE, without any training phase, or any signal akin to sunspots in the lab. By contrast, in the LtOE, the monetary steady state is the long run outcome of the aggregate dynamics of all but 3 of the 28 experimental economies. In treatment **O** (i.e. without a training phase), the economies always converge towards the steady state. For $\rho_2 = 5$, 2-cycle oscillations, with first order autocorrelation close to -1, arise after a training phase (i.e. in Treatment **T**), in two out of five cases (Groups 1 and 5). For $\rho_2 = 8$, such an outcome is observed only in Treatment **S**, once (Group 5), even if two other groups (3 and 7) display some fluctuations with a strong negative first order autocorrelation.

Third, *the LtOE display more variability than the LtFE*: both the aggregate time series display more volatility, and the individual decisions are more heterogeneous in the LtOE, compared to the LtFE. This suggests that coordination between a collection of heterogeneous players in the lab is easier in the LtF than in the LtOE.

We now take a closer look at the experimental results.

N.B.: For each experimental economy, the tables below report the type of equilibrium that is selected as a long run outcome of the dynamics in the experiments. *ADE* measures the average relative distance to this equilibrium, and is computed over the last 10 periods of each experiment (in absolute value). For instance, the ADE for an experimental economy is 0.1, aggregate savings are on average 10% away from their equilibrium values over the last 10 periods. The first order autocorrelation of average savings (denoted by ρ_s) is computed over the last 50 periods (* indicates that it is significant at 5%). The variance of aggregate savings among the N subjects (denoted by $Var(s_i)$) is evaluated in each period, and the tables report the average over the last 10 periods. In the LtFE, we compute the variance of the implied savings decisions, conditional on the price forecasts of the subjects. The earning efficiency ratios are expressed in percentage points, and measure the number of points earned by the subjects over the whole T periods of the experiments w.r.t. the maximum amount of points possible at the equilibrium (see also Subsection 4.6).

	LtFE				LtOE			
economy	1	2	3	4	1	2	3	4
type of equilibrium	steady state	steady state	steady state	steady state	steady state	steady state	steady state	steady state
<i>ADE</i>	0.000	0.000	0.001	0.000	0.055	0.195	0.029	0.073
ρ_s	-0.161	-0.082	-0.042	-0.074	0.313*	0.014	0.25	0.254
$Var(s_i)$	0.000	0.000	0.000	0.000	0.017	0.108	0.003	0.012
earnings efficiency ratio (all periods)	95.62	95.31	91.04	94.55	97.16	95.27	96.55	95.56

Table 2: $\rho_2 = 3$

	LtFE				LtOE			
economy	1	2	3	4	1	2	3	4
type of equilibrium	2-cycle	2-cycle	2-cycle	steady state	steady state	steady state	steady state	steady state
ADE	0.151	0.116	0.326	0.001	0.026	0.033	0.046	0.105
ρ_s	-0.966*	-0.962*	-0.941*	-0.329*	0.316*	0.592*	-0.024	-0.398*
$Var(s_i)$	0.008	0.002	0.011	0.000	0.00189	0.004	0.002	0.005
earnings efficiency ratio	81.78	69.23	83.89	96.78	98.1	92.4	94.05	91.34

	LtOE with training				
economy	1	2	3	4	5
type of equilibrium	2-cycle	steady state	steady state	steady state	2-cycle
ADE	0.854	0.07	0.042	0.057	1.528
ρ_s	-0.96*	-0.187	0.394*	-0.226	-0.845*
$Var(s_i)$	0.016	0.019	0.001	0.001	0.008
earnings efficiency ratio	96.9	91.31	96.57	92.6	93

Table 3: $\rho_2 = 5$

	LtFE				LtOE			
economy	1	2	3	4	1	2	3	4
type of equilibrium	2-cycle	2-cycle	2-cycle	steady state	steady state	steady state	steady state	steady state
ADE	0.041	0.169	0.022	0.532	0.04	0.089	0.039	0.062
ρ_s	-0.975*	-0.952*	-0.967*	-0.882*	-0.106	-0.134	0.527*	0.205
$Var(s_i)$	0.016	0.030	0.024	0.02	0.003	0.036	0.000	0.016
earnings efficiency ratio	77.85	87.86	80.4	98.03	93.43	88.52	81.12	96.45

	LtOE with training			
economy	1	2	3	4
type of equilibrium	steady state	steady state	steady state	steady state
ADE	0.042	0.054	0.111	0.023
ρ_s	0.072	0.186	0.47*	0.528*
$Var(s_i)$	0.000	0.006	0.072	0.001
earnings efficiency ratio	85.78	87.82	94.39	89.37

	LtOE with training and non-linear transformed savings (Tr. S)						
economy	1	2	3	4	5	6	7
type of equilibrium	steady state	steady state	steady state	steady state	2-cycle	steady state	steady state
ADE	0.153	0.068	0.34	0.071	17.901	0.039	0.362
ρ_s	0.193	0.032	-0.722*	0.001	-0.956*	0.681*	-0.663*
$Var(s_i)$	0.011	0.043	0.19	0.011	0.134	0.001	0.179
earnings efficiency ratio	69.04	84	88.5	84.89	74.67	85.73	76.6

Table 4: $\rho_2 = 8$

	$\rho_2 = 12$ (LtFE)				$\rho_2 = 13.5$ (LtFE)			
economy	1	2	3	4	1	2	3	4
type of equilibrium	2-cycle	2-cycle	2-cycle	2-cycle	2-cycle	2-cycle	2-cycle	2-cycle
ADE	0.512	0.257	0.225	0.222	0.046	0.081	0.047	0.041
ρ_s	-0.931*	-0.926*	-0.969*	-0.972*	-0.828*	-0.961*	-0.977*	-0.98*
$Var(s_i)$	0.147	0.006	0.001	0.000	0.007	0.000	0.002	0.019
earnings efficiency ratio	54.28	62.03	69.37	73.12	55.87	64.56	68.37	88.26

Table 5: $\rho_2 = 12$ and $\rho_2 = 13.5$.




Figure 2: Cumulative distribution of the relative distance of prices to equilibrium (either steady state or 2-cycle). Left panel: LtF versus LtO. Right panel: economies that converge to the steady state versus economies that converge to the 2-cycle, LtF and LtO experiments pooled together.

4.2 Aggregate convergence

Figure 2 reports the cumulative distribution of the relative distance of price to equilibrium in the experimental economies, by discarding the first 10 periods of each experiment. Figure 2a compares these distributions in LtFE versus LtOE, and a K-S test leads to reject the null hypothesis of equal distribution against the alternative hypothesis of distance values lower in the LtFE than in the LtOE (p-value = 0.0014). We can conclude that overall, *aggregate convergence is significantly better in LtFE than in LtOE*. Similarly, Figure 2b compares the distributions of distance to the monetary steady state versus the 2-cycle, by pooling LtFE and LtOE together. We can conclude that *aggregate convergence is significantly better when the dynamics converge towards the steady state, than towards the 2-cycle*.¹⁷

4.3 Coordination between subjects' decisions

We now take a closer look at the individual coordination of subjects' decisions in the experiments. In order to compare decision values that have the same order of magnitude, we report the standard deviations of the implied savings decisions, given subjects' price forecasts in the LtFE. Figure 3 displays the cumulative distribution of the standard deviations among the six individual

¹⁷The p-value of the same KS test is 0.0047.

decisions, measured in each period. Figure 3a compares the coordination between subjects in the LtFE versus the LtOE. Subjects' decisions are significantly more homogeneous in the LtFE than in the LtOE.¹⁸ In other words, subject's coordination is easier in the LtFE than in the LtOE, and *the LtOE display more heterogeneity between subjects than the LtFE*. This is striking from the comparison between the two designs with $\rho_2 = 3$, i.e. when the monetary steady state is the only equilibrium, and is stable. These economies converge almost perfectly on the steady state in the LtFE (this is the best coordination obtained among all experimental economies), but only converge to a (close) neighborhood of the steady state in the LtOE.

More homogeneous price predictions in comparison to savings decisions may be explained by two phenomena. First, we observe a bias towards round numbers in the LtOE. Recall that subjects make savings decisions with the help of a 2-D payoff, in which the savings decisions are discretized, but the instructions insist on the fact that they can submit any number (up to two digits). We find that, overall, 60% of the savings decisions are multiples of 5 (i.e. 50, 55, 60, etc.), and 47% are multiples of 10 (i.e. 50, 60, etc.). Pilot sessions using an A3 format payoff table with a thinner grid report the same type of decisions. By contrast, only 36% of price predictions are an integer, and they are concentrated at the beginning of the experiment, when subjects have only few past observations to make precise predictions. This tendency to submit round number can also be explained by the flatter payoff values in the neighborhood of the steady state under the LtOE than under the LtFE (see payoff tables in Appendix E). With flatter payoff values, the subjects have less monetary incentives to submit the exact equilibrium value. A second explanation to the more heterogeneous savings decisions with respect to price predictions could be strategic behavior from some subjects. Five subjects report in the questionnaire at the end of the experiment that they intentionally deviate from the average savings values in their experimental economy in an attempt to manipulate the return on savings. This is the case for instance in Group 2, with $\rho_2 = 3$: one subject reported that he/she made occasionally high savings decisions in an attempt to decrease the price, and increase the return on savings, in order to reach the bottom left part of the payoff table when the payoff is maximized. Those types of deviations are observed across treatments in the LtOE (for instance, based on questionnaire data, subject A2 in group 2 with $\rho_2 = 3$, or subject B3 in group 4 with $\rho_2 = 8$ and Treatment **S**). Finally, the difference in the nature of the experimental task (forecasting versus saving), and the implied cognitive load may also account for the observed differences between the two designs. This point is further developed in the conclusion of the paper.

¹⁸The corresponding p-value of the K-S test is lower than $2.2e - 16$.

Furthermore, Figure 3b compares the cumulative distribution of the standard deviations among individual decisions in the LtFE and in the LtOE, in each design, and in case of convergence to the steady state versus the 2-cycle. Coordination between subjects is significantly better on the steady state than along the 2-cycle, both in the LtFE and in the LtOE.¹⁹ The better coordination on the steady state compared to the 2-cycle can be at least partly explained by subjects' mistakes when entering their decisions in the experimental software: the likelihood of making mistakes when entering price predictions seems higher when alternatively entering a high and a low forecasts than when entering a constant number. Several subjects indeed reported in the questionnaire at the end of the experiment making typos, for instance in Group 1 of the LtFE, with $\rho_2 = 5$. However, the 2-cycle appears to be a long run outcome that is robust against individual deviations: even when the 2-cycle is temporarily "destroyed" after a subject's individual mistake, the dynamics settles down back to the 2-cycle after few periods. This is also the case in the LtOE: in Group 5 with $\rho_2 = 5$ and Treatment **T**, subject D3 reported in the questionnaire that he/she became "confused" around period 24/25, and made several typos when entering savings decisions, causing the amplitude of the cycle to diminish. However, towards the end of the session, the amplitude of the cycle increases again. We note that convergence back to the 2-cycle after a temporary deviation appears quicker in the LtFE than in the LtOE.

4.4 Estimation of individual forecasting rules in the LtFE

We then estimate forecasting rules on the individual price prediction data for the 120 subjects in the LtFE. We follow here standard practice in the related literature (see e.g. Heemeijer et al. (2009)). We discard the first 10 periods to allow for a learning phase, and estimate the general forecasting rule:

$$p_{i,t+1}^e = \alpha + \beta_{P_{t-1}} P_{t-1} + \beta_{P_{t-2}} P_{t-2} + \beta_{p_{i,t}^e} p_{i,t}^e + \beta_{p_{i,t-1}^e} p_{i,t-1}^e + \epsilon_{i,t} \quad (14)$$

where $p_{i,t+1}^e$ is the price forecast made by subject i at the beginning of period t for period $t + 1$, P_{t-1} the last observable price in period $t - 1$, P_{t-2} the price in period $t - 2$, $p_{i,t}^e$ the last price forecast made in period $t - 1$ for period t , $p_{i,t-1}^e$ the price forecast made in period $t - 2$ for period $t - 1$ and $\epsilon_{i,t}$ a noise term. We use the heteroskedasticity and autocorrelation consistent (HAC) estimator of the R package `sandwich` (Zeileis 2004). We use the Ljung Box

¹⁹Both in the LtOE and in the LtFE, the K-S tests with the alternative hypothesis that values of the standard deviations are lower on the steady state than on the 2-cycle gives a p-value smaller than $2.2e - 16$.




Figure 3: Cumulative distribution of the standard deviations of individual decisions in each period. Left panel: LtF (savings decisions implied by price forecasts) versus LtO (submitted savings decisions), all economies. Right panel: economies that converge to the steady state only versus 2-cycle, distinguishing LtF and LtO.

test for autocorrelation with 4 lags, which is consistent with the number of observations of each regression. We successively drop the non-significant variables and re-estimate the rule (14) until only significant variables remain present. We adopt a 5% confidence level for the whole econometric analysis.

The general rule (14) includes as a special case the second-order adaptive rule that is consistent with learning a 2-cycle pattern when it exists (see Equation (7), and Guesnerie & Woodford (1991)). A subject is said to use such a rule if the following constraints on the estimated coefficients result from the econometric estimation of (14): $\hat{\beta}_{p_{t-1}} + \hat{\beta}_{p_{t-1}^e} = 1$, $\hat{\beta}_{p_{t-1}}, \hat{\beta}_{p_{t-1}^e} \in]0, 1[$, $\hat{\beta}_{p_{t-2}} = \hat{\beta}_{p_t^e} = \hat{\alpha} = 0$. In the special case of $\hat{\beta}_{p_{t-1}} = 1$, $\hat{\beta}_{p_{t-2}} = \hat{\beta}_{p_t^e} = \hat{\beta}_{p_{t-1}^e} = \hat{\alpha} = 0$, the subject uses naive expectations, which delivers convergence to a two-cycle only in case of $\rho_2 = 5$. A stable AR(1) forecasting rule corresponds to $\hat{\beta}_{p_{t-1}} < 1$, $\hat{\beta}_{p_{t-2}} = \hat{\beta}_{p_t^e} = \hat{\beta}_{p_{t-1}^e} = 0$.

