UvA-DARE (Digital Academic Repository)

Learning to believe in simple equilibria in a complex OLG economy - evidence from the lab

Arifovic, J.; Hommes, C.; Salle, I.

DOI

[10.1016/j.jet.2019.05.001](https://doi.org/10.1016/j.jet.2019.05.001)

Publication date

2019

Document Version

Final published version

Published in

Journal of Economic Theory

License

Article 25fa Dutch Copyright Act

[Link to publication](#)

Citation for published version (APA):

Arifovic, J., Hommes, C., & Salle, I. (2019). Learning to believe in simple equilibria in a complex OLG economy - evidence from the lab. *Journal of Economic Theory*, 183, 106-182. <https://doi.org/10.1016/j.jet.2019.05.001>

General rights

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

Disclaimer/Complaints regulations

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

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Learning to believe in simple equilibria in a complex OLG economy - evidence from the lab [☆]

Jasmina Arifovic ^a, Cars Hommes ^{b,c,d}, Isabelle Salle ^{d,*}

^a *Simon Fraser University, Burnaby, BC, Canada*

^b *CeNDEF, Amsterdam School of Economics, University of Amsterdam, Netherlands*

^c *Tinbergen Institute, Netherlands*

^d *Bank of Canada, Ottawa, ON, Canada*

Received 28 July 2017; final version received 14 November 2018; accepted 10 May 2019

Available online 23 May 2019

Abstract

We set up a laboratory experiment to empirically investigate equilibrium selection in a complex economic environment. We use the overlapping-generation model of Grandmont (1985), which displays multiple perfect-foresight equilibria, including periodic and chaotic dynamics. The equilibrium selection problem is not solved under learning, as each outcome is predicted by at least one existing learning theory. We find that subjects in the lab systematically coordinate on an equilibrium despite the complexity of the environment. Coordination only happens on simple equilibria, in this case the steady state or the period-two cycle, a result which is predicted only if the subjects follow simple learning rules. This suggests that relevant perfect-foresight equilibria should be robust to the use of simple rules.

[☆] **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, the workshop on Behavioral Macroeconomics, Tinbergen Institute Amsterdam, October 26, 2015, the macro seminar at the University of Oregon, Eugene, May 6, 2016, the Canadian Economics Association Conference June 3–5, 2016, Ottawa, the Expectations in Dynamic Macroeconomics Models Conference, Dutch Central Bank, September 6–8, 2016, in particular the discussant Jan Tuinstra, the AEA meetings at Chicago, January 6–8, 2017, in particular the discussant Luisa Lambertini, and the EEA-ESEM meetings at Lisbon, August 22–25, 2017, for helpful comments and stimulating discussion. 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. The views in this paper are those of the authors and do not necessarily reflect those of the Bank of Canada.

* Corresponding author.

E-mail address: ISalle@bank-banque-canada.ca (I. Salle).

© 2019 Elsevier Inc. All rights reserved.

JEL classification: C90; D83; C62

Keywords: Laboratory experiments; Learning; Complex dynamics; Equilibrium selection

1. Introduction

From a theoretical standpoint, the self-fulfilling nature of expectations exposes dynamic general equilibrium models to indeterminacy. When multiple equilibria are possible, the selected one depends not only upon the economic structure, but on the beliefs that agents use to forecast prices (Benhabib and Farmer, 1999). Indeterminacy therefore poses a number of challenges for working with this class of models.

Conceptually, no combination of environmental structure and agent preferences alone can pin down expectations about the future. As a consequence, indeterminacy undermines the concept of rational expectations equilibrium and the predictive power of the model. When multiple equilibria exist, some may be suboptimal. For instance, some equilibria may imply high volatility in real variables that arises from random coordination devices or initial conditions, which is undesirable from the point of view of policy makers aiming to stabilize aggregate fluctuations. This multiplicity further creates both practical issues for comparative static analysis and additional considerations for robust policy design. Being able to predict equilibrium selection in this setting is therefore a critical issue for modeling and policy analysis.

Learning theory has been frequently advocated as a theoretical equilibrium selection device. The main idea is that only rational expectation equilibria that emerge as a long-run outcome of an adaptive learning process should be regarded as relevant (see Evans and Honkapohja (2001) for a comprehensive discussion). A problem with learning, however, is that ‘anything goes’: any equilibrium can be selected if the adaptive rule is suitably designed. 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 and Woodford, 1991; Evans and 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.

A theorist is then left with the crucial yet loosely-defined task of designing the belief formation process of the agents, with little guidance from theory and yet with major consequences for the model’s conclusions. Unsurprisingly, this challenge is accentuated if the model exhibits non-linear or even complex dynamics. In addition, allowing for heterogeneity of beliefs introduces further difficulty, as the process of coordination among agents has to be modeled in turn. Similarly, an empirical economist can pick an interesting equilibrium and fit the dynamics to the data, but what constitutes an interesting or relevant equilibrium remains a non-trivial question.

How agents learn to form and coordinate beliefs, and which equilibria are consequently selected and regarded as plausible, ultimately remain empirical questions. Collecting empirical evidence about agents’ processes of expectation formation and equilibrium selection can undoubtedly provide guidance when designing models of learning. Since this is difficult to do with most available data, economists have taken it to the lab. Lucas (1986) was the first to stress the 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)

One example of this approach comes from the pioneering work of Marimon et al. (1993) and Marimon and Sunder (1993, 1994). In these studies, the authors designed and used laboratory experiments with human subjects to empirically observe the process of equilibrium selection.

The goal of our paper, similarly, is to present laboratory experiments which can test different theories of learning and the resulting equilibrium selection in a *complex* environment. To achieve this, we conduct an experimental study within an OLG economy *à la* Grandmont (1985). There are many reasons for choosing this environment.

Importantly, this is a general-equilibrium environment with a pervasive multiplicity of equilibria, that possesses infinitely many long-run equilibria, including a steady state, cycles of all periods, and even chaotic dynamics. Additionally, since the model is fully deterministic, exogenous shocks play no role in cycle formation; complicated dynamics and multiple equilibria arise under perfect-foresight as soon as there is a strong conflict between the substitution and wealth effects of a change in the return on savings. By simply varying a single parameter, namely the risk aversion parameter in agents' utility functions, the complexity of the model can be tuned in order to produce various treatments with distinct multiplicities of equilibria. This feature considerably simplifies the lab implementation of the model, which undoubtedly represents a strong argument for the choice of the Grandmont OLG environment given the challenge of bringing general-equilibrium economies into lab experiments. What is more, this environment serves as an excellent and classical example of economic complexity, which has not yet been convincingly investigated in a laboratory setting.

A further advantage of the Grandmont model is that it has been extensively studied in the literature since the seminal work by Lucas (1972), and a wide range of learning predictions has been established to guide the construction of hypotheses. These learning theories, however, fail to deliver a clear prediction of equilibrium selection because 'anything goes': all equilibria can be selected under learning provided that agents use a suitable rule. As theory does not resolve equilibrium selection within this complex model, empirical insights gained through the present study may prove informative with respect to agents' coordinating behavior in the (considerably more complex) real world. Finally, the model has been designed so as to incorporate heterogeneous beliefs in a micro-founded general-equilibrium setting.

We use a learning-to-forecast experiment (LtFE), where the only degree of freedom is the belief formation process of the subjects; all other components are deterministic. This design allows the experimentalist to isolate the effects of expectations on the model dynamics (Marimon et al., 1993), and appears as the most natural way to implement a parsimonious, yet complex, lab-based model in which equilibrium selection depends only on self-fulfilling beliefs.

Within this framework, we design several experimental treatments and formulate two main hypotheses. Our first hypothesis relates to the possibility of coordination, amongst a group of individuals holding heterogeneous beliefs, precipitated entirely by repeated market interactions in such a complex environment. As the results will show, we always observe coordination on one of the existing perfect-foresight equilibria of the model, irrespective of the complexity of the

underlying model. This is already a remarkable result: our experiment is the first to document systematic coordination of beliefs in a chaotic environment, and the first in which *spontaneous* coordination on a 2-cycle equilibrium arises, even if the 2-cycle is unstable under learning. This means that, unlike in previous studies, subjects were able to reach a periodic equilibrium in the absence of exogenous fluctuations (see the discussion of related work by Marimon et al. (1993) below). Neither of these outcomes are obvious given the complexity of the underlying model, the heterogeneity in beliefs and the imperfect information that subjects have.

Our second hypothesis is that coordination is more likely to emerge on simple equilibria (such as a steady state) than on more complicated equilibria (i.e., higher-order cycles). This rather intuitive prediction is based on theories of learning and on existing empirical evidence, both of which suggest that subjects tend not to make use of information from more than a few periods prior. Even after considering various treatments with increasingly complicated dynamics, we consistently find aggregate convergence of prices and individual forecasts to the steady state or the 2-cycle, possibly after a long transition. Accordingly, we never observe selection of any higher-order cycles or more complicated equilibrium dynamics.

None of the learning theories predict entirely our experimental results, which underlines the relevance of empirical investigation through lab experiments. A necessary condition for equilibrium selection in the experiment is the weak E-stability criterion that predicts that only if the forecasting rule of the agents is *exactly* consistent with a steady state or a 2-cycle can these outcomes be achieved under adaptive learning. However, this criterion does not eliminate completely the multiplicity issue, and its prediction is not robust to misspecification or overparametrisation of the forecasting rule. Accordingly, we find that the subjects in the experiment only select for simple, weakly E-stable equilibria. This finding shows that, despite the complexity of the environment, subjects adopt simple rules (based on information from last period), but do not use higher order rules.

After the two main hypotheses have been addressed, we assess the robustness of our results to the nature of the experimental task by implementing a learning-to-optimize experiment (LtOE), where subjects explicitly make quantity decisions. This exercise is motivated by previous experimental results which showed that coordination is more challenging in a LtOE than in a LtFE. In the LtOE, we almost systematically find coordination on the monetary steady state. Those sessions also allow us to highlight possible explanations for the absence of coordination on the 2-cycle in the LtOE, namely strategic uncertainty, as subjects may prefer an allocation for which the payoff is constant over one for which it fluctuates, and a higher cognitive load. Finally, two additional sets of experimental sessions show that our results are robust to alternative designs of the experiment.

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 still differ in key ways. In the first of these, Van Huyck et al. (1994) employ an experimental coordination game with two efficient Nash equilibria in order to investigate the problem of equilibrium selection. The myopic best-response dynamics coincide with the chaotic quadratic map, while the interior equilibrium is stable under adaptive learning. In all their experimental sessions, subjects coordinate on the interior solution, in line with the prediction of adaptive learning.

An important difference between this experiment and our own is that our set-up has infinitely many perfect-foresight periodic cycles that arise as equilibrium outcomes of the model. These occur without the need to impose (a priori) any expectation rules, and can be stable under adaptive learning. As such, we aim to find out which of these periodic equilibria, if any, subjects may

coordinate on. By contrast, in Van Huyck et al. (1994), the chaotic dynamics are not an equilibrium of the coordination game, but result from the assumption of myopic best-response behavior. Furthermore, the authors do not address the question of whether or how subjects may coordinate on the best response, which seems especially difficult given the complicated dynamics at work.

The second closely related contribution is the work by Marimon et al. (1993), who were the first researchers to observe a form of coordination on 2-cycle dynamics in a laboratory experiment.¹ They use a design similar to our LtFE, but there are nevertheless several major differences worth noting. 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-population design, in which participants are randomly drawn from the pool to re-enter the market and form the new generation in each period. We use a single-population design, so that the resulting course of events in the experiment is the same as in the learning literature, especially the seminal contribution of Grandmont (1985), and we later show that our results are robust to an alternative design that introduces the overlapping generation friction.

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 generation size are removed.² Hence, these authors do not find evidence of 2-cycles arising spontaneously, a phenomenon which characterizes our own experimental results.