The distributions of the estimates of the parameters of the general rule for the 120 subjects are reported in Appendix C.1. First, we estimate the general rule (14) for all groups in the LtFE that converge towards a 2-cycle, i.e. 14 economies, for a total of 84 subjects. Table 6 classifies the subjects according to their corresponding rule. Given the learning predictions in theoretical models discussed in Section 2, the second-order adaptive rule, naive expectations and the AR(1) rule constitute our three benchmark rules in the case of a convergence to the 2-cycle, with 90% of the subjects falling into one of these categories. More than half of the

subjects (i.e. 47) follow a second-order adaptive rule, and we cannot reject the joint hypothesis $\hat{\beta}_{p_{t-1}} + \hat{\beta}_{p_{t-1}^e} = 1$, $\hat{\beta}_{p_{t-2}} = \hat{\beta}_{p_t^e} = \hat{\alpha} = 0$ at 5% for all but 6 of them. Almost all those subjects belong to economies with $\rho_2 = 8, 12$ or 13.5 , which is consistent with the theoretical predictions of second-order adaptive expectations (Guesnerie & Woodford 1991) and the observed convergence to the 2-cycle. Recall the stability condition of the 2-cycle under the adaptive rule (7) given in Sub-section 2.4: the stability of the 2-cycle depends on the weight put on P_{t-1} (w in Equation (7)). Intuitively, the 2-cycle is stable if the weight on the past observed price is not too high, and this weight has to be lower and lower when the 2-cycle becomes "more and more unstable". This theoretical prediction is in line with our estimations in the experiments. Coefficients on P_{t-1} are lower, the higher ρ_2 , and the average coefficient values for each ρ_2 treatment is always lower than the corresponding threshold \underline{w} .²⁰

A quarter of the subjects (20) use naive expectations, and they mostly belong to economies with $\rho_2 = 5$, for which the 2-cycle is stable under naive expectations, and stable under the adaptive rule (7) for all w values. For Group 2 with $\rho_2 = 5$, the results are less obvious, and this is also the group within this treatment in which the 2-cycle appears the most "noisy": 4 out of 6 subjects use an AR(1) rule without constant but with a coefficient on P_{t-1} close but significantly lower than unity. This may suggest that the learning process that converges to the 2-cycle, and to naive expectations has not been completed for those subjects.²¹ In most groups however, the results of the estimation are clear-cut. For instance, in Group 1 with $\rho_2 = 5$, all 6 subjects use naive expectations, which is theoretically consistent with the observed convergence to the two-cycle. In Group 4 with $\rho_2 = 13.5$, all subjects have second-order adaptive expectations, that delivers convergence to the two-cycle under learning. Looking at the corresponding time series in Appendix B.1, those two groups are also the groups that converge faster to the 2-cycle in their respective treatment. Figure 4a displays the estimated values of the coefficients associated to P_{t-1} against the ones of P_{t-1}^e for all the 84 subjects. It clearly shows that most points are scattered around the dashed line $y = 1 - x$, which corresponds to the second-order adaptive rule. We also observe a concentration of points around $(1, 0)$ for the case $\rho_2 = 5$, which corresponds to naive expectations, while the points are more scattered along the dashed line for higher values

²⁰More precisely, the average coefficient associated to P_{t-1} is 0.6603 in Group 1, 2 and 3 with $\rho_2 = 8$ ($\underline{w} \simeq 0.8$), 0.6016 ($\underline{w} \simeq 0.61$) across all groups with $\rho_2 = 12$ and 0.5583 ($\underline{w} \simeq 0.57$) when $\rho_2 = 13.5$. KS-tests indicate that these coefficients are all significantly lower than for $\rho_2 = 5$ (where the average value of this coefficient is 0.9076). Additionally, coefficients on P_{t-1} are significantly lower when $\rho_2 = 13.5$ than 8, but other pair-differences across treatments are not significant at 5 %.

²¹The point that the degree of homogeneity of forecasting rules among the subjects matters for the coordination on an outcome has been made by Marimon et al. (1993) who estimate similar forecasting rules in their LtFE.

		2nd order adaptive rule $\hat{\beta}_{p_{t-1}} + \hat{\beta}_{p_{t-1}^e} = 1$ $\hat{\beta}_{p_{t-2}} = \hat{\beta}_{p_t^e} = \hat{\alpha} = 0$	naive expectations $\hat{\beta}_{p_{t-2}} = \hat{\beta}_{p_t^e}$ $\hat{\beta}_{p_{t-1}} = 1$	stable AR(1) rule without constant $\hat{\beta}_{p_{t-1}} = \hat{\alpha} = 0$ $\hat{\beta}_{p_{t-1}} < 1$	stable AR(1) rule with constant $\hat{\beta}_{p_{t-2}} = \hat{\beta}_{p_t^e} = \hat{\beta}_{p_{t-1}^e} = 0$ $\hat{\beta}_{p_{t-1}} < 1, \hat{\alpha} > 0$	mixed rule – other combinations of significant variables
$\rho_2 = 5$	gp 1		6			
	gp 2	1	1	4		
	gp 3	1	4			1
$\rho_2 = 8$	gp 1	4	1			1
	gp 2	5	1			
	gp 3	5	1			
$\rho_2 = 12$	gp 1	3			1	2
	gp 2	5		1		
	gp 3	4	1		1	
	gp 4	2	2	1		1
$\rho_2 = 13.5$	gp 1	2	2			2
	gp 2	4	1	1		
	gp 3	5				1
	gp 4	6				
TOTAL		47 (56%)	20 (24%)	7 (8%)	2 (2%)	8 (10%)

Table 6: Distribution of forecasting rules among the subjects in the LtFE

of ρ_2 , which is line with the stability conditions discussed above.

The remaining 8 subjects (10%) use a mixed forecasting rule, half of them belong to Group 1 with $\rho_2 = 12$ and 13.5, in which we observe a particularly long transient with irregular movements in price before convergence to the 2-cycle.

The estimation of the rule (14) is less meaningful for the remaining 36 (out of 120) subjects, for which we observe a convergence to the steady state, as their predictions quickly become essentially constant over time (with the exception of Group 4, with $\rho_2 = 8$ where small oscillations persist throughout the experiment). Excluding this latter group, 26 subjects out of the remaining 30 have an implied value of savings (consistent with their price forecast) within 0.01 of the steady state value. In these groups, after 10 experimental periods, 98% of the price predictions are within the steady state value ± 1 , and 84% in a neighbourhood of 0.1. Therefore, for these 36 subjects, we perform the following exercise. We compute the long run estimated price level of Equation (14):

$$P^{**} \equiv \frac{\alpha}{1 - \hat{\beta}_{p_{t-1}} - \hat{\beta}_{p_{t-1}^e} - \hat{\beta}_{p_{t-2}} - \hat{\beta}_{p_t^e}} \quad (15)$$

We use the implied savings values for the LtFE in order to compare with the LtOE below. The distribution of the relative distance of the estimated long-run equilibrium to the steady state for the 36 subjects is reported in Figure 4b. Overall, we cannot reject the hypothesis that the long run savings level is equal to steady state.²²

²²The average relative distance to steady state of the estimated long run equilibrium of savings equals 0.003, and the p-value of a bilateral Wilcoxon signed rank test is 0.1252.




Figure 4: Outcomes of the estimations of individual forecasting rules. Left panel: Scatter plot of the estimated coefficients $\hat{\beta}_{p_{t-1}}$ and $\hat{\beta}_{p_{t-1}^e}$ in rule (14) for the economies that converge to a 2-cycle (i.e. all groups with $\rho_2 = 12$ and $\rho_2 = 13.5$ and Groups 1-2-3 with $\rho_2 = 5$ and $\rho_2 = 8$, 84 observations). The dotted gray line represents the locus of points for which $\hat{\beta}_{p_{t-1}} + \hat{\beta}_{p_{t-1}^e} = 1$. Right panel: frequency distribution of the distance between the steady value of average savings and the long run savings equilibrium M/p^{eq} implied by (14) for the economies which converge to the steady state (i.e. all groups with $\rho_2 = 3$ and Group 4 for $\rho_2 = 5$ and $\rho_2 = 8$, 36 observations).

4.5 Estimation of individual savings rules in the LtOE

We estimate savings rules from the individual savings decisions made by the subjects in the LtOE by following exactly the same procedure as for the estimations of the forecasting rules in the LtFE. We estimate the following behavioural rule for each participant:²³

$$s_{i,t} = \alpha + \beta_{s_{t-1}} s_{i,t-1} + \beta_{s_{t-2}} s_{i,t-2} + \beta_{R_{t-1}} R_{t-1} + \varepsilon_{i,t} \quad (16)$$

where $s_{i,t}$ is the savings decision made by subject i at the beginning of period t for period t , $R_{t-1} \equiv \frac{P_{t-1}}{P_{t-2}}$ the return on savings between period $t-2$ and period $t-1$ and $\varepsilon_{i,t}$ a noise term. We include two lagged values of the individual savings decisions because they are relevant along a 2-cycle, and we include R_{t-1} as this is the last observable return on savings that subjects have (and is displayed on their screen). The general rule (16) embeds a constant rule if the joint constraint $\hat{\beta}_{s_{t-1}} = \hat{\beta}_{s_{t-2}} = \hat{\beta}_{R_{t-1}} = 0$ results from the estimation, a stable AR(1) rule if $\hat{\beta}_{R_{t-1}} = \hat{\beta}_{s_{t-2}} = 0$ and $|\hat{\beta}_{s_{t-1}}| \in (0, 1)$, a stable AR(2) rule if $\hat{\beta}_{R_{t-1}} = 0$ and $|\hat{\beta}_{s_{t-2}} + \hat{\beta}_{s_{t-1}}| \in (0, 1)$.

²³For instance, Bao et al. (2016) estimate a similar rule with an AR(1) structure for quantity decisions in an LtOE in an asset pricing model.

		intercept only $\hat{\beta}_{s_{t-1}} = \hat{\beta}_{s_{t-2}} = 0$	stable AR(1) rule $\hat{\beta}_{s_{t-2}} = 0, \hat{\beta}_{s_{t-1}} \in]0, 1[$ $\hat{\beta}_{R_{t-1}} = 0$	stable AR(2) rule $\hat{\beta}_{s_{t-2}} \neq 0, \hat{\beta}_{s_{t-1}} \in [0, 1[$	mixed rule $ \hat{\beta}_{s_{t-2}} , \hat{\beta}_{s_{t-1}} \in [0, 1[$ $\hat{\beta}_{R_{t-1}} > 0$	mixed rule $\hat{\beta}_{R_{t-1}} < 0$
$\rho_2 = 3$	gp 1	2	1	1	2	
	gp 2	2	1	1	1	1
	gp 3	2	2	1		1
	gp 4	2		2		2
$\rho_2 = 5$	gp 1/O	2		2	1	1
	gp 2/O	2	1	1	1	1
	gp 3/O	3	1		1	1
	gp 4/O	1	1			4
	gp 1/T		4	2		
	gp 2/T	1	2	2		1
	gp 3/T	1	2	2	1	
	gp 4/T	1	3	1		1
	gp 5/T		1	5		
$\rho_2 = 8$	gp 1/O	1			2	3
	gp 2/O	2	2		1	1
	gp 3/O	2	1		1	2
	gp 4/O	3				2
	gp 1/T		3	2		1
	gp 2/T	1	3	2		
	gp 3/T	2	1			3
	gp 4/T	1	2	2		1
	gp 1/S	1	2	1	1	1
	gp 2/S		1		2	3
	gp 3/S			2	1	3
	gp 4/S	1	2			3
	gp 5/S			6		
	gp 6/S	1	1	2	1	1
	gp 7/S		1	2	1	2
TOTAL		34 (20%)	38 (23%)	40 (24%)	17 (10%)	39 (23%)

Table 7: Distribution of savings rules among the subjects in the LtOE. the groups that converge towards the 2-cycle are highlighted in bold.

We estimate a savings rule for the 168 subjects who participated in the LtOE, 150 in economies that converge or oscillate around the steady state, and 18 in economies that display regular up-and-down oscillations in a neighbourhood of the 2-cycle.²⁴

Appendix C.2 reports descriptive statistics of the individual savings time series, and the distributions of the estimates of the parameters of the general rule for the 168 subjects. First, we notice that the intercept is always significant, and for all but one subject is positive. A high concentration of the estimated values is observed between 0.5 and 0.6, which broadly corresponds to the steady state values of savings. Table 7 classifies the subjects according to their corresponding savings rule. By considering first the 150 estimates at the steady state, 94 subjects (63%) are characterized by an AR rule with significant intercept ($\hat{\beta}_{R_{t-1}} = 0$): 34 of them use a constant rule, 33 use an AR(1) rule, and 27 an AR(2) rule. For those 94 cases, we compute the estimated long run equilibrium value of savings from Rule (16) (with $\hat{\beta}_{R_{t-1}} = 0$) as $s^{**} \equiv \frac{\alpha}{1 - \hat{\beta}_{s_{i,t-1}} - \hat{\beta}_{s_{i,t-2}}}$.

²⁴Since savings decisions are more variable than price predictions in economies that converge to the steady state, the estimation of (16) is less problematic in this case than in the LtFE. In the LtOE, only 3 subjects have strictly constant savings decisions after the first 10 periods.

Figure 5a displays the frequency distribution of the relative distance of these long run savings equilibrium estimates to the steady value of the savings. The average relative distance equals -0.01 , and we cannot reject the null hypothesis that it equals zero.²⁵

The other 56 subjects among the steady state economies use a mixed rule, i.e. their savings decisions are well described by an AR rule with a significant reaction to past values of the return on savings R_{t-1} . Overall, the estimated coefficients associated to R_{t-1} are significantly negative.²⁶ If we assume that participants at least partly based their predictions of the return on savings on the last observed value R_{t-1} , negative coefficients are consistent with the offer curve displayed in the two-dimensional payoff table: a higher expected return on savings corresponds to a lower savings decision (bottom left corner of the payoff table), and reciprocally as long as the expected return is not too small.




Figure 5: Outcomes of the estimations of individual savings rules. Left panel: Frequency distribution of the relative distance between the steady value of the savings and the long run estimated savings equilibrium s^{**} for the economies which converge to the steady state and the 94 subjects that are characterized by an AR(1) or AR(2) rule. Right panel: Frequency distribution of the distance between the average of the 2-cycle values of e savings and the long run savings equilibrium s^{eq} in (16) for the economies which converge to the 2-cycle, 18 subjects.

In the three economies for which the dynamics corresponds to a (noisy) 2-cycle (i.e. for which the first order autocorrelation of the aggregate savings is close to -1), the 18 participants all use an AR(1) or an AR(2) rule, all the estimated coefficients associated to s_{t-1} are significantly negative, and all those associated to s_{t-2} are significantly positive.²⁷ These signs are consistent with a

²⁵A two-sided Wilcoxon rank sum test gives a p-value of 0.6481.

²⁶39 out of the 56 coefficients are significantly negative, and the p-value of the associated unilateral Wilcoxon signed rank test is 0.0244.

²⁷The average estimates of $\hat{\beta}_{s_{t-1}}$ and $\hat{\beta}_{s_{t-2}}$, respectively in group 1 with $\rho_2 = 5$ and Treatment **T**, in

two-cycle type of dynamics – see Hommes et al. (2013) for a discussion of AR rule associated to 2-cycle dynamics. We then compute the average relative distance of the long run estimated savings equilibrium s^{**} to the average value of savings along the 2-cycle for each of the three economies (see Figure 5b). This distance is very small for two of three groups, i.e. -0.0188 in group 1 with $\rho_2 = 5$ /Treatment **T** and -0.018 in group 5 with $\rho_2 = 8$ /Treatment **S**. It is larger for group 5 with $\rho_2 = 5$ /Treatment **T** (-0.1337). These estimates are fully consistent with the observed patterns in those groups: distances are always negative as the up-and-down oscillations never overshoot the theoretical 2-cycle, and we observe the wider oscillations in group 1 with $\rho_2 = 5$ /Treatment **T** and group 5 with $\rho_2 = 8$ /Treatment **S**.