The remainder of this paper is organized as follows. Section 2 introduces the underlying OLG model of the experiment and discusses its properties and learning dynamics. This section is quite technical and may be skipped over, as Section 3 then motivates and details the experimental design within the OLG model and our hypotheses based on learning predictions. Section 4 presents the experimental results, Section 5 provides estimates of individual forecasting rules, and Section 6 presents two additional sets of experimental sessions designed for robustness. Section 7 concludes. Appendices A–K contain details about all experimental sessions and treatments, their analyses and instructions.

2. The model

This section describes the underlying model of the experiment, a deterministic OLG economy *à la* Grandmont (1985). We recall Grandmont's result that the model has infinitely many perfect foresight equilibria and discuss the stability of these under learning.

2.1. The underlying OLG economy

Consider 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. Each agent lives for two periods, meaning that two generations coexist in each period: the young, and the old. Individuals receive an endowment $e_1 > 1$ of the consumption good when young, and $0 < e_2 < 1$

¹ To the best of our knowledge at the time of writing, they remain the only ones to have done so.

² See the five economies in Marimon et al. (1993, Figure 3, p. 89).

when old. These restrictions are sufficient – see Grandmont (1985, Assumption 1.d) – to result in the Samuelson case in the terminology of Gale (1973), in which young individuals seek to 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 $m_t = P_t s_t$. In the following period ($t + 1$), these individuals become the old generation; they now individually purchase goods from the newly-born young generation using all of their available savings. Purchases are made at the market-clearing price, denoted by P_{t+1} , while the aggregate money supply in the economy is held constant at some exogenously-specified $M > 0$.

Expressing the decision problem more formally, a young individual in a given period chooses his current consumption level³ c_t to maximize his expected utility function $U(c_t, c_{t+1}^e)$ over his two-period lifetime, subject to his current (when young) and expected (when old) budget constraints:

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

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

First, we derive the temporary equilibrium map as in Grandmont (1985). The model makes use of a separable utility function:

$$U(c_t, c_{t+1}) = V_1(c_t) + V_2(c_{t+1}) \quad (2)$$

where the functions $V_{1,2}(\cdot)$ are continuous, strictly increasing and concave on $[0, +\infty)$, twice continuously differentiable on $(0, +\infty)$, and 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 in terms of consumption:

$$V'_1(c_t) P_{t+1} = V'_2(c_{t+1}) P_t \quad (3)$$

or in terms of real money balances:

$$V'_1(e_1 - s_t) s_t = s_{t+1} V'_2(s_{t+1} + e_2). \quad (4)$$

Once the optimal savings and consumption decisions have been determined, it is possible to define a perfect-foresight equilibrium sequence of prices (or, equivalently, of real money balances) using the market clearing conditions in the markets for money and goods, respectively:

$$m_t = M, \quad s_t = \frac{M}{P_t} \quad \text{and} \quad s_{t+1} = \frac{P_t}{P_{t+1}} s_t \quad \text{for all } t. \quad (5)$$

At this stage, it is convenient to define the functions $v_1(s) \equiv s V'_1(e_1 - s)$, that maps $[0, e_1]$ onto $[0, +\infty)$; and $v_2(s) \equiv s V'_2(e_2 + s)$, which maps $[0, +\infty)$ onto itself. Since $v_1(\cdot)$ is strictly increasing, we know that $v_1^{-1}(\cdot)$ exists. The dynamics of s under perfect foresight can then be

³ Conversely, the same choice can be expressed as an individual choosing his real money balance s_t .

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}) , given the return on savings R_{t+1} . Equivalently, the dynamics of the model can be expressed in terms of a temporary equilibrium map in prices, given by a continuous map, denoted by G :

$$P_t = G(P_{t+1}^e) = \frac{M}{\chi(M/P_{t+1}^e)}. \quad (6)$$

The maps G and χ are topologically equivalent, since describing the dynamics of the model in terms of prices is equivalent to describing them in terms of savings.

A *perfect-foresight (periodic) equilibrium* is a (periodic) sequence of prices that is a solution of (6) 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 is a sequence (or orbit) of k prices $\{P_1, P_2, \dots, P_k\}$, such that $P_j = G^k(P_j)$ for $j = 1, \dots, k$, where G^k denotes the k^{th} iterate of the map G .⁴

2.3. Existence of infinitely many perfect-foresight equilibria

There are at most two steady states in this 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 are 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 ambiguous effect on consumption when young. These include periodic equilibria of any period, and infinitely many chaotic equilibria.⁵

As is often done in the related OLG literature, we make use of CRRA utility functions for our laboratory experiments:

$$V(c_1) = \frac{c_1^{1-\rho_1}}{1-\rho_1}, \quad V(c_2) = \frac{c_2^{1-\rho_2}}{1-\rho_2}, \quad (7)$$

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 further, the offer curve becomes bumpier, and the long-run price dynamics increase in complexity owing to 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.

⁴ Here we take k to be the smallest integer greater than one which satisfies these conditions.

⁵ Azariadis and Guesnerie (1986) show that if the model has an equilibrium cycle, it also has many sunspot equilibria on which expectations may coordinate. As we consider a deterministic environment, we do not address the existence of or coordination on sunspot equilibria.

2.4. Stability of perfect-foresight equilibria under learning

Which of these infinitely many equilibria are stable under learning? Grandmont (1985) distinguishes between *forward perfect-foresight dynamics* (when the map G in (6) is defined according to $P_{t+1}^e = P_{t+1}$) and *backward perfect-foresight dynamics* (when $P_{t+1}^e = P_{t-1}$ in (6)).⁶ It is a well-known result in this literature that 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, while backward dynamics are effectively fictitious to the extent they imply that time flows backwards.⁷

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 under backward perfect-foresight dynamics is also stable under *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. In this case, the equilibrium cycle is said to be *E-stable* (see Evans and Honkapohja (2001) for a detailed treatment of this literature).

Fig. 1 reproduces the bifurcation diagram of the backward perfect-foresight dynamics (i.e. under naïve expectations) from Grandmont (1985, p. 1030).⁸ 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 that is stable under backward perfect-foresight dynamics for each value of ρ_2 . For such an equilibrium of period k , as long as the forecasting rules of the agents are consistent with the k -periodicity of this equilibrium, it is equivalently stable under forward dynamics with 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 are (locally) stable under forward perfect-foresight dynamics. In particular, when there is a cycle of period three, it is well known that cycles of any periodicity co-exist with the period-three cycle as equilibrium solutions of the system. This occurs, for example, when ρ_2 is greater than 13.

For later use, we summarize and compare here the stability conditions obtained in the literature under different periodic equilibrium-consistent learning schemes. 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

$$|DG^k(P_k^*)| = \left| \prod_{i=1}^k DG(P_i^*) \right| < 1 \quad (8)$$

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 and Woodford (1991)), and is also equivalent to the *strong E-stability condition* under recursive learning, per Evans and Honkapohja (1995, Proposition 3, p. 197). Such a condition is defined as a stability criterion which is robust to over-parametrization of the forecasting rules of agents. Conversely, if an equilibrium is stable only when the agents' forecasting rule is exactly consistent with its periodicity, it

⁶ The 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.

⁷ A less extreme interpretation of backward perfect-foresight dynamics is that agents have naïve expectations.

⁸ This uses the same parametrization we employ in our experiments, which is $e_1 = 2$, $e_2 = 0.5$, $\rho_1 = 0.5$.

Fig. 1. Bifurcation diagram under backward perfect-foresight dynamics in Grandmont (1985, p. 1030). The red vertical lines indicate the five different ρ_2 values run in the different treatments of the experiment (see Section 3). (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

is said to be *weakly E-stable*. The weak E-stability criterion is, unsurprisingly, less stringent, requiring $DG^k(P_k^*) < 1$ 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 and Woodford (1991) consider a period- k adaptive expectation scheme with a constant parameter $0 < w < 1$:

$$p_{t+1}^e = wp_{t+1-k} + (1-w)p_{t+1-k}^e \quad (9)$$

This is consistent with a period $k \geq 1$ equilibrium, and reduces to backward perfect foresight when $w = 1$ and $k = 2$. 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 $w \rightarrow 1$.

In the special case of a steady state P^* (when $k = 1$), it is stable under rule (9) if and only if $DG(P^*) < 1$ or $DG(P^*) > \frac{2-w}{w}$. In the special case of a 2-cycle (when $k = 2$), its stability condition becomes $-\frac{(2-w)^2}{w} < DG^2(P_{1,2}^*) < 1$. If we define $d \equiv DG^2(P_{1,2}^*)$, then a 2-cycle is stable under rule (9) if and only if one of the two following conditions holds:

- i) $d \leq -1$ and $w \in (0, \underline{w})$, where $\underline{w} = \frac{4 - d - \sqrt{d(d-8)}}{2}$
- ii) $d \in (-1, 1)$

Two other learning mechanisms have been applied to this specific OLG economy, and predict different outcomes from the ones under adaptive learning. Bullard and Duffy (1998) use an heterogeneous agent version of this model with two population groups, in which agents forecast prices using evolutionary learning. A genetic algorithm (GA) selects for the lag k to be used in forming the next period's price forecast, such that $P_{i,t}^e(t+1) = P_{t-k}$. Because of this, their algorithm can, in principle, learn high-order cycles. They conduct numerical simulations of the OLG economy under the same range of ρ_2 values as in Fig. 1, and consistently observe convergence either to the monetary steady state or to the 2-cycle (when it exists, which is approximately when

$\rho_2 > 4$).⁹ Their result supports the intuition that agents tend to use simple forecasting rules and hence coordinate on simple — in the sense of low-order — equilibria.

The second learning mechanism is the so-called Sample-AutoCorrelation (SAC) learning (Hommes and Zhu, 2014), that has been applied by Hommes et al. (2013) to this OLG economy with $\rho_2 = 12$, i.e. when the dynamics in the backward 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. The authors find that only two outcomes — namely, (quick) convergence to the steady state or up-and-down oscillations akin to a ‘noisy’ 2-cycle — emerge as the result of such learning process.

As made clear by the above presentation, theory has its limitations in narrowing down the set of equilibria in that model. In other words, ‘anything goes’, that is, any equilibrium may be stable under some suitable learning process. As Table 1 and the discussion in Section 3.3 will make even clearer, the multiplicity issue is not fully eliminated by any learning stability criterion. Even when narrowing down the set of learnable equilibria, weakly E-stability or adaptive learning require first to discriminate between forecasting rules, as stability then requires agents to use a correctly specified learning rule consistent with the exact periodicity of the equilibrium. Therefore, in order to test which theories of learning are empirically relevant in this complex OLG environment in which multiplicity of equilibria is pervasive, we create and run a laboratory experiment. The following section details its design, hypotheses, and implementation.

3. Experimental design

The experiment employs a *single-population design* along with 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 governed by the OLG model described in Section 2. Each participant repeatedly plays the role of a ‘professional advisor’ working for one *young* individual in each of the T periods.¹⁰ 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 decisions, and subjects do not need to advise them. Our single-population design is motivated by its close relation to the adaptive learning literature, where the learning dynamics follows the intertemporal mapping between expected and realized prices in the model. Nonetheless, in Section 6, we show that our results are robust to a two-population design that preserves the overlapping generation frictions.

We implement two designs. In one design, called the learning-to-forecast experiment (LtFE), subjects’ roles are to submit price forecasts for their respective clients. Based on these forecasts, a computer calculates the young individuals’ conditional optimal savings decisions and imple-

⁹ They also find one 4-period cycle for one simulation at a specific value of ρ_2 , as well as two cases of non-convergence.

¹⁰ Note that the underlying model of the experiment is an infinite-horizon setting, but a constant supply of fiat money should be worthless in a finite implementation (see, for instance, Lim et al. (1994)). We choose not to address this issue in the LtFEs as subjects do not observe money holdings and focus on their price forecasting task. In the LtO condition, addressing this issue would require an additional price setting mechanism for the last period that would considerably complicate the design, and result in a larger disconnect with respect to the LtFEs, without obvious benefits. Indeed, cognitive limitations (see e.g. Hirota and Sunder 2007), as well as strategic uncertainty and the lack of common knowledge about rationality in a group of heterogeneous participants render unlikely that participants would engage in backward-induction reasoning over a hundred periods, and make their saving decisions accordingly. Indeed, we never observed coordination on very low savings in our experiment.

ments them. In the other design, called the learning-to-optimize experiment (LtOE), we test the robustness of our findings by keeping the environment the same but having subjects directly submit savings decisions (i.e., having them make forecasts only implicitly). We present both designs below.

3.1. The learning-to-forecast experiment (LtFE)

At the beginning of every period/generation t , each subject $i = 1, \dots, N$ has to submit a two-period-ahead price forecast $P_{i,t}^e(t+1)$ of the price $P(t+1)$. We assume that every member of the young generation then makes the optimal savings decision, conditional on the price forecast that he receives from his advisor. Using the CRRA utility functions (7) and combining the first-order condition (3) with the budget constraint (1), the optimal consumption $c_{i,t}$ of any young individual¹¹ is implicitly defined by

$$c_{i,t} + c_{i,t}^{(\rho_1/\rho_2)} \frac{P_{i,t}^e(t+1)^{[(\rho_2-1)/\rho_2]}}{P_t} = e_1 + e_2 \frac{P_{i,t}^e(t+1)}{P_t} \quad (10)$$

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

In treatments with high ρ_2 values, we use a transformation of the map G that governs the law of motion of the price. This is because, as ρ_2 increases, real money balances tend towards the bounds 0 and 2, which produces price ranges that are too large to be easily readable on the graph and table shown on subjects' screens (see Fig. 33 in Appendix I). Therefore, for $\rho_2 = 12$ and $\rho_2 = 13.5$ (see below), we map the subjects' price forecasts into the actual price values as follows

$$p = H(\tilde{p}) = 3.5^{\frac{\tilde{p}}{8}} - 1 \quad (11)$$

or, equivalently, as the inverse:

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

with $p, \tilde{p} \in [0, +\infty]$ being 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.¹³

Sequence of events In each generation t , once the N subjects have submitted their price forecasts, the corresponding level of consumption and savings of each young individual, together with the market clearing price $P(t)$, are solved numerically¹⁴ and displayed to the subjects. The consumption levels of the old individuals and their corresponding lifetime (two-period) utility values are determined simultaneously.

The economy then proceeds to the following period $t+1$, at which point the young individuals from t have become old and a new generation of young has been born. The process repeats up

¹¹ Or, conversely, his optimal real money balance $e_1 - c_{i,t}$.

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

¹³ In particular, the number of equilibria and their stability are unaffected, with the slope of any iterate of G and H being the same at any of their fixed points.

¹⁴ As condition (10) does not allow for a closed-form solution of optimal individual consumption levels.

until period T , which indicates the end of the experiment. 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 e.g. Bao et al. (2017) (see Hommes et al. (2005) for some motivations):

$$\max \left(1300 - \frac{1300}{49} (P_{i,t}^e(t+1) - P_{t+1})^2, 0 \right) \quad (13)$$

in which the payoff is maximal and equal to 1300 points in the case of perfect prediction, and equals zero if the prediction error is higher than 7 so as to avoid negative payoff and, together with the range of price values in the experiment and the exchange rate used, ensure reasonable average payoff levels even in case of rather poor forecasting performances. The timing of the payoff is two-periods ahead, as subjects only observe the realized price (and their forecast error) at the end of the following period. A table showing the associated payoff value for each possible forecast error was provided in the instructions.

3.2. The learning-to-optimize experiment (LtOE)

In the LtOE, we drop the assumption of optimal conditional savings given a price forecast, and instead ask the subjects to directly submit the savings decision of the 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 those forecasts. This is essentially because this design focuses on quantity decisions of subjects, and we did not wish to introduce a more demanding cognitive load by combining the tasks of forecasting and optimizing (see Bao et al. (2017) for more discussion). However, subjects are instructed (see Appendix H) that they face a two-stage decision process, and they first may 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 Fig. 33 in Appendix I) 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. Because the price transformation (11) that we use in the LtFE for high ρ_2 values impacts the return on savings, we only consider $\rho_2 = 3, 5$ and 8 for the LtOE (see Section 3.3).

Sequence of events At any generation t , once every subject $i = 1, \dots, N$ has submitted a savings decision $s_{i,t}$ for a member of the young generation, the market clearing price for the consumption good is given by $P_t = \frac{M}{\sum_{i=1}^N s_{i,t}}$. At the same time, the consumption of the old individuals and their

two-period utility are determined, after which the economy proceeds to period $t + 1$, and so forth until period T .¹⁵

Payoff Each subject earns a payoff based on the realized lifetime (two-period) utility of the individual over his *two-period* life. As in the LtFE, the timing of the payoff of any savings decision is then two periods ahead: a savings decision made for any member of the young generation in period t is rewarded at the end of period $t + 1$, once the consumption when old is revealed. In order to implement this payoff scheme in the lab, we use two transformations of the two period utility function U (with separable utility functions given by (7)). First, in order to rule out negative payoffs, we apply the following transformation of the utility:

$$\tilde{u} = \max \left(K \times (U(c_{i,t}, c_{i,t+1}) + C), 0 \right) \quad (14)$$

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.¹⁶ We use the payoff function (14) 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 inter-generational equity (Grandmont, 1985). This means that utility values along cycles, for instance along the 2-cycle and at the monetary steady state, may differ. In order to be consistent with the LtF design, where subjects' payoffs are maximized (and identical) along *any* perfect equilibrium path, we apply the following transformation of the payoff function:

$$\hat{u} = 1300 \times \left(\frac{\tilde{u}}{1300} \right)^\alpha \quad (15)$$

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 2-cycle is of the same order of magnitude as the payoff at the steady state. Without this transformation, the monetary steady state would be the only payoff-maximizing equilibrium over the two generations in the model, so that coordination on the steady state would be *ex ante* favored. Such a result is not a very useful one for our purposes. When $\rho_2 = 5$ and 8 (see Subsection 3.3), the payoff of any given savings decision s_t is given by (15).

Similarly to the LtFE, the instructions given to the subjects included a two-dimensional payoff table that reported the expected payoff to savings decisions for given (expected) values of the return on savings (see Appendix H). The optimal savings decisions conditional on each expected return on savings then correspond to the consumers' offer curve, the shape of which is unaffected by the transformations of the utility functions that we have considered.

¹⁵ We 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.

¹⁶ See Marimon et al. (1993) for a similar transformation of the utility function.

3.3. Treatments and hypotheses

We adopt the calibration used in Section 2: $e_1 = 2$, $e_2 = 0.5$, $\rho_1 = 0.5$. We then vary the parameter ρ_2 to define different treatments with increasingly complex equilibrium outcomes.¹⁷ The first hypothesis that we aim to test by bringing this OLG economy to the lab is whether or not coordination of a group of subjects with heterogeneous beliefs may happen at all as the result of repeated market interactions in such a complex environment.

Hypothesis 1 (*Selection of an equilibrium*). *Subjects coordinate their forecasts so that the resulting market price converges towards a perfect-foresight equilibrium of the underlying model.*

Given the cost of lab implementation, we can only investigate a few cases, and consider five treatments representing a variety of typical cases in non-linear dynamics: (i) a stable steady state ($\rho_2 = 3$), (ii) a stable 2-cycle ($\rho_2 = 5$), (iii) chaos with (unstable) 2^k cycles ($\rho_2 = 8$), (iv) chaos, with unstable cycles of any periodicity but three ($\rho_2 = 12$), and (v) a stable 3-cycle and unstable cycles of any periodicity ($\rho_2 = 13.5$). Furthermore, they are distinct enough from the bifurcation values to ensure clear-cut illustrations of those dynamics (see Fig. 1). Additionally, the case $\rho_2 = 8$ lies after but close to the limiting value of the cascade of period-doubling bifurcation route to chaos, so as to imply a wide range of existing equilibria along which price equilibrium values are within a range to be also implemented in the LtOE. Table 1 summarizes the stable equilibria under different learning theories from the literature (discussed in Section 2.4) for each of these five treatments.

Let us first consider E-stability. Recall that strong E-stability corresponds to stability under backward perfect foresight, which is the inverse of stability in the forward perfect-foresight dynamics. For $\rho_2 = 3$, the monetary steady state is strongly E-stable, and the map G does not have any other equilibrium.¹⁸ For $\rho_2 = 5$, the period-two cycle is the only strongly E-stable outcome, and the monetary steady state is weakly E-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. All other existing equilibrium cycles (i.e. those with periodicities which are multiples of 2^k , $k \geq 3$) are E-unstable.

For $\rho_2 = 12$ and $\rho_2 = 13.5$, the map G has chaotic dynamics with infinitely many periodic cycles¹⁹ and chaotic orbits, along with the monetary steady state. When $\rho_2 = 12$, none of those equilibrium cycles are strongly E-stable, while only the steady state and the 2-cycle are weakly E-stable. When $\rho_2 = 13.5$, the only strongly E-stable outcome is a period-three cycle, while the steady state and the 2-cycle are the only weakly E-stable equilibria.

The stability of an equilibrium k -cycle under the adaptive expectation rule (9) (conditional on the rule being consistent with periodicity k) depends on the weight w given to prices from the period $t + 1 - k$. Under a second order rule, only the steady state and the 2-cycle can be selected. Moreover, cycles of period 2^k are created by period-doubling bifurcations, so that the

¹⁷ The 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' screens, and it becomes impossible to construct a bounded path that never repeats any past value. However, it still leaves room for high order cycles.

¹⁸ In 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, when submitting savings decisions, subjects cannot submit a savings decision of zero, so the price level is always defined and the monetary steady state is the only feasible steady state.

¹⁹ For all periodicities but three for $\rho_2 = 12$, and including three for $\rho_2 = 13.5$.

Table 1
Summary of the stability of equilibria under theoretical learning predictions in the five treatments for different ρ_2 -values.