4.6 Earnings

In order to evaluate the efficiency of the participants in performing their experimental task, we take a closer look at the earnings efficiency ratios (reported in Tables 2-5) for each experimental economy. The earning efficiency ratios measure the average percentage of points earned by the subjects out of the maximum possible during the whole experiment. In the LtFE, at any period, the maximum points are 1300 in case of perfect prediction. In the LtOE, the maximum points are given by the transformed values of utility function on the payoff table.²⁸

Figure 6 reports the distributions of the earnings efficiency ratios in the experiments. As depicted in Figure 6a, efficiency is higher in the LtO than in the LtFE.²⁹ This can be explained by taking a closer look at the efficiency ratios at the steady state and along the 2-cycle. First, as shown in Figure 6b, the earnings efficiency ratios are significantly higher in case of convergence to the steady state than to the 2-cycle, and as detailed above, we observe many more 2-cycle dynamics in the LtFE than in the LtOE, in which most groups converge to the steady state.³⁰ Convergence is much quicker towards the steady state than towards the 2-cycle, as reflected by the long transient periods observed in the LtFE. Once at the steady state in the LtFE, the price dynamics is constant as the model is deterministic, subjects make perfect forecast and maximise their payoff. By contrast, the convergence towards the 2-cycle takes more periods, the transient phases are characterized by irregular price oscillations, without any clear pattern (see e.g. Groups 1 with $\rho_2 = 12$ and 13.5) and subjects make large forecasts errors. Therefore,

group 5 with $\rho_2 = 5$ and Treatment **T** and in group 5 with $\rho_2 = 8$ and Treatment **S** are -0.7172 , -0.3753 and -0.3646 , and 0.1646 , 0.4777 , and 0.5135 .

²⁸Recall that the utility functions have been transformed to allow similar payoff at the steady state and along the 2-cycle, see Section 3.

²⁹The p-value of the associated unilateral Wilcoxon rank sum test is 0.00444.

³⁰The p-value of the associated unilateral Wilcoxon rank sum test is 0.00003.




Figure 6: Distribution of the earnings efficiency ratios. Left panel: LtF (20 observations) versus LtO (28 observations); Middle panel: steady states (31 observations) versus 2-cycle (17 observations), LtFE and LtOE pooled together; Right panel: with (12 observations) versus without (16 observations) training in the LtOE.

the LtFE with $\rho_2 = 12$ and $\rho_2 = 13.5$ display the lowest earnings efficiency ratios across all experimental sessions.³¹ By contrast, most economies converge to the steady state in LtOE, the subjects' payoff is flat around the steady state (essentially between 0.5 and 0.6), and therefore near-maximized despite small persistent deviations from the optimal savings decision.

Figure 6c focuses on the LtOE, and compares economies without (Treatment **O**) and with an initial training phase (Treatments **T** and **S**). Training significantly increases the earnings efficiency ratios.³² During the 10 period training phase, aggregate savings, price and the corresponding returns on savings vary widely. We may then conjecture that the subjects experienced a broader range of values of the return on savings, and got more acquainted with the payoff table with than without the training phase. This may explain why training increases efficiency in the LtOE.

5 A behavioural explanation of the LtFE

To explain the LtFEs, this section presents a behavioral model in which agents switch between different forecasting rules based upon their relative performance.

Section 4.4 shows that many subjects use adaptive expectations as in Equation (7). Coordination on a steady state or on a 2-cycle in the LtFE may thus be explained by coordination on a very simple adaptive expectation rule. However, this rule does not provide an explanation of

³¹However, differences in earnings efficiency ratios are not significant at 5% across ρ_2 values for the groups that converge to the 2-cycle in the LtFE.

³²The p-value of the corresponding unilateral Wilcoxon rank sum test is $3.3e - 08$.

how subjects come to learn to use the order $k = 2$ and the coefficient w of the rule.

Hommes et al. (2013) provide an alternative explanation of coordination on a steady state or a two-cycle in Grandmont's OLG model, where agents learn the two parameters of an AR(1) forecasting rule through sample autocorrelation (SAC) learning. Under SAC learning, the perceived law of motion of the price is an AR(1) process:

$$P_t = \alpha + \beta(P_{t-1} - \alpha) + \epsilon_t \quad (17)$$

where ϵ is an *iid* error term, α, β are real numbers, α being the unconditional mean of the price and $\beta \in (-1, 1)$ the first order autocorrelation coefficient. Given (17), the two-period ahead price forecast that minimizes the MSE is:

$$P_{t+1}^e = \alpha + \beta^2(P_{t-1} - \alpha). \quad (18)$$

Agents do not know the true values of α and β but estimate them based on past price observations using two simple observable statistics: the sample average and the first-order sample autocorrelation. Hommes & Zhu (2014) present a constant gain version of sample auto-correlation (SAC) learning:

$$\begin{cases} \alpha_t = \alpha_{t-1} + \kappa(x_t - \alpha_{t-1}) \\ \beta_t = \beta_{t-1} + \kappa R_t^{-1} ((x_t - \alpha_{t-1})(x_{t-1} - \alpha_{t-1}) - \beta_{t-1}(x_t - \alpha_{t-1})^2) \\ \quad - \kappa^2 R_t^{-1} (x_0 - \alpha_{t-1} + \beta_{t-1}(x_t - \alpha_{t-1})^2) + \kappa^3 R_t^{-1} (\alpha_{t-1} - x_t) \\ R_t = R_{t-1} + \kappa ((x_t - \alpha_{t-1})^2 - R_{t-1}) - \kappa^2 (x_t - \alpha_{t-1})^2 \end{cases} \quad (19)$$

where $R_t \equiv \frac{1}{1+t} \sum_{i=0}^t (x_i - \alpha_t)^2$ is the price variance. The parameter κ is a positive gain.³³ Hommes & Zhu (2014) stress the behavioral interpretation of SAC-learning with constant gain: agents "guestimate" sample average and first-order auto-correlation from the observed time series. This interpretation explains why the empirical ACF in the lab price forecasts (see Figure 22) is indeed consistent with the strongly negative autocorrelation of prices. For groups converging to steady state, the sample average is close to the steady state price. For groups converging to the 2-cycle, the sample average tends toward the average of the price values along the 2-cycle and the first order sample autocorrelation is strongly negative, and in fact close to -1 . These observations are consistent with the definition of a first-order Consistent Expectation Equilib-

³³In the case of decreasing gain, we have $\kappa = \frac{1}{1+t}$.

rium (CEE) given in Hommes & Sorger (1998): a first-order CEE is a sequence of prices $(p_t)_{t=0}^{+\infty}$ that satisfies the actual law of motion of the price, for which α_t converge to the sample average and β_t to the sample first-order autocorrelation. In particular, if a 2-cycle $\{p_1, p_2\}$ is a CEE, then $\lim \alpha_t = (p_1 + p_2)/2$ and $\lim \beta_t = -1$.

In order to measure the relative contribution of the AR(1) rule under SAC learning and second-order adaptive expectations, we fit an heuristic switching model to the experimental data by extending Anufriev & Hommes (2012). We add these two rules to the switching model to test how much each one contributes to the explanation of the observed lab forecasts. The model includes six rules that cover the different types of behavior that have been identified in many previous LtFEs (Hommes 2011). These rules are: (i) an anchoring and adjustment rule $p_{t+1}^e = \frac{p_{t-1}^{av} + p_{t-1}}{2} + (p_{t-1} - p_{t-2})$ (p_{t-1}^{av} being the average of all past prices), (ii) a weak trend-following rule $p_{t+1}^e = p_{t-1} + 0.4(p_{t-1} - p_{t-2})$, (iii) a strong trend-following rule $p_{t+1}^e = p_{t-1} + 1.3(p_{t-1} - p_{t-2})$, (iv) second-order adaptive expectations $p_{t+1}^e = 0.5p_{t-1} + 0.5p_{t-1}^e$ (where we choose a weight $w = 0.5$, consistent with convergence to the 2-cycle in all ρ_2 treatments), (v) naive expectations $p_{t+1}^e = p_{t-1}$, and (vi) a SAC-learning rule (18). Each rule $i = 1, \dots, 6$ is evaluated in each period by its forecasting accuracy:

$$U_{i,t-1} = -(p_{t-1} - p_{i,t-1}^e)^2 + \eta U_{i,t-2} \quad (20)$$

And the fraction of each rule is updated according to a discrete choice model with asynchronous updating:

$$n_{i,t} = \delta n_{i,t-1} - (1 - \delta) \frac{\exp(\beta U_{i,t-1})}{\sum_{i=1}^n \exp(\beta U_{i,t-1})} \quad (21)$$

so that better performing rules spread into the population of rules at the expense of poor performers. Aggregate price expectations are the average of the six predictions weighted by their fraction n_i (Brock & Hommes 1997). We use $\beta = 0.3$, $\eta = 0.5$ and $\delta = 0.7$.³⁴ We report the results for three typical groups: convergence to the steady state ($\rho_2 = 5$, Gp. 4) in Figure 7, convergence to a "noisy" 2-cycle ($\rho_2 = 5$, Gp. 1) in Figure 8, and quick convergence to the 2-cycle ($\rho_2 = 13.5$, Gp. 4) in Figure 9. We report the experimental prices together with the heuristic switching model forecasts (left panel), the evolution of the fractions of the six rules (middle panel) and the evolution of α_t and β_t of the SAC rule (right panel).

³⁴This calibration is close to the one used in Anufriev & Hommes (2012) but implies more sensitivity of the switching process, which appears consistent with the experimental data. The qualitative picture that we discuss from these simulations is not affected by changes in these parameter values. For the gain coefficient in Equation (19), we use $\kappa = 0.2$.




Figure 7: $\rho_2 = 5$, Gp. 4 (quick convergence to the steady state).
NB: Right panel, steady state price in dotted red line.




Figure 8: $\rho_2 = 5$, Gp. 1 (convergence to a "noisy" 2-cycle).
NB: Right panel, average price on the 2-cycle in dotted red line.




Figure 9: $\rho_2 = 13.5$, Gp. 4 (quick convergence to the 2-cycle). Left panel: heuristic switching model forecasts versus experimental prices; Middle panel: evolution of the share of each of the six rules; Right panel: evolution of β and α of the SAC rule, average price on the 2-cycle in dotted red line.

We observe first, from the left panels, that the forecasts from the heuristic switching model fit the experimental data quite nicely. From the middle panels, we see that the SAC learning rule (blue dashed line) is always dominant, i.e. is always the best performing forecasting rule, at the beginning of the experiment (typically within the first twenty periods), regardless of whether the experiment converges on the steady state or the 2-cycle. It means that the appearance of either the steady state or the 2-cycle as long run outcome of the experiment always start with coordination on a simple AR(1) rule. However, the evolution of the values of α_t and β_t differ (right panels): in the case of the steady state, we observe $\alpha_t \rightarrow p^*$ and $\beta_t \rightarrow 1$ (Figure 7), while we observe $\alpha_t \rightarrow \frac{p_1 + p_2}{2}$ and $\beta \rightarrow -1$ in Figures 8 and 9 along the 2-cycle.

In the case of the steady state (Figure 7), agents in the simulations learn the steady state value

and the price converges quickly (all rules give then a perfect prediction as the price is constant, so that all fractions converge to $1/6$). By contrast, in the case of the 2-cycle, agents progressively update the first autocorrelation parameter and learn the strong negative autocorrelation pattern of the price which amplifies the up and down oscillations of their forecasts and the realized prices. Either naive expectations (in the case of $\rho_2 = 5$, Figure 8, middle panel) or the second order adaptive expectations (in the case of $\rho_2 = 13.5$, Figure 9, middle panel) then dominates and the price dynamics settles on the 2-cycle. After convergence to the 2-cycle, naive and adaptive expectations deliver the same, perfect prediction and their fraction converges to $1/2$. Notice that the two trend rules or the anchoring rule perform particularly badly in this experimental environment, and their fraction are close to zero.

We conclude that an initial coordination on a simple AR(1), whose parameters are updated with observed average and first-order autocorrelation, followed by coordination on a simple second-order adaptive rule once the up-and-down pattern of the prices has been learned, provides an interesting behavioral explanation of the LtFEs.

6 Conclusions

This experimental study adds to the literature on equilibrium selection, and provides an empirical test of learning predictions. We design an experiment in the well-know complex OLG environment first studied by Grandmont (1985), which exhibits infinitely many periodic, and even chaotic equilibria, together with a steady state. Adaptive learning theory supports most of these equilibria as stable outcomes of some suitable learning process. We compare two designs of the experiment: a *learning-to-forecast* design in which subjects submit price predictions and real money balances are optimally computed, and a *learning-to-optimize* design in which subjects directly submit savings decisions. Our lab experiments form an empirical test of the most likely outcomes of the coordination process of a group of individuals in a complex environment.

We find two major results. First, subjects tend to coordinate on simple equilibria: they either select the monetary steady state or the 2-cycle, possibly after a long transient, but never coordinate on any other, more sophisticated equilibria. This shows that subjects are able to coordinate in a complex environment, but only simple equilibria may be regarded as empirically relevant, even if many other outcomes are theoretically possible under specific assumptions about adaptive learning. It is worth stressing that this is the first lab experiment that reports spontaneous coordination on a 2-cycle equilibrium of a group of individuals without any exogenous

shocks or sunspot announcements.

The second major finding is the differences between the LtF and the LtO designs. While subjects mostly coordinate on the 2-cycle in the LtFE, even if it is unstable under learning, they mostly coordinate on the (unstable) steady state in the LtOE. We provide three potential explanations of these differences.

A first explanation could be a *framing effect*. Cautious or conservative behaviour may appear more natural when it comes to making savings decisions than to making forecasts and tracking a time series pattern. Relatively stable savings decisions from one period to the next drives the dynamics towards the steady state by driving the return on savings towards unity.³⁵ By contrast, as discussed in Section 3.3, the forecasting task is akin to pattern recognition for which cycles of low periods are more likely to be spotted and tracked by human subjects.

A second explanation could be the difference in *cognitive load* implied by the two experimental tasks. We report a significantly higher cognitive load in the LtO than in the LtF design using two measures: the cumulative distribution of individual decision times and the length of the instructions before starting the experiment across the two designs (see Figure 10). Subjects read the instructions, complete the quiz and make their decisions in the LtFE more quickly despite the more complicated equilibrium on which they often coordinate (the 2-cycle) vis-à-vis the steady state in most LtOE³⁶ This may suggest that the more sophisticated the experimental task and the higher the implied cognitive load, the simpler the equilibrium subjects are likely to coordinate on, and the simpler their behavioural rules.

A third explanation, namely *risk aversion*, occurred to us when studying the questionnaire that subjects have to fill in at the end of the experiment. The experimental environment is deterministic, but uncertainty arises from others' actions as long as there is no perfect coordination between the six subjects. Even if the payoff function is designed in the LtOE to give rise to the same average payoff values at the steady state and along the 2-cycle, subjects may perceive differently the 2-cycle payoff outcome. Along the 2-cycle, the *average* payoff over two periods is the same as the steady state payoff, but this requires that all subjects will not deviate and coordinate on the 2-cycle. By contrast, at the steady state, subjects get the same payoff

³⁵This could also explain why savings exhibit a stronger negative autocorrelation in Treatment **S**, where the 2-cycle values are rescaled to less extreme values, than in Treatments **T** and **O** with $\rho_2 = 8$.