Expectations	homogeneous					heterogeneous	
	forward	backward	strong	weak	adaptive	SAC	GA
	perfect foresight Grandmont (1985)		<i>E</i> -stability Evans and Honkapohja (1995)		expectations Guesnerie and Woodford (1991)	learning Hommes et al. (2013)	learning Bullard and Duffy (1998)
$\rho_2 = 3$	none	SS	SS	SS	SS	SS iff $\beta = 0$	SS
$\rho_2 = 5$	SS	2-cycle	2-cycle	SS 2-cycle	SS 2-cycle $\forall w$	SS iff $\beta = 0$ 2-cycle	SS 2-cycle
$\rho_2 = 8$	2^k -cycles	none	none	SS 2-cycle 4-cycle	SS 2-cycle (if $w < 0.8$) 2^k -cycles ($k > 1$) (if w low enough)	SS iff $\beta = 0$ noisy 2-cycle ($\beta \simeq -1$)	SS 2-cycle
$\rho_2 = 12$	all cycles (period $\neq 3$)	none	none	SS 2-cycle	SS 2-cycle (if $w < 0.61$) any cycle ($k \neq 3$, if w low enough)	SS iff $\beta = 0$ noisy 2-cycle ($\beta \simeq -1$)	SS 2-cycle
$\rho_2 = 13.5$	all cycles (period $\neq 3$)	3-cycle	3-cycle	SS 2-cycle 3-cycle	SS 2-cycle (if $w < 0.57$) 3-cycle $\forall w$ any cycle $k > 3$ (if w low enough)	SS iff $\beta = 0$ noisy 2-cycle ($\beta \simeq -1$)	SS 2-cycle

Notes: SS stands for the monetary *Steady State*, SAC learning for *Sample AutoCorrelation* learning, GA for *Genetic Algorithm*. The stability of any cycle under adaptive expectations is conditional on agents using an adaptive rule consistent with the cycle’s periodicity. Results under GA learning are obtained through numerical simulations. We highlight in bold the predictions which are supported by our experimental evidence (see Section 4).

derivative of the second iterate of the map G , $DG^2(P_{1,2}^*)$ at the 2-cycle is always negative. For $\rho_2 = 5$, since the 2-cycle is strongly E-stable — which is equivalent to $DG^2(P_{1,2}^*) \in (-1, 0)$ — the 2-cycle is stable for any w value. When $\rho_2 = 8, 12$ or 13.5 , we have $DG^2(P_{1,2}^*) < -1$. Given our calibration, the stability thresholds $\underline{w}$ are as follows:

$$\underline{w} \simeq \begin{cases} 0.8 & \text{when } \rho_2 = 8 \\ 0.61 & \text{when } \rho_2 = 12 \\ 0.57 & \text{when } \rho_2 = 13.5 \end{cases}$$

The increasing complexity of the model dynamics under these five treatments aims to test the following hypothesis:

Hypothesis 2 (*Selection of simple equilibria*). *Coordination on higher-order cycles or complicated dynamics, if any exist, is less likely than coordination on simple equilibria, such as a steady state or a cycle of low periodicity.*

Previous results in LtFEs (albeit in much simpler linear environments) have documented the use of simple forecasting rules that involve past information with only a few lags (typically one or two). If this is true here as well, the learning theories reviewed in Section 4 predict that the resulting selected equilibrium should be a lower-order cycle (or even the steady state). This is supported by findings from cognitive psychology, in which the sequence-learning literature concludes that humans are only good at learning patterns of up to a handful of prior observations (see Spiliopoulos (2012) and the references therein). Forecasting can be viewed as akin to sequence prediction, where a period k -cycle is akin to a pattern of length k ; under such an assumption, subjects are also more likely to coordinate on lower-order cycles.

Consequently, this prediction arises at least in part from the limits of subjects' memories and cognitive capacities. This would suggest the use of higher-order adaptive rules to be an unlikely occurrence. Conversely, the relative naïveté needed to sustain complex dynamics could result in higher forecasting errors than would otherwise be the case, leading to a conscious decision to aim for more elaborated forecasting rules.

Finally, we use the LtOE to assess the robustness of our hypothesis testing procedure with respect to the experimental task.

Hypothesis 3 (*Robustness to a learning-to-optimize procedure*). *The coordination pattern and resulting equilibria (if any) selected in the learning-to-forecast experiments are also observed in the learning-to-optimize experiments.*

The robustness of our results to a LtO design is an interesting and relevant exercise because, with very few exceptions (see, e.g., Evans and McGough (2018)), the learning literature is entirely focused on the expectations component of dynamic macroeconomic models. This treats the process of optimal decision making conditional on such expectations as a trivial one. In reality, however, agents have to make economic decisions based on implicit forecasts and therefore effectively employ a two-step decision process.

As such, we treat the robustness to LtO design of predictions drawn from learning theories as a matter of empirical concern. Previous experimental evidence does suggest that optimizing is a more complicated task than forecasting, and generates noisier aggregate outcomes as a result (see the trial sessions discussed in Marimon et al. (1993) in an OLG model; see Bao et al. (2013, 2017) in, respectively, a cobweb and an asset-pricing model). Our experimental environment is

Table 2
Summary of the baseline experimental treatments and designs.

ρ_2	3	5	8	12	13.5
<i>Equilibrium values</i>					
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}
<i>LtFE</i>					
<i>Nb. of periods</i>	50	100	100	100	100
<i>Nb. of groups</i>	4	4	4	4	4
<i>Exchange rate</i>	0.00027 (1300 points = 0.35E)				
<i>Maximum payoff</i>	1300 points (0.35E) per period on any perfect foresight equilibrium				
<i>LtOE</i>					
<i>Nb. of periods</i>	100	100	100		
<i>Nb. of groups</i>	4	4	4		
<i>Initial price P_0</i>	50	50	15		
<i>Exchange rate</i>	0.00045	0.00039			
<i>Parameters of the payoff function</i>					
C	0	1	15.25		
K	500	350.5	71.92		
α	NA	4	25		
<i>Payoff on the SS</i>	975 (0.44E)	723 (0.28E)	639 (0.26E)		
<i>Payoff on the 2-cycle</i>	NA	$\frac{40+1310}{2} = 675$ (0.26E)	$\frac{15+1301}{2} = 658$ (0.27E)		

Notes: The equilibrium price values for $\rho_2 = 12$ and 13.5 in the LtFE correspond to the transformed values given by Equation (11). The exchange rate between experimental currency and euros is higher in the LtOE than in the LtFE in order to ensure comparable earnings for subjects with respect to the relative length of the experimental sessions. The initial price in the LtFE is determined by the first submitted forecasts due to the two-period ahead structure. In the LtOE, it has to be initialized to determine the first return on savings in the end of period 1. SS stands for *Steady State*. The parameters of the payoff function in the LtOE correspond to Equations (14) and (15).

more complicated than those in related studies, and none of them explicitly examines equilibrium selection in a setting with multiple potential equilibria (let alone in the presence of complex dynamics). How the outcomes of our LtFE differ from those in a LtOE in such a complex environment is then an open question.

3.4. Implementation

The experiment was programmed in Java using the software package PET²⁰ and was run at the CREED laboratory at the University of Amsterdam over the periods of November–December 2014, February–May 2015 and April–June, 2018. A total of 192 subjects were recruited from the CREED subject pool²¹ to participate in the 32 baseline experimental economies of $N = 6$ subjects each. Table 2 reports the main features of these different treatments and designs, along with the equilibrium values of prices and savings at the monetary steady state and on

²⁰ The PET software was developed by AITIA, Budapest under the FP 7 European project *CRISIS*, Grant Agreement No. 288501. It is available at <http://www.aitia.ai/en/web/iaws/downloads>.

²¹ The CREED subject pool is composed of students from all fields, and includes both undergraduates and graduates.

the perfect-foresight 2-cycle. We ran 4 economies (groups of 6) per treatment, for a total of 20 LtF economies and 120 subjects, and 12 LtO economies and 72 subjects for each of the values $\rho_2 = \{3, 5, 8\}$. Each experimental economy with $\rho_2 = 3$ was run for $T = 50$ generations/periods, as pilot observations indicated a very quick stabilization in the LtFE, and we ran all the other treatments for $T = 100$ periods.

The computer interfaces of the LtFE and the LtOE are provided in Appendix I, while the instructions, payoff tables, and questionnaires are contained in Appendix G for the LtFE and Appendix H for the LtOE. Subjects received a detailed description of the OLG environment underlying the experiment, along with their experimental task and their payoff. Following Marimon et al. (1993), we refer to the consumption good as ‘chips’. The participants were given the opportunity to read the instructions at their own pace, and were then asked to fill in a quiz on paper. The instructors then checked that each individual subject was able to correctly answer all of the questions. If an incorrect answer had been given, the experimenter privately explained to the participant what the correct one was. Only when all participants had answered every question correctly was the experiment started.

This procedure allows us to be reasonably certain that every subject understood both 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’ payoffs were expressed in points, which were converted into euros at the end of the experiment at an exchange rate given in the instructions; mean participant earnings came out to 23.6 euros. Each experimental session lasted around 2 hours on average, including an average of 40 minutes to complete the instructions and questionnaire. These times, however, exhibited strong disparities across treatments (see below for details).

4. Experimental results

Do subjects achieve coordination on a perfect-foresight equilibrium in this complex environment with a pervasive multiplicity of equilibria? and if yes, which one do they select? The qualitative features of our experimental results, summarized in Figs. 2–6, together with the quantitative measures in Table 3 hereafter, provide clear-cut answers to these questions.

The figures display the observed aggregate saving levels (the real money balances) in the 20 groups of the LtFEs and the 12 groups of the LtOEs respectively (with one graph per ρ_2 -value treatment, and one line per group). For all 32 experimental economies separately, the dynamics of the individual price forecasts or saving decisions together with the realized price or savings are reported in the figures in Appendices A and B.

Table 3 quantifies these experimental outcomes for each group along five dimensions, that we successively detail in the sections below. The first three dimensions – namely i) the type of equilibrium selected, ii) the *average relative distance of the price to equilibrium* (ARDE, in absolute value and percentage points), and (iii) the first-order autocorrelation ρ_s of aggregate savings, are used in Section 4.1 along with Figs. 2–6 to characterize aggregate behavior and price convergence. In Section 4.2, we then discuss the coordination of individual saving decisions using the *average relative standard deviation* (denoted by $\sigma(s_i)/\mu_{s_i}$) of individual saving decisions among the six subjects. Finally, in Section 4.3, we investigate the efficiency of the subjects’ decisions in terms of earnings using the *earnings efficiency ratio* (EER), i.e. the ratio (in percentage points) of realized payoffs during the whole experiment to the maximum amount of points possible in equilibrium.

Fig. 2. Real money balances time series for $\rho_2 = 3$. Note: The dashed red line corresponds to the monetary steady state.

Fig. 3. Real money balances time series for $\rho_2 = 5$. Note: The lower and upper blue dotted lines correspond to the perfect-foresight 2-cycle, the middle dashed red line to the monetary steady state.

4.1. Convergence of aggregate behavior

Figs. 2–6 already illustrate the first three main findings from the experiment. We start by highlighting them before providing formal support. First, in line with Hypothesis 1:

Finding 1 (*Systematic equilibrium selection*). *In all LtF experimental economies, the price (or, equivalently, aggregate savings) converges towards a perfect-foresight equilibrium.*

Convergence by the sole force of repeated market interactions already constitutes a remarkable feature of our experiment since spontaneous, systematic equilibrium selection was not assured, *a priori*, in such a complex experimental environment given (initially) heterogeneous beliefs.

Second, in line with Hypothesis 2:

(a) The 4 LtF groups (3 2-cycles, 1 steady state)

(b) The 4 LtO groups (4 steady states)

Fig. 4. Real money balances time series for $\rho_2 = 8$. Note: See Fig. 3.

(a) The 4 LtF groups (4 2-cycles)

Fig. 5. Real money balances time series for $\rho_2 = 12$. Note: See Fig. 3.

(a) The 4 LtF groups (4 2-cycles)

Fig. 6. Real money balances time series for $\rho_2 = 13.5$. Note: See Fig. 3.

Table 3

Summary statistics of the baseline experimental economies.