³⁶The average individual decision time in the LtFE is 19.9 seconds while it is 24.71 in the LtOE, and it took on average 33.5 minutes for the subjects to read the instructions and answer the quiz in the LtFE, while it took them on average 42.4 minutes to do so in the LtOE. The p-value of the corresponding unilateral K-S test is less than 2e-16 in the two cases. Note that the instructions are slightly longer (half of page) in the LtOE but it probably cannot account alone for almost 10 minutes more.

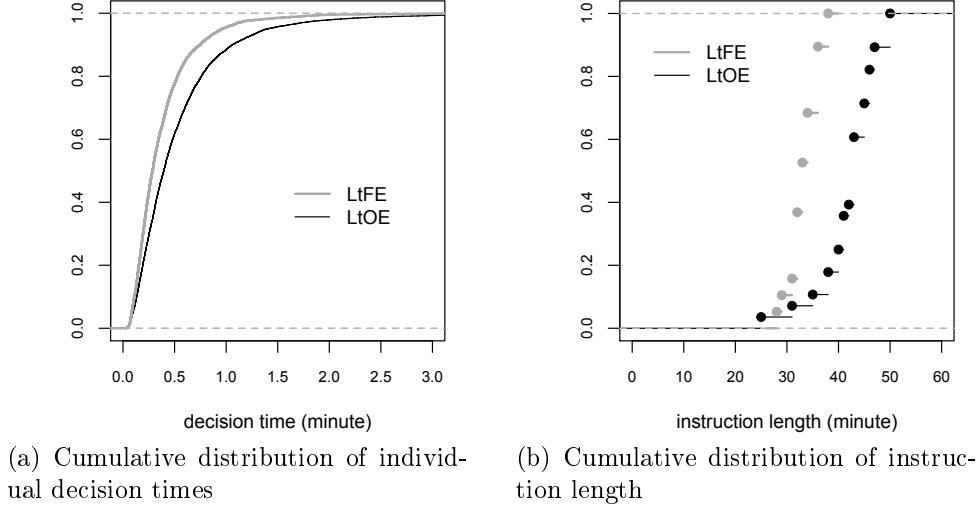


Figure 10: Measurement of cognitive loads in LtFE vs. LtOE

every period, provided that the others will keep on playing the same strategy. In this sense, coordination on the steady state may appear less risky to the subjects. Indeed, in the questionnaire, subjects exposed to the large variations in return on savings during the first 10 training periods reported several times that they were trying to "secure a smooth payoff", "hold on to an equilibrium situation", "have a sure payment", or "avoid fluctuations". In the LtFE, we do not observe such a risk aversion for price cycles because pattern recognition is likely to offset it by making coordination on a 2-cycle easier.

The differences between the LtF and the LtO designs may suggest that insights can be only derived from the LtO, as in economic models, agents are optimizers. However, we believe that this would be a misleading interpretation of our results. Instead, our results show that the differences between LtF and LtO depend on the underlying framework. We find that within a complex OLG framework, LtO favors simpler equilibria. Previous experimental work shows that LtO may lead to more complex behavior than LtF. For example, in the negative feedback cobweb experiments (Bao et al. 2013), convergence to steady state in LtO is slower than in LtF, while in the positive feedback asset-pricing experiment (Bao et al. 2016), volatility of asset prices is significantly higher in LtO than in LtF.

Whether our experimental results in terms of equilibrium selection remain valid in a larger class of contexts, e.g. in multi-dimensional expectation systems, constitutes an interesting follow-up question that is left for future research.

References

- Anufriev, M. & Hommes, C. (2012), ‘Evolutionary Selection of Individual Expectations and Aggregate Outcomes in Asset Pricing Experiments’, *American Economic Journal: Microeconomics* **4**(4), 35–64.
- Arifovic, J., Evans, G. & Kostyshyna, O. (2014), Are sunspots Learnable? An Experimental Investigation in a Simple General Equilibrium Model. Manuscript, Simon Fraser University.
- Azariadis, C. & Guesnerie, R. (1986), ‘Sunspots and Cycles’, *Review of Economic Studies* **53**(5), 725–37.
- Bao, T., Duffy, J. & Hommes, C. (2013), ‘Learning, forecasting and optimizing: An experimental study’, *European Economic Review* **61**, 186–204.
- Bao, T., Hommes, C. & Makarewicz, T. (2016), ‘Bubble formation and (in)efficient markets in learning-to-forecast and -optimize experiments’, *Economic Journal* . forthcoming.
- Brock, W. & Hommes, C. (1997), ‘A Rational Route to Randomness’, *Econometrica* **65**(5), 1059–1096.
- Bullard, J. & Duffy, J. (1998), ‘Learning And The Stability Of Cycles’, *Macroeconomic Dynamics* **2**(01), 22–48.
- Camerer, C. (2003), *Behavioral Game Theory*, Princeton University Press.
- Duffy, J. & Fisher, E. O. (2005), ‘Sunspots in the Laboratory’, *American Economic Review* **95**, 510–529.
- Evans, G. & Honkapohja, S. (1995), Adaptive learning and expectational stability: An introduction, in A. K. . M. Salmon, ed., ‘Learning and Rationality in Economics,’ , Oxford: Basil Blackwell., pp. 102–126.
- Evans, G. W. & Honkapohja, S. (2001), *Learning and Expectations in Macroeconomics*, Princeton University Press.
- Gale, D. (1973), ‘Pure exchange equilibrium of dynamic economic models’, *Journal of Economic Theory* **6**, 12–36.
- Gardini, L., Hommes, C., Tramontana, F. & de Vilder, R. (2009), ‘Forward and backward dynamics in implicitly defined overlapping generations models’, *Journal of Economic Behavior & Organization* **71**(2), 110–129.
- Grandmont, J. (1985), ‘On Endogenous Competitive Business Cycles’, *Econometrica* **53**(5), 995–1045.
- Guesnerie, R. & Woodford, W. (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, pp. 110–133.
- Heemeijer, P., Hommes, C., Sonnemans, J. & Tuinstra, J. (2009), ‘Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation’, *Journal of Economic Dynamics & Control* **33**, 1052–1072.

- Heemeijer, P., Hommes, C., Sonnemans, J. & Tuinstra, J. (2012), ‘An Experimental Study on Expectations and Learning in Overlapping Generations Models’, *Studies in Nonlinear Dynamics & Econometrics* **16**(4), 1–49.
- Hommes, C. (2011), ‘The heterogeneous expectations hypothesis: Some evidence from the lab’, *Journal of Economic Dynamics & Control* **35**, 1–24.
- Hommes, C. & Sorger, G. (1998), ‘Consistent expectations equilibria’, *Macroeconomic Dynamics* **2**, 287–321.
- Hommes, C., Sorger, G. & Wagener, F. (2013), Consistency of Linear Forecasts in a Nonlinear Stochastic Economy, in G. B. et al., ed., ‘Global Analysis of Dynamic Models in Economics and Finance’, Springer-Verlag Berlin Heidelberg.
- Hommes, C. & Zhu, M. (2014), ‘Behavioral learning equilibria’, *Journal of Economic Theory* **150**(C), 778–814.
- Lucas, R. E. (1986), ‘Adaptive Behavior and Economic Theory’, *The Journal of Business* **59**(4), S401–S426.
- Marimon, R., Spear, S. E. & Sunder, S. (1993), ‘Expectationally Driven Market Volatility: An Experimental Study’, *Journal of Economic Theory* **61**(1), 74–103.
- Marimon, R. & Sunder, S. (1993), ‘Indeterminacy of Equilibria in a Hyperinflationary World: Experimental Evidence’, *Econometrica* **61**(5), 1073–107.
- Marimon, R. & Sunder, S. (1994), ‘Expectations and Learning under Alternative Monetary Regimes: An Experimental Approach’, *Economic Theory* **4**(1), 131–62.
- Spiliopoulos, L. (2012), ‘Pattern recognition and subjective belief learning in a repeated constant-sum game’, *Games and Economic Behavior* **75**(2), 921–935.
- Van Huyck, J. B., Cook, J. P. & Battalio, R. C. (1994), ‘Selection Dynamics, Asymptotic Stability, and Adaptive Behavior’, *Journal of Political Economy* **102**(5), 975–1005.
- Woodford, M. (1990), ‘Learning to Believe in Sunspots’, *Econometrica* **58**(2), 277–307.
- Zeileis, A. (2004), ‘Econometric Computing with HC and HAC Covariance Matrix Estimators’, *Journal of Statistical Software* **11**(10), 1–17.

A Summary of the experimental treatments

ρ_2	3	5	8	12	13.5
s^*	0.562	0.5387	0.5246	0.5166	0.5148
P^*	17.78	8.91	0.72	17.1	16.208
$\{s_1^*, s_2^*\}$	NA	{1.1823, 0.1203}	{1.4614, 0.0094}	{0.0001, 1.4981}	{0.0002, 1.497}
$\{P_1^*, P_2^*\}$	NA	{4.06, 39.9}	{0.26, 40.39}	{11.0835, 60.71}	{10.288, 66.5}

LtFE					
ρ_2	3	5	8	12	13.5
exchange rate	0.00027 (1300 pts = 0.35E)				
T	50	100	100	100	100
nb. of observations	4	4	4	4	4

LtOE					
ρ_2	3	5	8	12	13.5
nb. of observations	4	9	15	NA	NA
α	NA	4	25	NA	
P_0	50	50	15	NA	
exchange rate	0.00045	0.00039		NA	
payoff on the steady state	975 pts 0.44E	723 pts 0.28E	639 pts 0.26E	NA	
payoff on 2-cycle	NA	$\frac{40+1310}{2} = 675\text{pts}$ = 0.26E	$\frac{15+1301}{2} = 658\text{pts}$ = 0.27E	NA	
C	0	1	15.25	NA	
K	500	350.5	71.92	NA	

Table 8: Summary of the experimental treatments and designs

N.B.: The equilibrium price values for $\rho_2 = 12$ and 13.5 in the LtFE corresponds to the transformed values given by Equations (12). The exchange rate between experimental currency and euros is higher in the LtOE than in the LtFE in order to ensure reasonable earnings for subjects with respect to the relative length of the experimental sessions.

B Individual and aggregate dynamics in the experimental economies

B.1 Experimental economies – LtFE

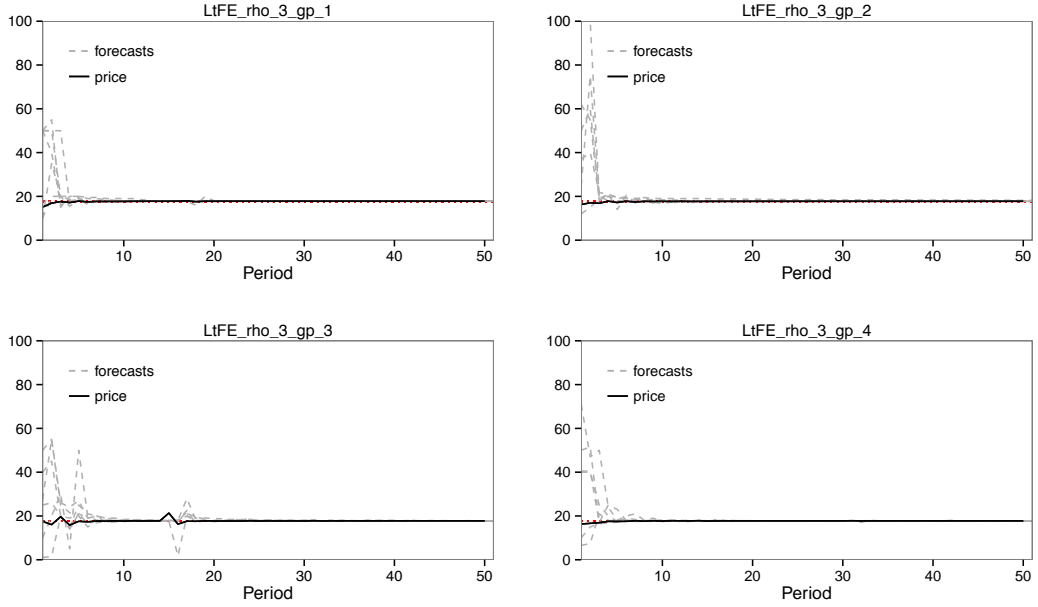


Figure 11: $\rho_2 = 3$ (steady state)

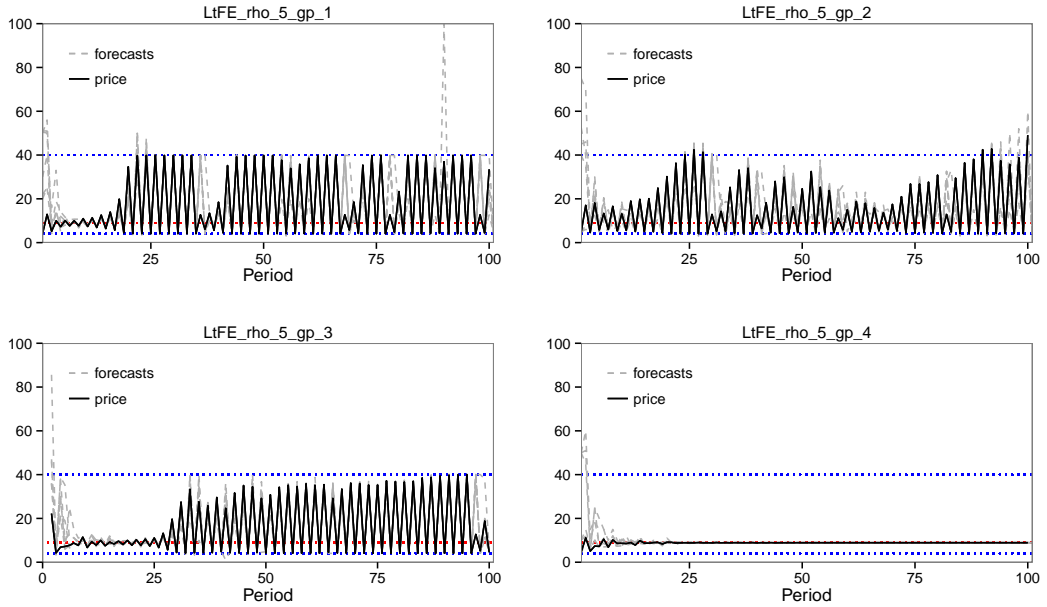


Figure 12: $\rho_2 = 5$ (2-cycle)