Group	1	2	3	4	1	2	3	4
$\rho_2 = 3$								
	LtFE				LtOE			
Equilibrium	SS	SS	SS	SS	SS	SS	SS	SS
ARDE	0.01	0.17	0.12	0.03	4.13	12.48	6.26	6.76
ρ_s	−0.17	−0.07	−0.13	−0.37	0.39	−0.51	0.55	0.31
$\sigma(s_i)/\mu_{s_i}$	0.05	1.39	0.54	0.13	11.71	32.68	13.55	15.68
EER	95.6	95.3	91	94.6	97.2	95.3	96.6	95.6
$\rho_2 = 5$								
	LtFE				LtOE			
Equilibrium	2-cycle	2-cycle	2-cycle	SS	SS	SS	SS	SS
ARDE	10.19	12.44	7.99	0.09	3.13	4.24	3.96	7.16
ρ_s	−0.98	−0.94	−0.94	−0.33	0.18	0.65	−0.03	−0.72
$\sigma(s_i)/\mu_{s_i}$	7.78	8.82	0.37	0.00	5.91	12.6	8.57	11.09
EER	81.8	69.2	83.9	96.8	98.1	92.4	94.1	91.3
$\rho_2 = 8$								
	LtFE				LtOE			
Equilibrium	2-cycle	2-cycle	2-cycle	SS	SS	SS	SS	SS
ARDE	3.31	19.23	5.85	46.75	3.89	6.18	4.74	6.82
ρ_s	−0.95	−0.95	−0.96	−0.87	−0.17	−0.28	0.58	0.17
$\sigma(s_i)/\mu_{s_i}$	2.42	14.8	10.22	19.7	10.43	8.79	1.8	18.45
EER	77.9	87.9	80.4	98	93.4	88.5	81.1	96.5
$\rho_2 = 12$ (LtFE only)					$\rho_2 = 13.5$ (LtFE only)			
Equilibrium	2-cycle	2-cycle	2-cycle	2-cycle	2-cycle	2-cycle	2-cycle	2-cycle
ARDE	7.03	1	0.48	0.47	2.24	1.58	0.78	0.5
ρ_s	−0.93	−0.96	−0.96	−0.97	−0.95	−0.97	−0.96	−0.96
$\sigma(s_i)/\mu_{s_i}$	22.54	11.32	2.92	2.06	7.38	7.1	4.8	2.3
EER	54.3	62	69.4	73.1	55.9	64.6	68.4	88.3

Notes: For each experimental group the table reports: the selected equilibrium (SS stands for *Steady State*); ARDE measures the *Average Relative Distance of the price to this Equilibrium* (in absolute value and percentage points) over the last 25 periods of each group— for instance, if the ARDE is 10, aggregate prices are on average 10% away from their equilibrium value over the last 25 periods; $\sigma(s_i)/\mu_{s_i}$ is the relative standard deviation of savings (in percentage points) among the N subjects, averaged over the last 25 periods (in the LtFEs, we compute this quantity from the savings decisions implied by the price forecasts of the subjects); ρ_s is the first-order autocorrelation of aggregate savings over the last 25 periods; EER stands for *Earnings Efficiency Ratio*, expressed in percentage points, and measures the number of points on average earned by the subjects over the whole periods of the experiments w.r.t. the maximum amount of points possible in equilibrium.

Finding 2 (*Characterization of the selected equilibria*). In all LtF experimental economies, the price converges towards either the monetary steady state or the 2-cycle.

Phrased differently, we do not observe any chaotic long-run dynamics or cycles of period three or higher periods, but we do witness cycles of two periods in addition to the steady state. We shall stress the novelty of this result: our experiment is the first to observe *spontaneous* coordination on a 2-cycle without the use of any training phase or signals akin to sunspots.

Third, Finding 3 provides some insight into the outcomes in the 12 LtOEs (see Figs. 2b, 3b and 4b):

Finding 3 (Robustness to LtOE). *In the LtOE, the monetary steady state is the only selected perfect-foresight equilibrium.*

We now support these three findings by quantitative measures from Table 3. To compute those measures, we use the last 25 periods of each experimental economy to discard occasionally long out-of-equilibrium transient phases, while considering enough data points not to overweigh the effects of temporary individual deviations. We primarily use the ARDE to assess convergence to an equilibrium: an ARDE ϵ means that the average relative distance of the price to that equilibrium over the last 25 periods is $\epsilon\%$. Hence, the lower this number (in percentage points), the more accurate the convergence to the corresponding equilibrium. The selected equilibrium reported in Table 3 for each economy is the one that returns the lowest ARDE.

Although somewhat arbitrary, a formal definition of convergence may help summarize our results. For instance, one could define convergence as 10%-convergence, which would be satisfied for 27 out of the 32 economies. Three out of the remaining five still lie within a 12.5%-neighborhood, which can be considered as near-equilibrium behavior (the ARDE for those groups are 10.19, 12.44 and 12.48%). The last two groups (with ARDE-values 19.23 and 46.75%) deserve a closer look before being classified. As we detail below, they in fact correspond to a case of slow convergence to a 2-cycle, and a ‘noisy’ steady state.

Let us first detail the outcomes from the 20 LtFEs for the five different ρ_2 -values (see Fig. 2a, 3a, 4a, 5a and 6a). When $\rho_2 = 3$, the steady state is the only perfect-foresight equilibrium and it is stable.²² Unsurprisingly, the corresponding four experimental economies very quickly and almost perfectly converge to the steady state, as testified by ARDE values below 0.5% in all four economies.

In the 16 LtF economies with multiple perfect-foresight equilibria, we only observe two instances of convergence to the unstable steady state. In the first case ($\rho_2 = 5$, Group 4), convergence occurs quickly and is almost perfect (ARDE = 0.09%). The second case ($\rho_2 = 8$, Group 4) corresponds to the instance of ‘noisy’ convergence mentioned above, as the price oscillates around the steady state with strong negative autocorrelation ($\rho_s = -0.87$).²³

The remaining 14 economies converge towards the 2-cycle, with first-order autocorrelation of aggregate time series close to -1 (≈ -0.95) regardless of whether the 2-cycle is stable (when $\rho_2 = 5$) or unstable (for higher values of ρ_2). Convergence is accurate, as quantified by low ARDE values (less than 3.3% in 8 economies; less than 8% in 11 economies). This is especially the case in the eight economies of the strongly unstable treatments ($\rho_2 = 12$ and 13.5), where the only selected equilibrium is the (unstable) 2-cycle and convergence is near-perfect for 7 out of 8 groups (the corresponding ARDE values are at most 2.25%).

These results seem to suggest that the more unstable the steady state, the more likely the selection of the 2-cycle, whether stable or not. Admittedly, group experiments can only provide a small number of data points but, bearing in mind this inherent limitation of the method, our experimental results provide a clear-cut picture of equilibrium selection in this complex environment.

²² Hereafter, we refer to a stable equilibrium as stable in the backward perfect-foresight dynamics as depicted in Fig. 1, which also corresponds to stability under naïve expectations and to strong E-stability (see Section 2). Similarly, by ‘more unstable’, we mean a higher (in absolute terms) value of the slope of the derivative of the map G at this equilibrium.

²³ In fact, the rather high ARDE value of 47.75% is partly due to the particularly low price steady state value in this treatment – 0.72 – but the price in this group constantly oscillates between a minimum of 0.3 and a maximum of 1.58, with a close-to-equilibrium average of 0.76.

Furthermore, we observe from the figures that convergence to the 2-cycle itself sometimes only happens after a long transition: 6 of the observed groups took at least 50 periods to do so, see Groups 1 and 4 in the $\rho_2 = 12$ treatment, Groups 1, 2, and 3 in the $\rho_2 = 13.5$ treatment and Group 2 in the $\rho_2 = 8$ treatment. This latter group is the instance of particularly long convergence mentioned previously: the ARDE over the last 25 periods is the largest among the 2-cycles (19.23%) because convergence occurs only in the last 10 periods – if computed only over the last 10 periods, the ARDE drops to 8.58%, and even to 2.04% over the last 5 periods.

One may wonder whether there is any statistical difference in the goodness of convergence between the two outcomes, a steady state or a 2-cycle. To do so, Fig. 7a reports the cumulative distribution of the ARDE values of each experimental economy. Hence, each dot or triangle represents one ARDE value from Table 3. Following the distinction established in Table 3, we separate the LtF economies that converge to the steady state from those that converge to the 2-cycle. A K-S test leads us to reject the null hypothesis that these economies have equal distributions of this distance value, in favor of the alternative hypothesis that economies converging to the steady state have lower average distance (p-value = 0.009).²⁴ We therefore conclude that in the LtFE *aggregate convergence is significantly better when the dynamics converge towards the steady state than when they converge towards the 2-cycle*.

We next take a closer look at the 12 LtOEs. The ARDE values support the visual impression conveyed by Figs. 2b, 3b and 4b that all economies in the LtO condition converge towards the steady state: 6 out of 12 economies display an ARDE below 5%; all but one below 7.16% and only one ($\rho_2 = 3$, Group 2) converges in a neighborhood slightly larger than 10% (12.48%).

Yet, the convergence on the steady state in the LtOE seems somewhat weaker than in the LtF treatment. While the experimental economies converge almost perfectly to the steady state in the LtFE (with ARDE values typically lower than 0.5%), they only converge to a (close) neighborhood of the steady state in the LtOE (with ARDE values typically below 7%). This difference is also notable from a visual inspection of Fig. 2b, 3b and 4b²⁵: the LtOEs appear to experience more price volatility than do the LtFEs. A K-S test on the distributions of the ARDE values (that are also displayed in Fig. 7a) reveals that this visual impression is statistically significant (the p-value of the unilateral K-S test is 0.004).

We conclude that, *even though all LtO groups select the steady state, convergence is significantly better in the LtFEs that do so*. Section 6.1 below develops additional elements of discussion on the absence of 2-cycle equilibrium in the LtO condition, in contrast to the LtF condition.

To summarize so far, taking the LtO and the LtF treatments together, more than half of the economies (19 out of 32) converge within a 5%-neighborhood of this equilibrium, all but 5 converge within 10% of this equilibrium and all but 2 are still within 12.5%, which we can define as near-equilibrium behavior. The two remaining economies correspond to a case of particularly slow but successful convergence to a 2-cycle, and a ‘noisy’ convergence to the steady state. Therefore, we conclude that, *in this complex lab environment, only simple equilibria, namely the monetary steady state or the 2-cycle, are selected as coordination devices, and may, accordingly, be viewed as empirically most relevant*.

Before, turning to individual participants’ data, we conclude the discussion of the aggregate convergence outcomes in light of learning predictions summarized in Table 1. In order to give

²⁴ We use K-S tests throughout the analysis, but the results are robust to the use of Wilcoxon rank sum tests instead.

²⁵ To see that, compare in particular Figs. 2a and 2b when $\rho_2 = 3$ and the monetary steady state is both stable and the only available equilibrium. See also the figures reported in Appendices A and B.

Fig. 7. Cumulative distributions of experimental statistics. **Note:** Each point on the three graphs represents the numbers given in, respectively, row 2, 4 and 5 of Table 3: the ARDE and the relative standard deviation (RSD) of individual savings are averaged over the last 25 periods of each economy and the Earning Efficiency Ratios over the whole experiment. We separate the LiF economies that converge to the steady state from those that converge to the 2-cycle, and the LiO economies according to the classification established in Table 3.

an accurate prediction of the experimental results, a learning selection device should pick up the steady state or the 2-cycle as a stable outcome of the learning dynamics. In that respect, the first striking observation is that none of the reviewed learning theory exactly predicts the equilibrium selection in our experiment. The GA learning yields the closest predictions but we do not observe any coordination on the steady state when the complexity of the model increases (increasing ρ_2 -values), while the study by Bullard and Duffy (1998) does. Furthermore, neither backward nor forward perfect-foresight dynamics can predict the outcomes in the experimental economies. The two observed outcomes – the steady state and the 2-cycle – are not strongly E-stable (except when $\rho_2 = 3$ for the steady state, and when $\rho_2 = 5$ for the 2-cycle), they are only *weakly E-stable* in all treatments. Hence, the criterion of *weak E-stability* appears as a *necessary condition* for an equilibrium to be selected in our experiments, but it is *not a sufficient condition*: indeed, not all weakly E-stable cycles are observed in our experiments; the 4-cycle when $\rho_2 = 8$ and the 3-cycle when $\rho_2 = 13.5$ are never selected (see again Table 1).

Selection of the monetary steady state or the 2-cycle regardless of their E-stability makes a clear case for simple forecasting rules, involving — at most — information from period $t - 1$ on. Indeed, from Table 1, convergence to the 2-cycle or the steady state is only predicted if subjects use some appropriately-weighted second-order adaptive rule or econometric learning based on AR(1) rules. These types of forecasting rules are more sophisticated than naïve expectations, but simpler than higher-order adaptive processes. Section 5 below discusses this point in greater detail based on subject-level data series, which also supports Hypothesis 2. Together, these observations give us Finding 4:

Finding 4 (*Learning predictions versus experimental observations*). *In light of the experimental data,*

- (a) *No learning criterion exactly predicts our experimental outcomes.*
- (b) *Weak E-stability is a necessary but non-sufficient condition for an equilibrium to be selected.*
- (c) *The criterion of strong E-stability does not predict the selected equilibria.*

- (d) *The selected equilibria are consistent with adaptive or learning rules using only information from period $t - 1$ (second-order adaptive rules).*
- (e) *In the LtOE, the monetary steady state is the selected perfect-foresight equilibrium, regardless of its stability under learning.*

4.2. Coordination between individual decisions

The previous subsection has shown that aggregate savings in both the LtFEs and LtOEs converge to simple equilibria, but does this imply that individual decisions are coordinated on these equilibrium values or does some degree of heterogeneity remain? A simple and intuitive measure of individual coordination is the relative standard deviation of (implied) individual savings decisions, $\sigma(s_i)/\mu_{s_i}$. The smaller this number, the lower heterogeneity and therefore the stronger coordination between individual decisions. Table 3 reports the average over the last 25 periods of this value for each experimental economy. Fig. 7b also plots the cumulative distributions of those values (each dot or triangle being a value reported in Table 3), by distinguishing between LtFEs converging to the steady state, LtFEs converging to the 2-cycle, and the LtOEs (all converging to the steady state). Our main results can be summarized as follows:

Finding 5 (Coordination between individual decision).

- (a) *Subjects coordinate their forecasts better on the steady state than on the 2-cycle.*
- (b) *Subjects coordinate their forecasts better than they coordinate their savings decisions.*

Let us start by the LtFE converging to the steady state: the individual coordination is almost perfect, with heterogeneity less than 1.39% in 5 out of 6 groups; only in the ‘noisy’ converging group ($\rho_2 = 8$, group 4) heterogeneity is somewhat larger (19.7%). For the 14 LtF economies converging to the 2-cycle, more individual heterogeneity is observed, but overall individual coordination is still rather high, with relative standard deviations less than 5% for 6 out of 14, and less than 10% for 10 out of 14 economies, and only one case (Group 1, $\rho_2 = 12$) slightly larger than 20%. As is the case with the previous subsection, in the LtFEs, we find that *subjects coordinate significantly better on the steady state than on the 2-cycle*: the K-S test with the alternative hypothesis that standard deviations are lower in the steady state than on the 2-cycle decisively rejects its null hypothesis with a p-value of 0.0076.

The better coordination on the steady state as compared to the 2-cycle can be at least partly explained by subjects’ mistakes when entering their decisions in the experimental software. Intuitively, the likelihood of making mistakes when entering price predictions should be higher when alternating between high and low forecasts than when entering a constant number. Several subjects indeed reported in their post-experiment questionnaires that they had made typos.²⁶ However, the 2-cycle appears to be a long-run outcome that is robust against individual deviations; even when it is temporarily ‘disturbed’ after a subject’s individual mistake, the dynamics settle back down to the 2-cycle after a few periods.

²⁶ For instance, in Group 1 of the LtFE with $\rho_2 = 5$.

As for the comparison between LtF and LtOEs, Fig. 7b contrasts the coordination between subjects in the LtFE converging to the steady state versus those in the LtOE. We find that *the LtOEs display more heterogeneity between subjects than do the LtFEs*.²⁷

More homogeneous price predictions as opposed 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 two-dimensional payoff matrix in which the savings decisions are discretized, but that the instructions insist on the fact that they can submit any number (up to two decimal places). We find that, overall, 60% of the savings decisions are multiples of 5 (50, 55, 60, etc.), and 47% are multiples of 10 (50, 60, etc.). This is unlikely to result from the numbers used in the payoff table; pilot sessions using a table in A3-format with a finer grid report the same type of decisions. By contrast, only 36% of price predictions are integer values, and most of these at the beginning of the experiment (when subjects have less historical information). This tendency to submit round numbers may be a consequence of payoff values being less sensitive to decision accuracy near the steady-state under the LtOE than under the LtFE (see payoff tables in Appendix H). With flatter payoff values, the LtOE subjects have less of a monetary incentive to refine their savings decisions.

A second explanation for the more heterogeneous savings decisions as compared to price predictions could be strategic behavior by certain subjects. Five subjects reported in the post-experiment questionnaire that they intentionally deviated from the average savings values in their group 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. In doing so, they sought to increase the return on savings so as to reach the payoff-maximizing region of the payoff table, even though those attempts resulted in payoff losses. We now take a closer look at the participants' earnings.

4.3. Participants' earnings

In order to evaluate the efficiency of the participants in performing their experimental task, we make use of the earnings efficiency ratio (EER). In a given period of the LtFE, making a perfect prediction yields a maximum payoff of 1300 points — an amount which declines as the absolute error increases. In the LtOE, the maximum points are given by the transformed values of the utility function on the payoff table.

Table 3 returns the EER for each experimental economy and Fig. 7c reports the cumulative distributions of those values in the experiments. We see that the LtFEs converging to steady state as well as the LtOEs (all converging to steady state) yield very high numbers for efficiency, almost always larger than 90% or even 95%. By contrast, the LtFEs converging to the 2-cycle display lower efficiency rates, around 80% for $\rho_2 = 5$ and $\rho_2 = 8$ and even lower values — around 60% for $\rho_2 = 12$ and $\rho_2 = 13.5$. These differences are statistically significant: *the ratios are significantly higher in the case of convergence to the steady state than of convergence to the 2-cycle*.²⁸

This result is easily explained by the convergence times recorded in the experiments: convergence is much quicker towards the steady state than towards the 2-cycle, reflected in the

²⁷ The corresponding p-value of the one-sided K-S test is 0.0039. For the sake of completeness, we also compare the individual coordination of the LtFE converging to a 2-cycle to the LtOE converging to a steady state. A K-S test does not reject the null hypothesis that these economies display similar levels of heterogeneity (p-value of 0.0733).

²⁸ The corresponding p-value of the one-sided K-S test is 0.0002.

long transition periods observed when $\rho_2 = 8, 12$ or 13.5 . Once at the steady state, the deterministic structure of the model stabilizes price dynamics, and subjects easily make perfect (payoff-maximizing) forecasts.

As it happens, efficiency ratios do not differ significantly between the steady state LtFEs and the LtOE. ²⁹ In the LtOEs, all economies converge to the steady state, around which utility is less sensitive to individual decisions (essentially for savings values between 0.5 and 0.6). Utility in the LtOEs therefore tends to be nearly maximized despite some heterogeneity and small persistent deviations from the optimal savings decision.

By contrast, along the transition phases towards the 2-cycle in the LtFEs, price values oscillate in an irregular manner without any clear pattern (see, e.g., Group 1 in both $\rho_2 = 12$ and 13.5), and subjects make large forecast errors. Therefore, the economies with $\rho_2 = 12$ and $\rho_2 = 13.5$ display the lowest earnings efficiency ratios across all experimental sessions (see again Table 3). ³⁰

After discussing overall participants' efficiency in performing their experiment task, we now take an in-depth look at individual time series and seek to answer the following questions: how do subjects in the lab learn to coordinate on the steady state or the 2-cycle and which behavioral rules do they use?

5. Estimation of individual forecasting rules

This section focuses on the analysis of individual behavioral rules in the LtFEs. This is because the forecasting rules can be directly linked to the theoretical predictions derived from the learning literature presented in Table 1. Therefore, we fit forecasting rules to the price prediction data for the 120 LtFE subjects in order to check consistency with those learning predictions. ³¹

As a baseline, for each participant to the LtFE, we estimate the general forecasting rule:

$$p_{i,t+1}^e = \alpha + \beta_{P_{t-1}} P_{t-1} + \beta_{P_{t-2}} P_{t-2} + \beta_{p_{i,t}^e} p_{i,t}^e + \beta_{p_{i,t-1}^e} p_{i,t-1}^e + \epsilon_{i,t} \quad (16)$$

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 (from 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), and use the Ljung-Box test for autocorrelation with 4 lags. ³² Following a variant of the backward stepwise regression procedure, we successively drop the non-significant variables and re-estimate (16) until only significant variables remain present. We adopt a 5% confidence level for the whole econometric analysis.

The general rule (16) allows for the selection of equilibrium cycles up to period three, which is the highest periodicity consistent with strong E-stability in our treatments (see Table 4). However, in light of the theoretical predictions discussed in Section 2 and the experimental data

²⁹ The p-value of the two-sided KS-test is 0.9607.

³⁰ However, differences in earnings efficiency ratios are not significant at 5% across ρ_2 values.

³¹ A similar exercise, whose results are reported in Appendix F, is performed on the individual savings choices. Admittedly, a link between microeconomic behaviors and equilibrium selection is more tedious to establish in the LtOE, as the learning literature has only been concerned with expectation formation. Yet, our main result below also applies to the LtOEs. In particular, we establish a connection between estimated individual savings strategies and aggregate experimental data as in the LtFE case, and we also document the use of simple heuristics when subjects are directly tasked with economic decisions.

³² This is consistent with the number of time periods observed at each data point.

Table 4

Distribution of forecasting rules among the 84 subjects in the LtFE that converge to the 2-cycle.

	2nd order adaptive rule $\hat{\beta}_{p_{t-1}} + \hat{\beta}_{p_{t-1}}^e = 1$ $\hat{\beta}_{p_{t-2}} = \hat{\beta}_{p_t}^e = \hat{\alpha} = 0$	naïve expectations $\hat{\beta}_{p_{t-2}} = \hat{\beta}_{p_t}^e = \hat{\beta}_{p_{t-1}}^e = \hat{\alpha} = 0$ $\hat{\beta}_{p_{t-1}} = 1$	stable AR(1) rule without constant $\hat{\beta}_{p_{t-1}}^e = \hat{\alpha} = 0$ $\hat{\beta}_{p_{t-1}} < 1$	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		4	1		1
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	4				2
Gp 2	5		1		
Gp 3	4	1		1	
Gp 4	2	2	1		1
$\rho_2 = 13.5$					
Gp 1	4				2
Gp 2	4	1	1		
Gp 3	5				1
Gp 4	6				
TOTAL	49 (58%)	21 (25%)	4 (5%)	1 (1%)	9 (11%)

Note: A stable AR(1) forecasting rule corresponds to $\hat{\beta}_{p_{t-1}} < 1$ and $\hat{\beta}_{p_{t-2}} = \hat{\beta}_{p_t}^e = \hat{\beta}_{p_{t-1}}^e = 0$.

highlighted in Section 4.1, we shall focus on the second-order adaptive rule (Equation (9)) and naïve expectations as our two benchmark rules. Those two rules correspond to special cases of the general rule (16). A subject is said to use a second-order adaptive rule if he uses a forecasting heuristic of the form:

$$p_{i,t+1}^e = \beta P_{t-1} + (1 - \beta) p_{i,t-1}^e + \epsilon_{i,t}, \beta \in (0, 1)$$

which corresponds to the following constraints on the estimated coefficients result from the econometric estimation of (16): $\hat{\beta}_{p_{t-1}} + \hat{\beta}_{p_{t-1}}^e = 1$, $\hat{\beta}_{p_{t-1}}, \hat{\beta}_{p_{t-1}}^e \in [0, 1]$, and $\hat{\beta}_{p_{t-2}} = \hat{\beta}_{p_t}^e = \hat{\alpha} = 0$.

In the special case of $\hat{\beta}_{p_{t-1}} = 1$ and $\hat{\beta}_{p_{t-2}} = \hat{\beta}_{p_t}^e = \hat{\beta}_{p_{t-1}}^e = \hat{\alpha} = 0$, the subject's heuristic takes the form:

$$p_{i,t+1}^e = P_{t-1} + \epsilon_{i,t}$$

which corresponds to naïve expectations. This forecasting heuristic selects the 2-cycle only in case of $\rho_2 = 5$.