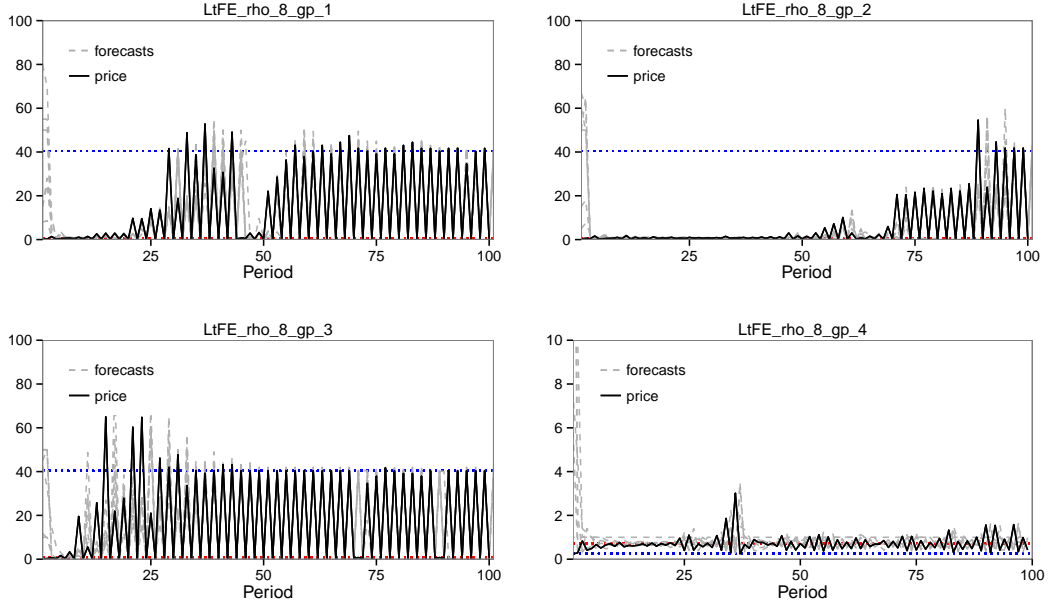


Figure 13: $\rho_2 = 8$ (noisy 32 cycle)

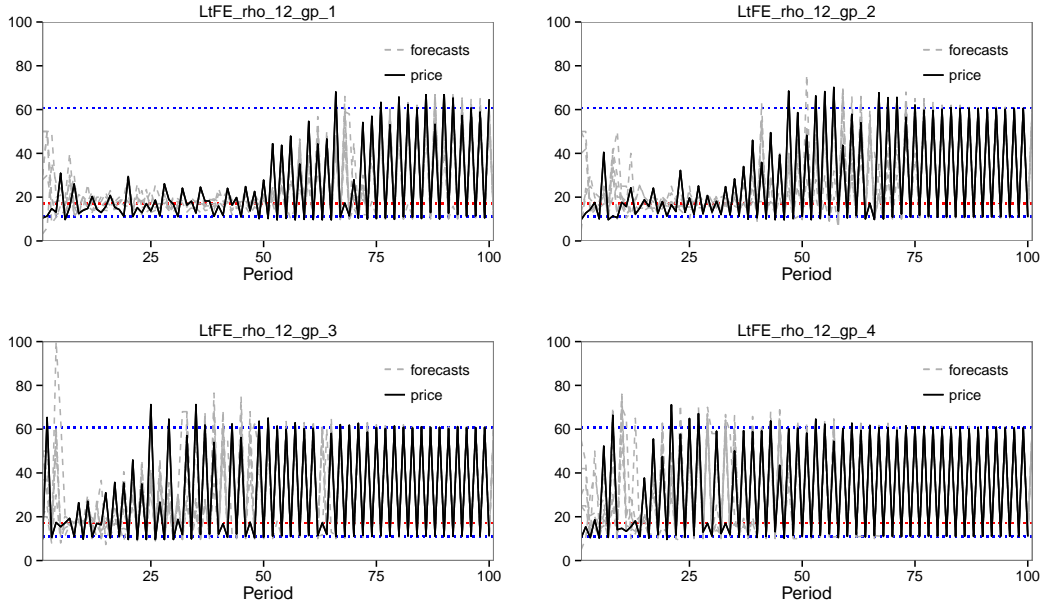


Figure 14: $\rho_2 = 12$ (chaotic dynamics)

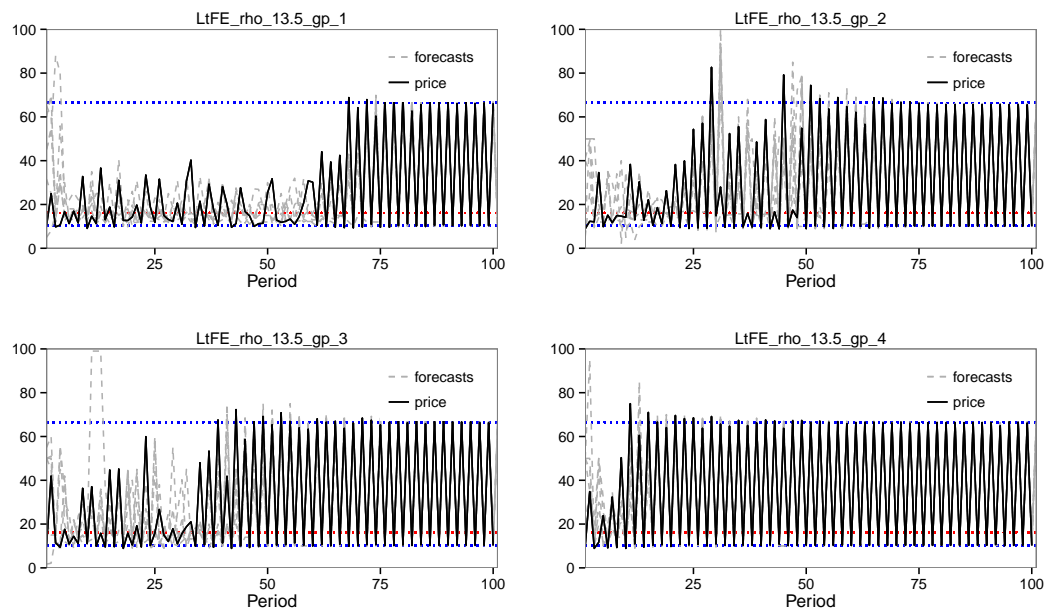


Figure 15: $\rho_2 = 13.5$ (3-cycle)

B.2 Experimental economies – LtOE

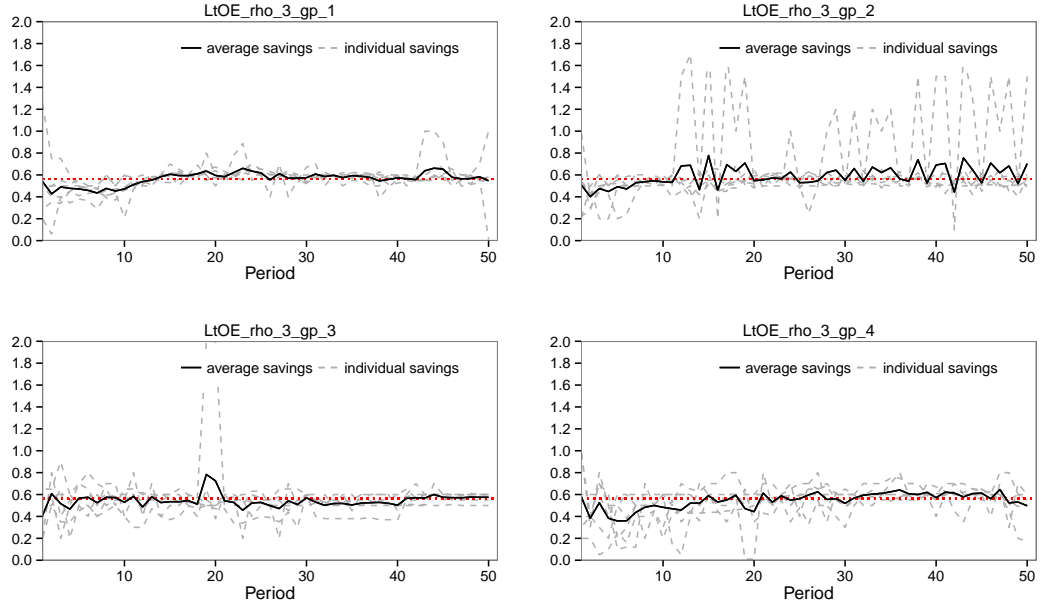


Figure 16: $\rho_2 = 3$ (steady state)

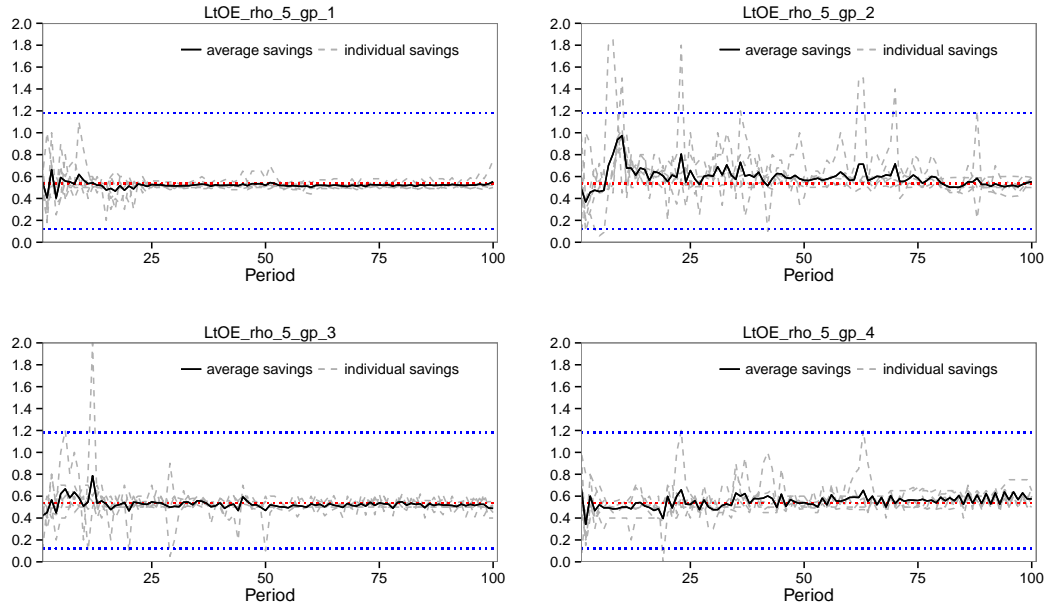


Figure 17: $\rho_2 = 5$ (2-cycle)

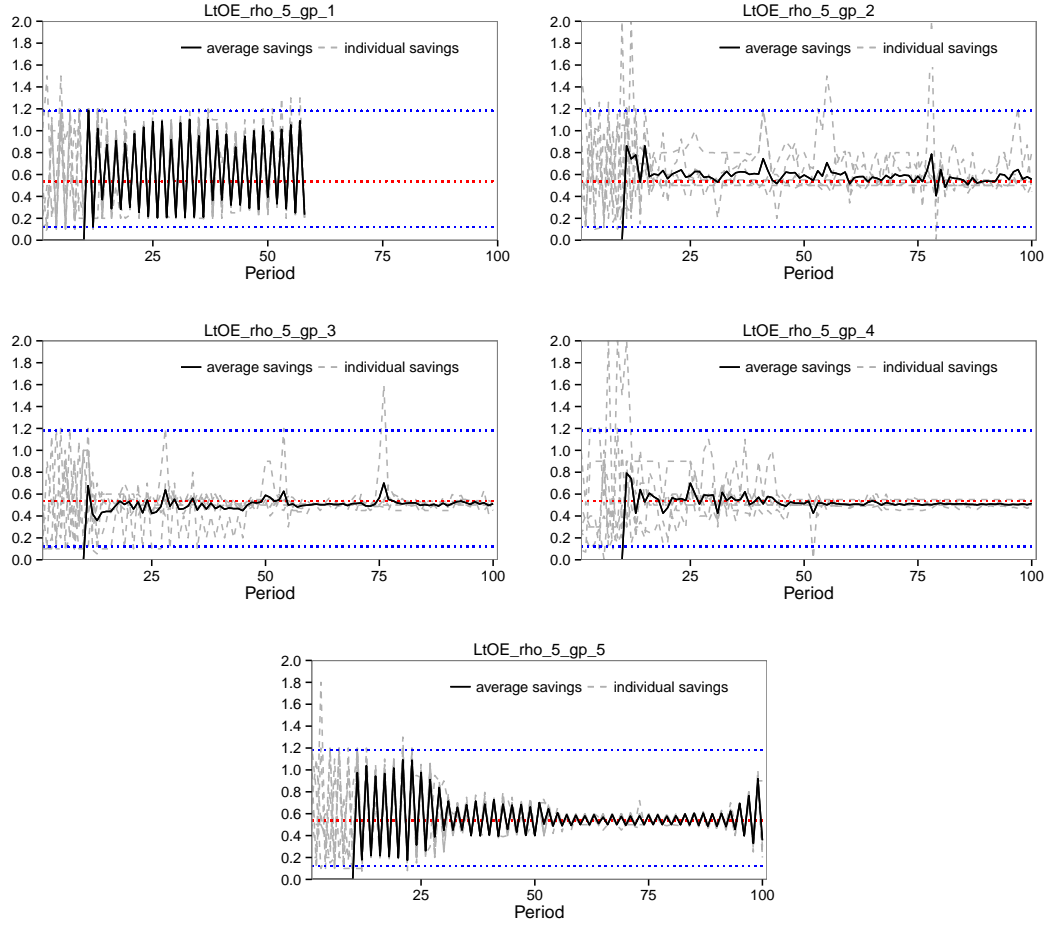


Figure 18: $\rho_2 = 5$ (2-cycle) **with training**. N.B.: a computer crashed at period 58 in group 1.

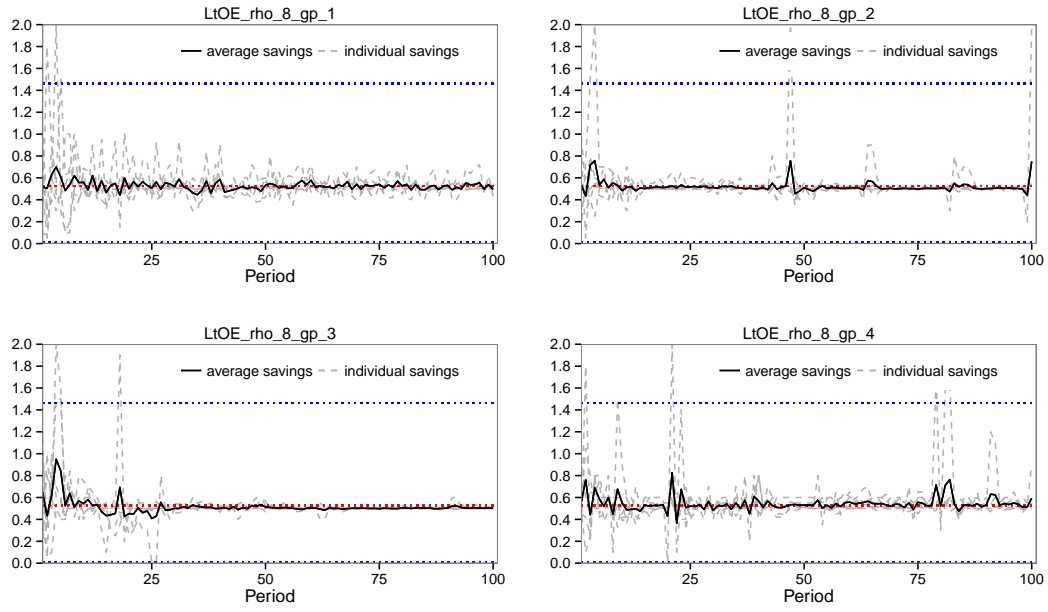


Figure 19: $\rho_2 = 8$ (32-cycle)

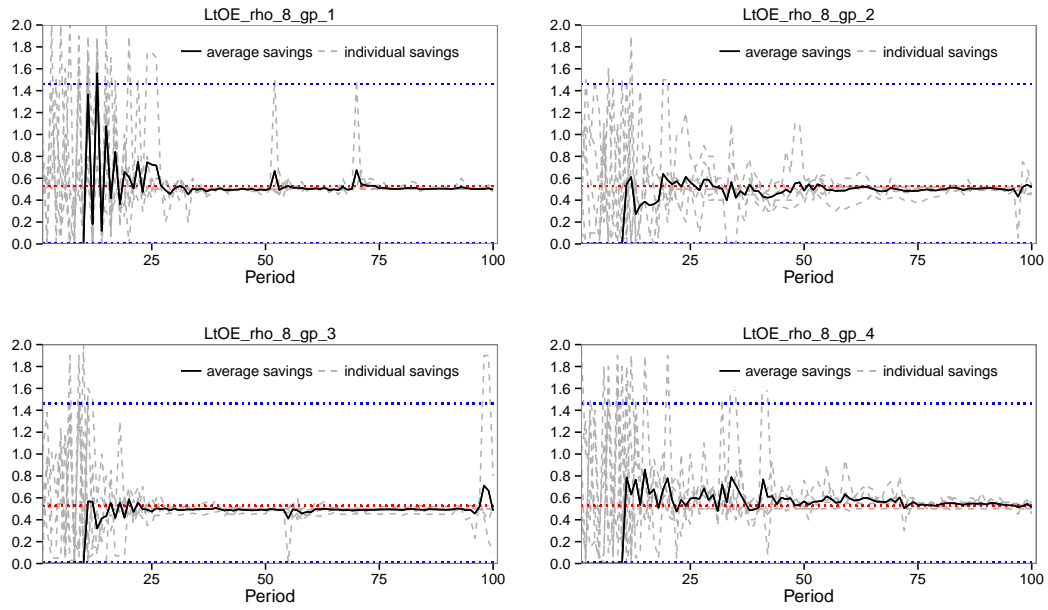


Figure 20: $\rho_2 = 8$ (32-cycle) **with training**

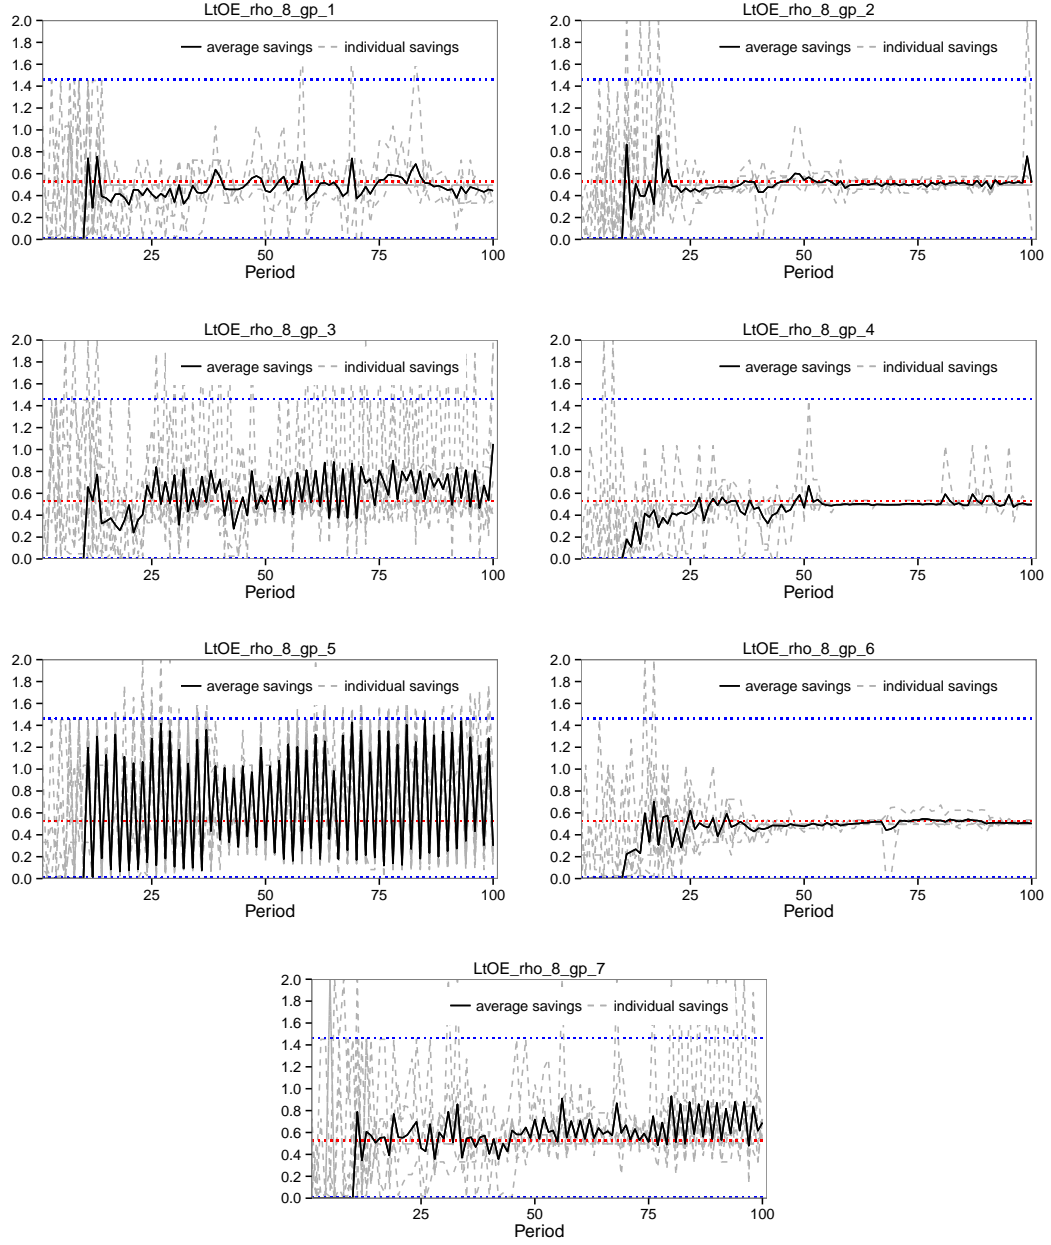


Figure 21: $\rho_2 = 8$ (32-cycle) **with training and non-linear transformation of savings** $\in [1, 100]$

C Additional results

C.1 LtFE

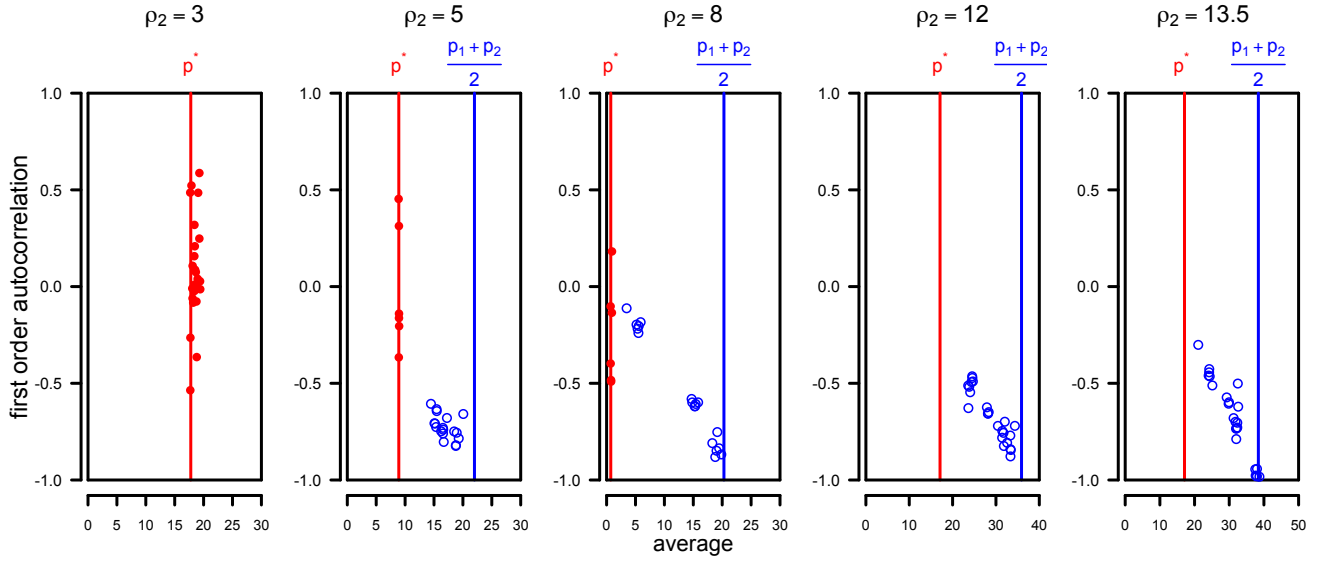


Figure 22: Sample average and first order sample autocorrelation of individual forecasts, LtFE, discarding the first 10 periods. In red: convergence to the steady state, in blue: convergence to the 2-cycle.

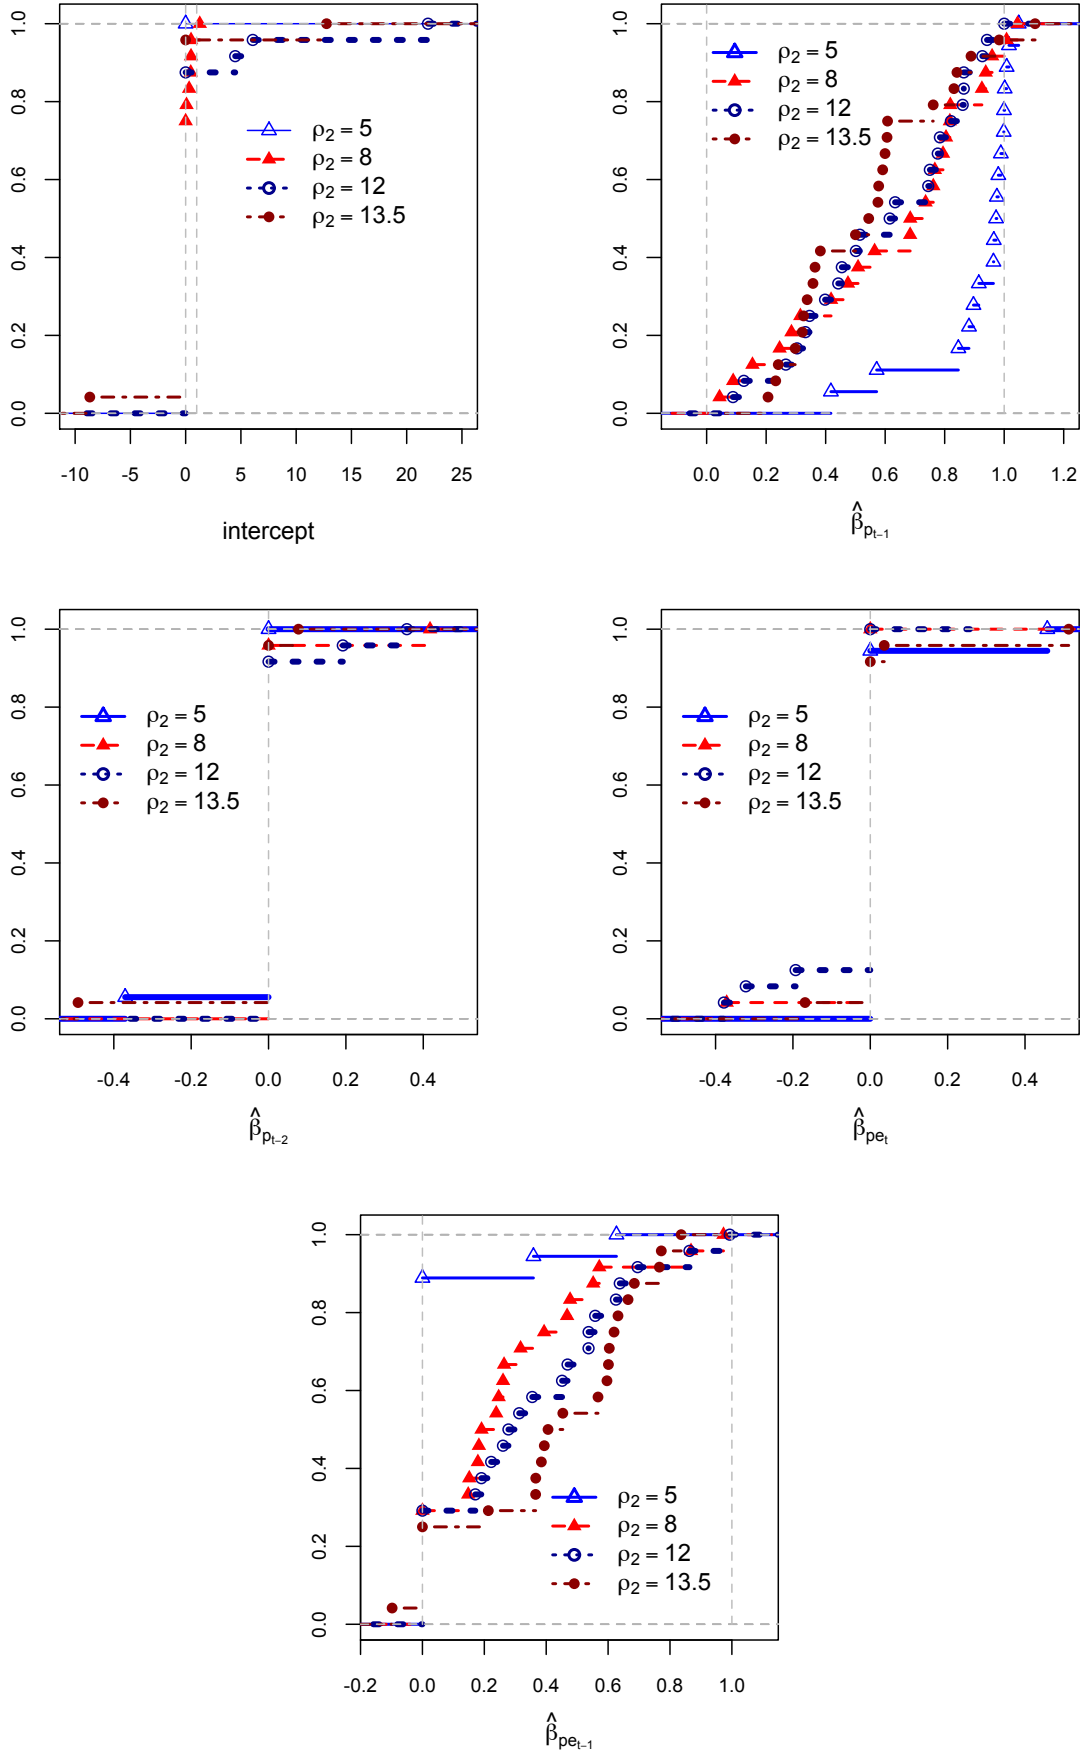


Figure 23: Cumulative distribution of the estimated coefficients in Equation (14) for the economies that converge to a two-cycle (i.e.⁵³all groups with $\rho_2 = 12$ and $\rho_2 = 13.5$ and Groups 1-2-3 with $\rho_2 = 5$ and $\rho_2 = 8$, 84 observations).