We first estimate the general rule (16) in the 14 economies that converge towards the 2-cycle, for a total of 84 subjects. Table 4 classifies these subjects according to the learning rule that best

matches their behavior.³³ As expected, coefficients on information of lag 2 are never significant, and more than 80% of the subjects use a second-order adaptive rule or naïve expectations.

More than half of the subjects (49 out of 84) follow a second-order adaptive rule, and we fail to 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 4 of them. Almost all of these 49 subjects are part of economies with $\rho_2 = 8, 12$ or 13.5 , which is consistent with the theoretical learning predictions and the observed selection of the 2-cycle. For instance, all the subjects in Group 4 with $\rho_2 = 13.5$ have *second-order adaptive expectations*, which deliver convergence to the 2-cycle.³⁴

More can be said using the estimates of the coefficients in these second-order adaptive rules. Recall the stability condition of the 2-cycle under the adaptive rule (9) given in Subsection 2.4: the stability of the 2-cycle depends on the weight given to P_{t-1} (w in Equation (9)). Intuitively, the 2-cycle is stable if the weight on the past observed price is not too high, and this weight should be lower, the more unstable the 2-cycle. Our estimates appear to support this theoretical prediction: treatments with higher ρ_2 values tend to be associated with lower weights on prior price observations. Furthermore, the average weighting coefficient values for each ρ_2 treatment are always lower than the corresponding stability threshold $\underline{w}$.³⁵

A further quarter of the subjects (21 out of 84) use naïve expectations, with most being part of economies with $\rho_2 = 5$. Recall that, at this ρ_2 value, the 2-cycle is stable under naïve expectations (in fact, it is stable for all values of w of the adaptive rule (9)). For instance, all six subjects in Group 1 with $\rho_2 = 5$ use naïve expectations.

The above discussion is summarized in Fig. 8a, which plots the estimated coefficients of P_{t-1}^e against those of P_{t-1} for each of the 84 subjects. Most of the points are scattered around the dashed line $y = 1 - x$, which corresponds to the second-order adaptive rule (i.e., $\hat{\beta}_{p_{t-1}} + \hat{\beta}_{p_{t-1}^e} = 1$). We also observe a concentration of points around $(1, 0)$ for the case $\rho_2 = 5$ (indicating naïve expectations), while higher values of ρ_2 are associated with a more even dispersion along the line (which is line with the stability conditions discussed above).

The remaining 14 subjects (less than 20% of the total from economies that converge to the 2-cycle) use a mixed forecasting rule. Some of these subjects are part of Group 1 with $\rho_2 = 12$ and 13.5 , where we observe a particularly long transition, with irregular price movements, before convergence to the 2-cycle. This description leads us to the following result:

Finding 6 (*Individual behavioral rules and group behaviors*). *The individual estimates of forecasting rules reveals the wide use of second-order heuristics and, therefore, bridge the gap between theoretical predictions of individual learning behaviors and experimental data on aggregate price patterns.*

Despite our results, evidence for second-order forecasting heuristics is sparse in the related literature. While several studies have documented the use of first-order adaptive rules, especially

³³ The whole distributions of the parameter estimates are reported in Appendix E.

³⁴ This group, where all subjects use the same forecasting rules, is also the one that displays the fastest convergence toward the 2-cycle in this treatment. The point that the homogeneity of forecasting rules within a group impacts coordination on outcomes has been made by Marimon et al. (1993), who estimate similar forecasting rules in their LtFE.

³⁵ More precisely, the average weighting coefficients associated with P_{t-1} are 0.6603 in Groups 1, 2 and 3 with $\rho_2 = 8$ ($\underline{w} \simeq 0.8$); 0.6016 across all groups with $\rho_2 = 12$ ($\underline{w} \simeq 0.61$); and 0.5583 when $\rho_2 = 13.5$ ($\underline{w} \simeq 0.57$). K-S 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 when $\rho_2 = 8$, but other pair-differences across treatments are not significant at 5%.

Fig. 8. Outcomes of the estimations of individual forecasting rules. Note: Left panel: Scatter plot (84 observations) of the estimated coefficients $\hat{\beta}_{p_{t-1}}$ and $\hat{\beta}_{p_{t-1}}^e$ in Rule (16) for the subjects in 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 and 3 with $\rho_2 = 5$ and $\rho_2 = 8$). 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 (36 observations) of the distance between the steady-state value of average savings and the long-run savings equilibrium M/P^* implied by (16), for participants to the economies which converge to the steady state (i.e. all groups with $\rho_2 = 3$, and Group 4 for $\rho_2 = 5$ and $\rho_2 = 8$).

in linear frameworks involving one-period-ahead forecasts (Heemeijer et al., 2009; Hommes, 2011), the only study that finds evidence for second-order adaptive rules is Marimon et al. (1993) (and only with a much smaller sample of subjects and periods). By contrast, in a related framework to ours, Marimon and Sunder (1995) do not find strong evidence for the use of adaptive expectations, and Bernasconi and Kirchkamp (2000) instead highlight the possibility of *inertial expectations*, characterized by significant and positive intercepts in the estimated behavioral rules.

The estimation of rule (16) is less meaningful for the 36 remaining subjects — who select the steady state — as their predictions quickly become more or less constant over time. With the exception of Group 4 with $\rho_2 = 8$ (where small oscillations persist throughout the experiment), we find after only ten periods that 98% of price predictions fall in the range $[P^* - 1, P^* + 1]$, and 84% fall in the range $[P^* - 0.1, P^* + 0.1]$ (where P^* is the steady-state price). Therefore, for each of these 36 subjects, we compute the long-run estimated price level of Equation (16) to be

$$P^{**} \equiv \frac{\alpha}{1 - \hat{\beta}_{p_{t-1}} - \hat{\beta}_{p_{t-1}}^e - \hat{\beta}_{p_{t-2}} - \hat{\beta}_{p_t}^e} \quad (17)$$

The distribution of distances of the implied long-run savings equilibria from the steady-state level is reported in Fig. 8b. Based on these data, we do not reject the null hypothesis that the long-run savings level is equal to that in the steady state.³⁶

To conclude so far, we have established the systematic selection of simple equilibria both in the LtF and the LtO conditions in this complex lab environment, and consistency between

³⁶ The average relative distance from the estimated long-run savings equilibria to the steady state equals 0.003, and the p-value of a two-sided Wilcoxon signed rank test is 0.1252.

the observed prevailing equilibria, the individual behaviors of participants and the theoretical learning predictions. Before the conclusive remarks, the next section presents the outcomes from two additional sets of experiments designed to test the robustness of our findings and provide further explanations of our results.

6. Robustness checks

This section presents two additional robustness sessions. Subsection 6.1 discusses a set of additional experiments that aim to shed further light on the absence of 2-cycle selection in the LtOE in contrast to the LtFE, while Subsection 6.2 reports on a two-population version of our experiment that preserves the original overlapping-generation metaphor of the model.

6.1. LtOEs with 2-cycle training phase

In order to further test the robustness of the two-cycle, we ran 9 additional LtO sessions with an initial 10-period ‘training phase’ during which individual subjects played against five ‘robots’ that were perfectly coordinated on the 2-cycle (see Marimon et al. 1993; Duffy and Fisher 2005; Arifovic et al. 2014 for similar designs). We also ran 7 extra training sessions with a non-linear savings transformation in order to make the gap between the savings points of the 2-cycle less extreme in the case of $\rho_2 = 8$. Fig. 9 displays the aggregate price series and Table 5 summarizes the quantitative measures of these experiments in the same way as for the baseline experimental sessions in Section 4, while the exhaustive data for each group are deferred to Appendix C.

In total, out of the 16 robustness economies that we ran with these modifications, we observed three cases of up-and-down oscillations in subjects’ savings decisions in phase with the 2-cycle (with strong negative autocorrelation, see Table 5) induced by the training phase. Despite up-and-down oscillations, Gp. 5 with $\rho_2 = 5$ ends up in a closer neighborhood to the steady state than the 2-cycle (the corresponding ARDE value is 17.28%). As for the two others, namely Gp. 1 with $\rho_2 = 5$ and Gp. 5 with $\rho_2 = 8$ and the transformed savings values, the distance to the equilibrium 2-cycle values remained substantial, in contrast to the patterns observed in the LtFE, as shown by the much higher ARDE values reported in Table 5 (39.1% and even 55.2%). The other 13 groups converge towards the steady state, with no substantial difference with respect to the baseline experimental sessions. We therefore conclude that Finding 3 is a robust result, in that even groups in the LtOE who were ‘trained’ to follow the 2-cycle did not converge to it.

Finding 7 (*Robustness of the selection of the steady state in the LtOEs*). Almost all LtOEs converge to the steady state; the few instances of up-and-down oscillations in phase with the two-cycle induced by a training phase are hard to sustain and remain further away from the two-cycle than in the LtFE.

These robustness sessions also allow us to establish a plausible explanation for the absence of 2-cycle selection in the LtOE: namely, *strategic uncertainty*. We determined from post-session questionnaires that subjects exposed to the large variations in return on savings during the first ten training periods were concerned about volatility; statements regarding motivations included the desire to ‘secure a smooth payoff’, ‘hold on to an equilibrium situation’, ‘have a sure payment’, and ‘avoid fluctuations’.

The payoff schemes for these experiments were designed to generate the same payoff value in every period in the steady state equilibrium as for the *two-period average* along the 2-cycle.

Fig. 9. Real money balances time series in LtOEs with training. Note: See Fig. 3. The first 10 training periods are not displayed. Group 1 with $\rho_2 = 5$ prematurely crashed in period 58 due to a server issue.

However, the two payoffs are no longer equivalent once strategic uncertainty is taken into account. Strategic uncertainty is inherent to any group experiments requiring coordination between six players. Even though the experimental environment is deterministic, uncertainty arises from others' actions: coordination on the steady state may be preferable because strategic uncertainty weighs heavier along the 2-cycle in the LtO treatment.

This is a consequence of the strong asymmetry of payoff along the 2-cycle in the LtO treatment (about 1300 and 40 points over 2 periods) that is absent from the steady state allocation in the LtOE (that provides about 700 points per period) and from any perfect-foresight equilibrium payoff in the LtF treatment, including the 2-cycle (that provides 1300 points per period). If one subject deviates from the low saving decision on the 2-cycle, the high payoff (1300) is not realized and the average pay-off over the two periods is much lower than after a one-time deviation from the steady state. To see this clearly, one should look at the payoff table (see Appendix H): if the next period's return on savings is lower than its two-cycle value (i.e. 9.9), the payoff for any saving decision quickly drops below 1300 points. The effect of strategic uncertainty is even exacerbated at the 2-cycle if one recalls that subjects tend to make more mistakes due to con-

Table 5
Summary statistics of the 16 LtO experimental economies with training.

Group	1	2	3	4	5	1	2	3	4	
With training phase										
Equilibrium	$\rho_2 = 5$					$\rho_2 = 8$				
	2-cycle	SS	SS	SS	SS	SS	SS	SS	SS	
	ARDE	39.1	8.39	5.56	5.97	17.28	4.32	5.58	8.43	2.51
	ρ_s	−0.77	−0.38	0.27	−0.05	−0.84	0.31	0.03	0.45	0.43
	$\sigma(s_i)/\mu_{s_i}$	20.9	21.9	9.85	5.13	6.6	2.51	7.8	13.7	5.95
	EER	91	91.6	91.5	92.8	93.3	88.5	85.8	89.4	88.5
Group	1	2	3	4	5	6	7			
$\rho_2 = 8$: with training phase and non-linear transformation of savings										
Equilibrium	SS	SS	SS	SS	2-cycle	SS	SS			
ARDE	13.14	4.78	22.8	6.4	55.2	2.56	22.55			
ρ_s	0.67	−0.1	−0.65	0.02	−0.62	0.72	0.72			
$\sigma(s_i)/\mu_{s_i}$	26.7	12.98	56.85	10.47	61.66	7.6	48.43			
EER	81.1	87.8	69	84.9	74.6	85.7	76.6			

Notes: See Table 3.

fusion or mistakes when entering successively a high and a low number than a constant steady state decision, as explained in Section 4.2.