C.2 LtOE




Figure 24: Average and first order autocorrelation of individual savings decisions, LtOE, discarding the first 10 periods. In red: convergence to the steady state, in blue: convergence to the 2-cycle.

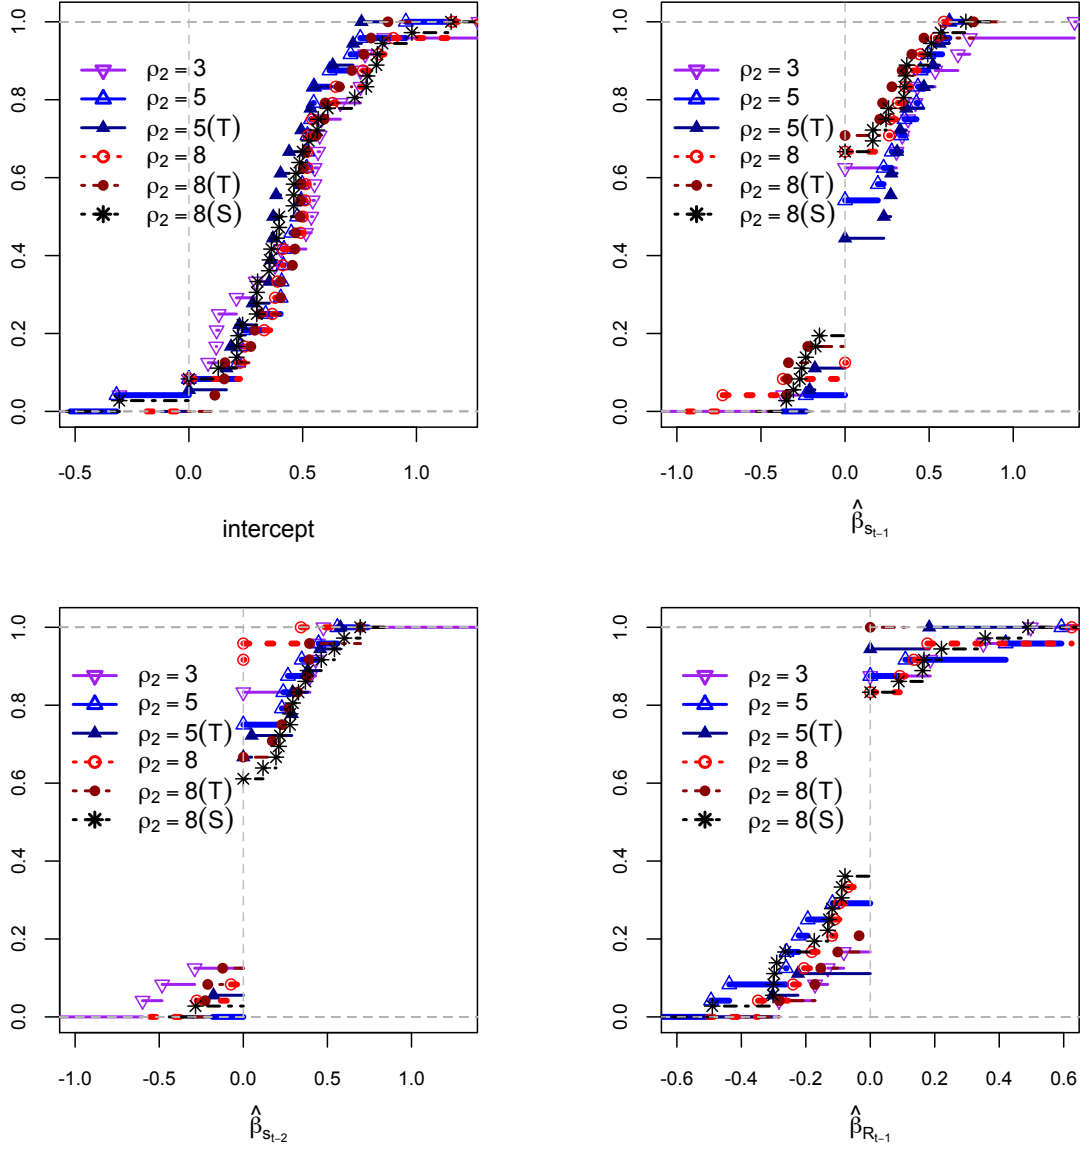


Figure 25: Cumulative distribution of the estimated coefficients in Equation (16) for groups that converge to steady state, 150 observations.

D Instructions of the LtFE for $\rho_2 > 3$ [$\rho_2 = 3$]

Welcome! The experiment is anonymous, the data from your choices will only be linked to your station ID, not to your name. If you follow these instructions carefully, you can earn a considerable amount of money. You will be paid privately in cash at the end of the experiment, after all participants have finished the experiment. Before the payment, you will be asked to fill out a short questionnaire. On your desk you will find a calculator and scratch paper, which you can use during the experiment. Before starting the experiment, you have to answer the questions at the end of the instructions 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. If you have any questions, please raise your hand, and we will come to you and answer your question privately.**

Information about the experimental economy

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 a consumption good. We will refer to this consumption good as chips. 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 receive an income of 200 chips. The old generation does not work any more, and therefore only receives a smaller income of 50 chips. These incomes are fixed and identical across all individuals from the same generation.

Young individuals can choose to *consume* only part of their 200 chips, and to *save* the rest to consume more in the next period, when they will be old. In each period, a young individual then consumes:

$$\text{consumption of chips when young} = 200 - \text{number of chips saved}$$

To carry the saved chips to the next period, the young individual converts these chips into money, by selling them to the old individuals at the current price in the chips market. The savings of a young individual in money then equals:

$$\text{savings in money} = \text{number of chips saved} \times \text{current price of the chips}$$

Once old, in the next period, an individual spends all his money to buy as many chips as his savings can buy from the new young individuals, at the prevailing price for chips. The amount of consumption of chips of an old individual then equals:

$$\text{consumption of chips when old} = 50 + \frac{\text{savings in money}}{\text{price of the chips when old}}$$

The price of chips is always determined in such a way that the chips saved by the young individuals can be exactly bought by the monetary savings 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 six 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 number of chips young individuals will be selling to the old ones *in the current period*, and therefore **the price of the chips in the current period: the higher your price forecast for the next period, the more chips the young individuals save and the more chips to buy in the market in the current period, and the lower the realized price of chips in the current period**. 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**. As for old people, they do not need your forecasts, as they just consume the number of chips their savings can buy.

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 [50] periods or generations. **At the beginning of each period, you have to submit a prediction 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 its realized value (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. For this period only, you are asked to submit two forecasts: your price forecast for the current period (period 1) and for the next period (period 2). Once all participants have submitted their two price forecasts, all young individuals decide 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 can now observe your forecast error for period 1. You are then entering period 2.

From period 2 to the end of the experiment (period 100 [50]), **you have to submit a single forecast of the price in the next period**. At 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 save 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 for period 2. You are then entering period 3. This sequence of events takes place in each of the 100 [50] periods of the experiment.

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

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

All these elements can be relevant to make 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 bottom left part of the screen. 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. At the bottom of the screen there is a status bar telling you when you can enter your prediction and when you have to wait for other participants.

Information about your payoff

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

$$\text{Your earnings} = \max \left[1300 - \frac{1300}{49} (\text{your forecast error})^2, 0 \right]$$

There is a payoff table with the instructions. It 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 $(6 - 5.7)^2 = 0.3^2 = 0.09$, and your payoff is $\max(1300 - \frac{1300}{49} \times 0.09 = 1298, 0) = 1298$ points. If your prediction of the price was 32 and the realized price is 42, your squared error is $(42 - 32)^2 = 10^2 = 100$, and your payoff is $\max(1300 - \frac{1300}{49} \times 100 = -1353, 0) = 0$, and you do not earn any point.

The sum of your prediction scores over the different periods is shown in the bottom right of the screen. At the end of the experiment, your cumulative payoff over all 100 [50] periods is computed, and converted into euro. **For each 1300 points you make, you earn 0.35 euros. This will be the only payment from this experiment, you will not receive a show-up fee on top of it.**

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

Questionnaire

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

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

...
3. Suppose that in a period, your prediction for the market price was 40, and the market price turns out to be 45.5, how many points do you earn in this period?

...
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?

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

...
6. Suppose all forecasters like you are predicting at the beginning of period 12 a "high" price for period 13, would you say that:
 - (a) The price in period 13 is likely to be high;
 - (b) The price in period 13 is likely to be low;
 - (c) The price in period 12 is likely to be high;
 - (d) The price in period 12 is likely to be low;
 - (e) Forecasts of the price for period 13 do not influence the price in period 13;
 - (f) Forecasts of the price for period 13 do not influence the price in period 12.

N.B.: **multiple answers are possible.**

Pay-off table for the price forecasting task

Your payoff : $\max \left[1300 - \frac{1300}{49} (\text{your forecast error})^2, 0 \right]$

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	error ≥ 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		

E Instructions of the LtOE for $\rho_2 = 3$ [$\rho_2 = 5$] $\{\rho_2 = 8\}$ /tr = S/

General information about the experiment

Welcome! The experiment is anonymous, the data from your choices will only be linked to your station ID, not to your name. If you follow these instructions carefully, you can earn a considerable amount of money. You will be paid privately in cash at the end of the experiment, after all participants have finished the experiment. Before the payment, you will be asked to fill out a short questionnaire. On your desk you will find a calculator that you can use during the experiment. Before starting the experiment, you have to answer the questions at the end of the instructions 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. If you have any questions, please raise your hand, and we will come to you and answer your question privately.

Information about the experimental economy

You participate in a market for a consumption good. We will refer to this consumption good as *chips*. In every period, two generations of individuals – the young and the old – trade chips. 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 receive an income of **200 /100/ chips**. The old generation does not work any more, and therefore only receives a smaller income of 50 /80/ chips. These incomes are fixed and identical across all individuals from the same generation.

Young individuals can choose to *consume* only part of their 200 /100/ chips, and to *save* the rest to consume more than their 50 /80/ chips in the next period, when they will be old. In each period, a young individual then consumes:

$$\text{consumption of chips when young} = 200/100/ - \text{quantity of chips saved}$$

You work for a Professional Saving Advisor Bureau, and you have to decide in each period the quantity of chips a young individual will save. In each period, there are six young individuals, each of them follows the savings decision of a professional advisor. Each advisor is played by a participant like you.

To carry 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:

$$\text{savings in money} = \text{number of chips saved} \times \text{current price of the chips}$$

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

and the more chips the old individuals can purchase back 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:

$$\text{consumption of chips when old} = 50/80/ + \frac{\text{savings in money}}{\text{price of the chips when old}}$$

Your savings decision influences what the individual consumes *both when young* in the current period, and *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:

$$\text{return on savings} = \frac{\text{current price (when young)}}{\text{future price (when old)}}$$

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 decision.** However, **you should make a forecast of the return on savings of the next period to guide your savings decision in the current period.**

Information about your task as an advisor

The savings advisor bureau exists for 50 $\{ \{ /100 / \} \}$ periods or generations. Each individual lives for two periods, consumes and saves when young, and consumes when old. At the beginning of each period, you have to submit a savings decision for a young individual. **Your payoff depends on the consumption of chips of this individual both when young and 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 will have become old.**

The experiment starts at period 1. From period 1 to the end of the experiment (period 50 $\{ \{ /100 / \} \}$), you have to make a savings advice. Once all participants have entered their savings decision in period 1, all young people consume and save chips, all old individuals trade the money they are initially endowed with against the saved chips of the 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 advice for period 2, young individuals consume 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 both in period 1 (when young) and 2 (when old), and therefore the corresponding payoff of your savings decision made in period 1. You are then entering period 3. This sequence of events takes place in each of the 50 $\{ \{ /100 / \} \}$ periods of the experiment.

The computer interface is mainly self-explanatory. When making your savings decision at any period, the following information will be displayed in the table (right panel of the computer screen):

- 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 6 advisers from the beginning of the experiment (period 1) up to the previous period;
- Your savings decisions from the beginning of the experiment (period 1) up to the previous period;
- The corresponding consumption of chips when young 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.

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

All these elements can be relevant to make your savings decision but it is up to you to determine how to use this information. In each period, **the return on savings you need to forecast for the next period and the savings decision you need to make for the current period will be displayed on your screen with question marks (?)** to help you.

When submitting your savings decision, use a decimal point if necessary (not a comma). For example, if you want to save 15.05 chips, type 15.05. At the bottom of the screen there is a status bar telling you when you can enter your savings decision and when you have to wait for other participants.

Information about your payoff

In each period, **your payoff depends on the quality of your savings decisions. The higher utility the individual you are advising gets from his consumption when young *and* when old**, the higher the quality of your savings decisions, and **the higher your payoff**. You do not need to calculate his utility, and hence your payoff yourself. **There is a payoff table on your table**. 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 should use this payoff table to choose your savings decision *in the current period* (horizontal axis) according to your forecast of the return on savings *in the next period* (vertical axis)**. 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 130 or 140, but you may submit 141.2. 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.

Example 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 $\frac{10}{20} = 0.5$, this person consumes $200 - 90 = 110/100 - 90 = 10/$ when young, and $50 + 0.5 \times 90 = 95/80 + 0.5 \times 90 = 125/$ when old, and your payoff is 772[422]{356}/329/ 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 $\frac{10}{5} = 2$ and this person consumes $50 + 2 \times 90 = 230/80 + 2 \times 90 = 260/$ when old, and your payoff is 1002[630]{475}/394/ points.

The sum of your payoff from your savings advices over the different periods is shown in the bottom right of the screen. At the end of the experiment, your cumulative payoff over all 50[{/100/}] periods is computed, and converted into euro. **For each 1000 points you make, you earn 0.5 euros. This will be the only payment from this experiment, you will not receive a show-up fee on top of it.**

You now have to fulfil the questionnaire below on the last page of these instructions. We will make sure that every subject has filled out the questionnaire with the correct answers for each of the seven questions before starting the experiment.

If you have any questions, please ask them now!

Questionnaire

1. If you enter period 6, for which period do you need to forecast the return on savings to make your savings decision?

...

2. If you make a savings decision at the beginning of period 9, in which period will you observe your corresponding payoff?

...

3. If you advise to save 150 chips, how many chips will the individual consume when young?

...

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?

...

5. Suppose you forecast that the return on savings will be 9.5, how many chips should you advise to save? Use your payoff table!

...

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?

...

7. Suppose you have decided for a young individual to save 100 chips in a given period.

- (a) The young individual will consume $100 + 50 = 150$ chips when old.
- (b) You do not know yet how many chips the individual will consume when old.
- (c) The consumption of the individual when old will depend only on the price of the chips in the next period.
- (d) The consumption of the individual when old will depend on both the price of the chips in the current and in the next period, and his savings when young.
- (e) You know the current price of the chips when making a saving decision.

N.B.: **multiple answers are possible.**

Pay-off table $\rho_2 = 3$

Your savings decision

	1	10	20	30	40	50	60	70	80	90	100	110	120	130	140	150	160	170	180	190	200
0.01	411	382	350	316	281	244	207	168	127	84	39	0	0	0	0	0	0	0	0	0	0
0.05	413	398	380	361	340	318	293	267	238	207	174	137	97	54	5	0	0	0	0	0	0
0.1	415	417	417	414	408	398	386	371	352	331	306	277	244	207	164	115	59	0	0	0	0
0.2	419	454	484	507	522	530	533	530	522	508	490	466	438	404	364	316	261	193	109	0	0
0.3	423	488	544	586	615	633	643	644	639	627	609	586	556	521	479	430	372	303	216	98	0
0.4	426	521	598	653	691	715	727	729	724	711	691	666	634	596	552	500	440	368	279	159	0
0.5	430	552	647	712	755	780	793	794	787	772	750	722	688	648	601	547	485	411	320	197	0
0.6	434	581	691	763	808	834	845	845	835	818	793	763	726	684	635	580	515	440	347	223	0
0.7	438	609	731	808	854	879	888	885	873	853	826	794	755	711	660	603	537	460	367	242	0
0.8	442	635	768	847	893	916	923	918	903	881	852	817	777	731	679	621	554	475	381	255	0
0.9	446	660	801	882	927	948	952	944	927	903	872	836	794	747	694	634	566	487	392	265	0
1	450	684	831	913	956	975	977	967	948	921	889	851	808	759	705	645	576	496	400	273	0
1.1	453	707	859	941	982	998	997	985	964	936	902	863	819	770	715	653	584	503	407	279	0
1.2	457	728	885	966	1005	1018	1015	1001	978	949	913	873	828	778	722	660	590	509	412	284	0
1.3	461	749	909	988	1025	1036	1031	1014	990	959	923	882	836	785	728	665	595	514	416	288	0
1.4	464	768	931	1008	1042	1051	1044	1026	1000	968	931	889	842	790	733	670	599	518	420	291	0
1.5	468	787	951	1027	1058	1065	1056	1036	1009	976	938	895	847	795	738	674	603	521	423	294	0
1.6	472	804	970	1044	1073	1077	1066	1045	1017	982	943	900	852	799	741	677	606	524	425	296	0
1.7	475	821	987	1059	1085	1088	1075	1053	1023	988	948	904	856	803	744	680	608	526	427	298	0
1.8	479	838	1004	1073	1097	1097	1083	1059	1029	993	953	908	859	806	747	683	611	528	429	300	0
1.9	483	853	1019	1085	1107	1106	1090	1066	1034	998	957	911	862	808	750	685	613	530	431	301	0
2	486	868	1033	1097	1117	1114	1097	1071	1039	1002	960	914	865	811	752	687	614	531	432	303	0
3	521	988	1135	1176	1178	1162	1136	1103	1066	1024	980	931	880	824	763	697	624	540	440	310	0
4	553	1070	1194	1217	1208	1185	1153	1117	1077	1034	988	938	886	829	768	701	627	543	443	312	0
5	584	1128	1231	1241	1225	1197	1163	1125	1083	1039	992	942	889	832	770	703	629	545	444	314	0
6	613	1172	1255	1257	1235	1204	1168	1129	1087	1042	994	944	890	833	771	704	630	546	445	314	0
7	641	1205	1272	1267	1242	1209	1172	1132	1089	1043	996	945	891	834	772	705	631	546	446	315	0
8	668	1230	1285	1274	1247	1212	1174	1133	1090	1045	997	946	892	835	773	706	631	546	446	315	0
9	692	1251	1294	1279	1250	1215	1176	1135	1091	1045	997	946	892	835	773	706	631	547	446	315	0
10	716	1267	1302	1283	1253	1216	1177	1136	1092	1046	998	947	893	835	773	706	632	547	446	316	0

Your forecast of the return on savings
(Number of chips bought when old with one chip saved when young)

Your savings decision

	1	5	10	20	30	40	50	60	70	80	90	100	110	120	130	140	150	160	170	180	190	200
0.05	5	5	5	5	5	5	5	5	5	6	7	10	7	6	5	5	5	5	5	5	5	5
0.075	5	5	5	5	5	5	5	6	6	6	7	7	12	12	7	6	6	6	5	5	5	5
0.1	5	5	5	5	5	6	7	8	10	11	13	14	15	20	15	14	12	11	9	7	6	5
0.2	5	5	5	7	14	26	42	59	76	90	100	106	108	105	99	90	78	64	49	34	20	6
0.3	5	5	6	17	44	83	126	165	196	217	228	229	222	208	189	166	140	114	86	60	35	8
0.4	5	5	8	37	98	170	236	288	321	338	339	328	308	281	249	215	179	143	108	74	43	10
0.5	5	5	13	70	168	268	349	401	428	433	422	398	366	328	287	244	201	159	119	82	47	11
0.6	5	6	20	113	248	368	451	497	512	504	480	445	403	357	309	261	214	169	126	86	50	11
0.7	5	7	31	164	331	461	540	574	576	556	521	477	428	376	324	272	222	174	130	89	51	11
0.8	5	9	44	220	412	544	614	636	625	594	550	499	445	389	333	279	227	178	132	90	52	12
0.9	5	11	62	279	489	617	675	684	662	622	571	515	456	397	339	283	230	180	134	91	53	12
1	5	14	82	340	559	680	725	721	690	642	586	526	464	403	344	286	232	182	135	92	53	12
1.1	5	17	105	400	624	734	766	751	711	658	597	534	470	407	346	289	234	183	135	92	53	12
1.2	5	22	130	458	682	780	799	775	728	669	605	540	474	410	349	290	235	183	136	93	53	12
1.3	5	27	157	514	734	819	826	794	741	678	612	544	477	412	350	291	236	184	136	93	53	12
1.4	5	33	186	567	780	852	849	809	751	685	616	547	480	414	351	292	236	184	137	93	53	12
1.5	5	40	217	618	820	880	867	821	759	691	620	550	481	415	352	293	237	185	137	93	54	12
1.6	5	48	248	665	856	903	883	831	766	695	623	552	483	416	353	293	237	185	137	93	54	12
1.7	5	57	280	708	888	924	895	839	771	699	626	554	484	417	353	293	237	185	137	93	54	12
1.8	5	66	313	749	917	941	906	846	775	702	627	555	485	418	354	294	237	185	137	93	54	12
1.9	5	77	346	787	941	957	915	852	779	704	629	556	485	418	354	294	238	185	137	93	54	12
2	5	88	379	822	964	970	923	856	782	706	630	557	486	418	354	294	238	185	137	93	54	12
3	6	230	678	1051	1087	1036	960	878	795	714	636	560	488	420	356	295	238	186	137	93	54	12
4	9	399	897	1153	1131	1056	970	884	798	716	637	561	489	420	356	295	238	186	137	93	54	12
5	15	563	1044	1203	1149	1064	974	886	799	717	637	561	489	420	356	295	238	186	137	93	54	12
6	23	710	1141	1229	1157	1067	976	886	800	717	637	561	489	421	356	295	238	186	137	93	54	12
7	36	834	1206	1243	1162	1069	977	887	800	717	637	561	489	421	356	295	238	186	137	93	54	12
8	51	936	1250	1252	1164	1070	977	887	800	717	637	561	489	421	356	295	238	186	137	93	54	12
9	71	1019	1281	1257	1166	1070	977	887	800	717	637	561	489	421	356	295	238	186	137	93	54	12
10	93	1087	1303	1261	1167	1071	977	887	800	717	637	561	489	421	356	295	238	186	137	93	54	12
20	415	1354	1364	1268	1169	1071	978	887	800	717	637	561	489	421	356	295	238	186	137	93	54	12

Your forecast of the return on savings
(Number of chips bought when old with one chip saved when young)

Pay-off table $\rho_2 = 8$

Your savings decision

	1	5	10	20	30	40	50	60	70	80	90	100	110	120	130	140	150	160	170	180	190	200
0.005	10	10	10	10	10	10	10	10	10	10	10	10	11	15	15	11	10	10	10	10	10	10
0.006	10	10	10	10	10	10	10	10	10	10	10	10	10	11	11	15	15	11	10	10	10	10
0.007	10	10	10	10	10	10	10	10	10	10	10	10	10	10	11	11	15	15	11	11	10	10
0.008	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	11	11	15	15	11	10	10
0.009	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	11	15	15	11	10	10
0.01	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	11	15	15	11	10	10
0.02	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	11	15	15	11	10	10
0.03	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	11	11	11	11	15	11
0.04	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	11	15	10
0.05	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	11	15	15	10
0.06	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	11	15	15	10
0.07	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	11	15	15	10
0.08	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	11	15	15	10
0.09	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	11	11	11	15	15	10
0.1	10	10	10	10	10	10	10	10	10	10	10	10	10	10	11	11	11	12	15	15	11	11
0.2	10	10	10	10	10	10	10	11	14	19	28	39	50	60	67	70	70	66	60	51	39	21
0.3	10	10	10	10	10	11	19	37	68	103	135	158	170	171	165	152	135	116	96	75	55	26
0.4	10	10	10	10	12	26	68	132	194	241	266	271	262	243	218	191	163	135	108	83	59	28
0.5	10	10	10	10	21	76	172	265	328	355	356	338	310	277	241	206	173	141	112	86	60	28
0.75	10	10	10	23	143	334	467	520	519	489	445	397	348	301	257	216	179	145	115	87	61	29
1	10	10	10	94	373	573	638	626	581	525	466	409	355	305	260	218	180	146	115	87	61	29
1.25	10	10	13	239	584	713	713	665	601	536	472	412	357	306	260	218	180	146	115	87	61	29
1.5	10	10	24	415	730	786	747	680	609	539	474	413	358	307	260	218	180	146	115	87	61	29
1.75	10	10	53	579	821	824	762	687	612	541	475	414	358	307	260	218	180	146	115	87	61	29
2	10	10	104	711	876	843	769	690	613	541	475	414	358	307	260	218	180	146	115	87	61	29
3	10	25	461	973	952	866	777	693	614	542	475	414	358	307	260	218	180	146	115	87	61	29
4	10	109	788	1043	965	870	778	693	614	542	475	414	358	307	260	218	180	146	115	87	61	29
5	10	279	978	1064	968	870	778	693	614	542	475	414	358	307	260	218	180	146	115	87	61	29
6	10	485	1077	1072	969	870	778	693	614	542	475	414	358	307	260	218	180	146	115	87	61	29
8	10	828	1155	1076	970	871	778	693	614	542	475	414	358	307	260	218	180	146	115	87	61	29
10	10	1028	1178	1076	970	871	778	693	614	542	475	414	358	307	260	218	180	146	115	87	61	29
25	291	1247	1192	1077	970	871	778	693	614	542	475	414	358	307	260	218	180	146	115	87	61	29
50	1070	1252	1192	1077	970	871	778	693	614	542	475	414	358	307	260	218	180	146	115	87	61	29
75	1250	1252	1192	1077	970	871	778	693	614	542	475	414	358	307	260	218	180	146	115	87	61	29
100	1287	1252	1192	1077	970	871	778	693	614	542	475	414	358	307	260	218	180	146	115	87	61	29
125	1297	1252	1192	1077	970	871	778	693	614	542	475	414	358	307	260	218	180	146	115	87	61	29
150	1300	1252	1192	1077	970	871	778	693	614	542	475	414	358	307	260	218	180	146	115	87	61	29
200	1302	1252	1192	1077	970	871	778	693	614	542	475	414	358	307	260	218	180	146	115	87	61	29

Number of chips bought when old with one chip saved when young)

Your forecast of the return on savings

F Computer interfaces




Figure 26: Subjects' computer interface (LtFE)

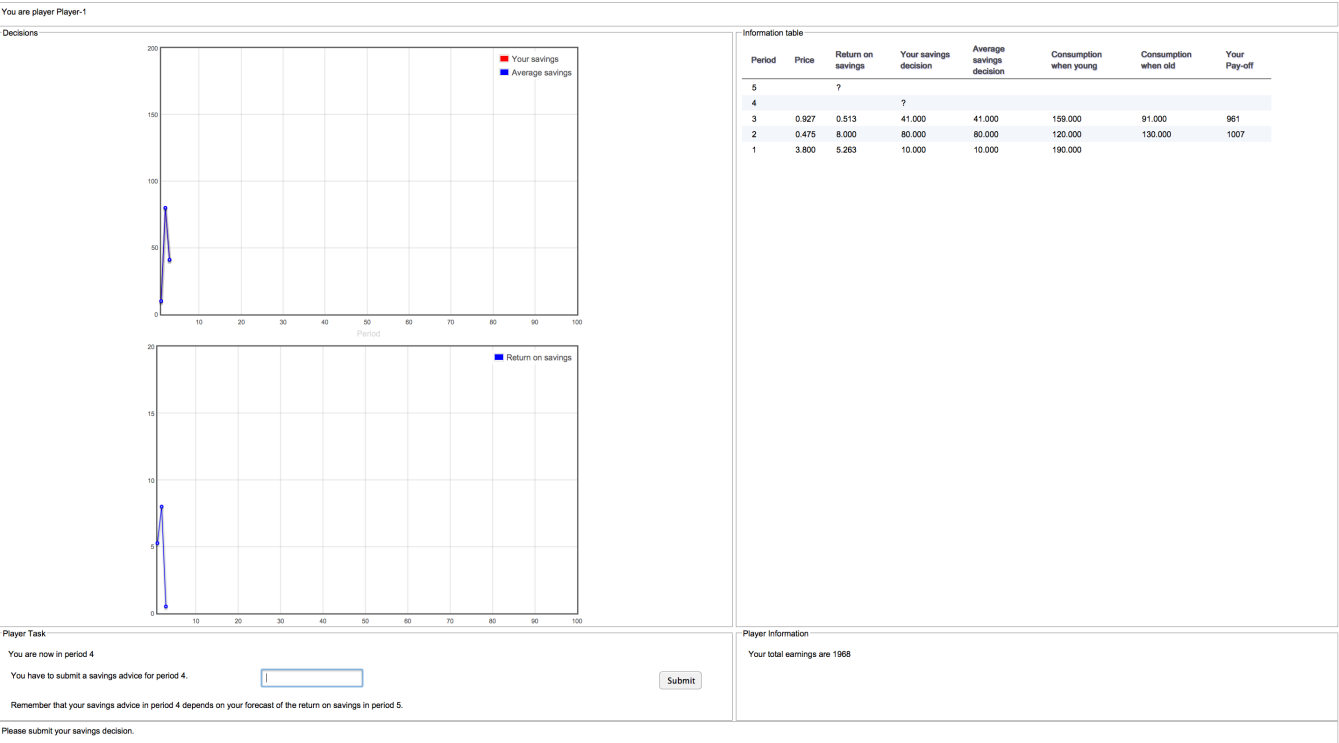


Figure 27: Subjects' computer interface (LtOE)