By contrast, the steady state in the LtOE provides a payoff structure comparable to the 2-cycle in the LtFE, in the sense that the payoff stream is constant, and potential losses due to one-time individual deviations are limited. As a result, subjects may prefer an allocation for which the payoff does not fluctuate (namely coordination on the steady state in the LtOE or any perfect-foresight equilibrium in the LtFE) to one that does and involves potentially larger payoff losses.

Two additional elements may play a role in the prevalence of the monetary steady state in the LtOE: a framing effect and subjects’ cognitive load. First, cautious or conservative behavior may appear more natural when it comes to making savings decisions than when making forecasts and tracking a time series pattern. Relatively stable savings decisions from one period to the next drive the dynamics towards the steady state by quickly pushing the return on savings towards unity.

Second, the cognitive load implied by the two experimental tasks is different. We report a significantly higher cognitive load in the LtO than in the LtF design according to two measures: the cumulative distribution of individual decision times (Fig. 10a) and the length of the instructions (Fig. 10b). Subjects read the instructions, complete the quiz and make their decisions more quickly in the LtFE despite the more complicated equilibrium on which they often coordinate (i.e., the 2-cycle).³⁷ This salient difference between the two designs constitutes a serious candidate for an explanation of the absence of selection of the 2-cycle in the LtOE. Our result may suggest that the more sophisticated the experimental task and the higher the implied cognitive load, the simpler the subjects’ behavioral rules and the simpler the selected equilibrium.

Before concluding, we take a look at the robustness of our results to another variation of the experimental design that is commonly used across experimental and learning OLG studies.

³⁷ The 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 0.0001 in the two cases. Note that the instructions are slightly longer (by around half a page) in the LtOE, but this factor alone likely cannot account for the additional 9 minutes.

Fig. 10. Measurement of cognitive loads in the 20 LtFE vs. the 28 LtOEs. Note: The 12 baseline LtO groups and the 16 robustness groups with training are pooled together on the figures.

6.2. Two-population design

We present here another second set of experiments to test whether our results established in a single-population design are robust to a setup that preserves the overlapping generation metaphor of the underlying model and genuinely involves two distinct populations of alternatively young and old agents. One reason for this investigation is the common use of such a two-population design in related OLG experimental studies (see, *inter alia*, Aliprantis and Plott 1992; Marimon and Sunder 1993) as well as in the GA learning model of Bullard and Duffy (1998). This latter study shares two interesting similarities with our group experiment: it only results in the selection of the steady state or the two-cycle, no matter the ρ_2 -value, and it implements an heterogeneous-agent version of the OLG model, instead of the aggregate temporary equilibrium mapping between expected and realized prices that is used in the adaptive learning literature.

For these reasons, we ran robustness LtF and LtO sessions using a two-population version of our experimental design. Instead of $N = 6$ subjects who make forecasts/savings decision in every period, we introduce a design where two groups of six subjects each make decisions every other period. Thus, one group makes decisions in every *odd* period, and the other every *even* period. Note that population sizes of six ensure that the market influence of each subject remains the same, and we doubled the exchange rate to ensure comparable individual earnings between the two designs. The complete instructions are given in Appendices J and K.

We ran four two-population groups of two representative treatments of the LtFE, namely $\rho_2 = 5$ that gives rise to an unstable steady state or a stable 2-cycle and $\rho_2 = 13.5$, that displays the richest set of potential equilibria. We also conducted four LtO sessions for $\rho_2 = 5$. Given the extreme payoff values along the 2-cycle in the LtOE (see Table 2), we added four more LtO groups with a *flattened payoff* (see Appendix K) so as to render the payoff between the odd and the even generations less asymmetric, and provide closer incentives to the two populations of subjects to coordinate on the 2-cycle.

Fig. 11 presents the behavior of real money balances in all those groups and Table 6 shows the same statistics as for the baseline treatments (for all group figures, see Appendix D).

Fig. 11. Real money balances time series in the two-generation experimental economies. Note: See Fig. 3.

Our results in the baseline treatments are fully robust to the two-population implementation:

Finding 8 (*Robustness to the two-population design*).

- (a) *In all but one experimental economies, the price level converges to a perfect foresight equilibrium.*
- (b) *In all but one LtF economies, subjects coordinate on the 2-cycle.*
- (c) *In all LtO economies, the price, or equivalently aggregate savings, converge towards or in a neighborhood of the steady state.*

Taking first a closer look at the two-population LtFEs, except the third group of $\rho_2 = 13.5$, all groups converge to the 2-cycle equilibrium values. Comparing with Table 3, both convergence of the price level (through the ARDE values) and coordination between subjects' forecasts (through the relative standard deviation values) are even better in the two-population design than in the one-population design.

The case of $\rho_2 = 5$ is particularly clear: the ARDE values are all below 5%, three out of four are below 1%, and all the relative standard deviations are on average close to zero. The three converging groups with $\rho_2 = 13.5$ display similar high levels of convergence and coordination as in the single-population design, but coordination seems quicker: we do not observe any long chaotic transient as in the single-population design (compare the first 50 periods in Figs. 11b and 6a).

By contrast, the third group fails to coordinate on any perfect-foresight equilibrium: the time series (green line of Fig. 11b) does not display any clear pattern, the ARDE value is particularly high (almost 37%), while the autocorrelation is negative (-0.43), but not nearly close to -1 , subjects are poorly coordinated (the relative standard deviation is the highest observed across all sessions, as high as 60.4%) and the average price value over the last 25 periods (20.92) is almost twice lower than the one that would prevail along the 2-cycle (37.27).

Digging into individual data indicates that this coordination failure is due to a single subject who displayed confusion, entered random forecasts (reaching an earning efficiency ratio lower than 8%), and hindered the group coordination. This example of coordination failure shows that coordination between a group of 6 or even 12 subjects in such a complex environment is far from trivial but yet achieved in all but one of our 64 experimental groups (over all treatments and designs), which is quite remarkable.

Lastly, the two-population LtOEs tend to coordinate on the steady state: the ARDE values are of the same order of magnitude as in the one-population design, and 5 out of the 8 groups display ARDE values below 10%. In the three other groups (namely Group 3 of the same payoff and Groups 1 and 3 of the flattened payoff), the ARDE values are higher (20.3 to 27.9%, see Table 6), but additional computations reveal that the average aggregate savings over the last 25 periods are respectively equal to 0.535, 0.564 and 0.497, which corresponds to a small neighborhood (less than 10%) of the steady state value of 0.538. A closer look at Figs. 11c and 11d, together with those numbers, reveal that these groups converge near the steady state but with clear up-and-down oscillations akin to a dampened or attenuated two-cycle (to see that, look at the strong negative autocorrelation values of aggregate savings in Table 6, a feature absent from the one-population design in Table 3).

End-questionnaire data tell us that this feature is due to participants competing with the other generation for the higher returns on savings by lowering their own savings decisions to enjoy a higher-than-one return.³⁸ Interestingly, they coordinated on lower-than-steady-state savings, but not quite as low as coordination on the 2-cycle would have required. This discrepancy seems to have two explanations. First, subjects reported concerns about utility losses when attempting to bring the return down by savings very little. Second, participants reported concerns about the

³⁸ More than one third of the participants to those three groups made statements when describing their strategies such as: 'drive supply to around 40 for our team and keep it there. Keep previous period's return around 1.5 as most people use it to make decisions. This way, the other team gets return around 0.6, save 70 and still get around 550, no too little so it does not motivate trade war'; 'most points when return high, so try to keep my savings decision low, but not too low not too loose too many points.'; 'I wanted to get the price high in the round period so I advised to save little but not too little otherwise my utility would be lower'; 'try to pull down the average savings of my team'.

Table 6
Summary statistics of the experimental economies in the two-group design.

Group	1	2	3	4	1	2	3	4
LtFE								
$\rho_2 = 5$				$\rho_2 = 13.5$				
<i>Equilibrium</i>	2-cycle	2-cycle	2-cycle	2-cycle	2-cycle	2-cycle	2-cycle	2-cycle
<i>ARDE</i>	0.72	0.06	0.16	2.16	2.33	0.2	36.91	0.13
$\sigma(s_i)/\mu_{s_i}$	0.78	0.00	0.06	1.32	5.22	2.61	60.4	1.13
ρ_s	−0.96	−0.96	−0.96	−0.95	−0.91	−0.96	−0.43	−0.96
<i>EER</i>	89.7	90.7	91.4	91.9	81.6	83.8	55.7	91.2
LtOE ($\rho_2 = 5$)								
Same payoff				Transformed payoff				
<i>Equilibrium</i>	SS	SS	SS	SS	SS	SS	SS	SS
<i>ARDE</i>	5.42	5.78	27.9	6.7	26.3	3.07	20.3	7.58
$\sigma(s_i)/\mu_{s_i}$	10.66	6.37	13.39	7.04	8.67	9.52	17.26	9.01
ρ_s	−0.08	−0.9	−0.89	−0.88	−0.77	−0.75	−0.93	−0.81
<i>EER</i>	97	96	93.9	97.4	98.8	99	99.1	99.5

Notes: See Table 3.

other generation, who would then receive a much lower payoff than theirs, and would therefore have an incentive to reply by lowering their savings and driving the return towards one. Our effort, through a flattened payoff table, to provide a more generous payoff to the high savings/low returns generation did not eliminate this inter-generational concern, and has led to similar results as for the other LtO groups.

To conclude, our results are fully robust to a design that preserves the overlapping generation friction as common in related experimental and learning studies. Again, remaining cautious given the few data points that a group experiment allows us to collect, it even seems that the two-population design favors coordination on a 2-cycle more than the one-population design does.

7. Conclusion

This experimental study adds to the literature on equilibrium selection when self-fulfilling beliefs lead to indeterminacy in the model, and additionally provides an empirical test of learning predictions when multiple equilibria are possible. We design an experiment in an heterogeneous-agent version of the well-known complex OLG environment first studied by Grandmont (1985). This environment exhibits infinitely many periodic — and even chaotic — equilibria, along with the monetary steady state. Existing theoretical contributions provide little guidance into the process of equilibrium selection because any of these equilibria can emerge as a stable outcome of a suitable expectation formation process. Hence, the equilibrium selection problem requires empirical insights.

We can broadly summarize our experimental results as follows. In all (but one) experimental economies, prices converge to a neighborhood of a simple perfect-foresight equilibrium. This equilibrium is either a steady state or a two-cycle. Our learning-to-forecast experiments are the first example of spontaneous coordination on a two-cycle in the lab. By contrast, learning-to-optimize experiments do not converge to a two-cycle, not even after up-and-down oscillations induced by an initial training phase, but rather converge to a steady state. This may be due to strategic uncertainty or differences in cognitive load between the two designs. Our results are ro-

bust to an alternative, two-population, design that preserves the overlapping generation friction in the lab.

Comparing the experimental data to theoretical predictions of forecasting rules, our experiment provides evidence on behaviors in a complex environment, which has not yet been extensively investigated in a laboratory setting. We find that subjects use simple belief-formation processes by tracking low-order patterns and only considering very recent observations. Simple behavioral rules, together with repeated market interactions, enforce convergence to simple equilibria even though the set of possible outcomes is large and complicated. To revisit a quote from Lucas (1986, pp. S424–S425) first shown in the introduction, this empirical result supports the idea that ‘*all but a few equilibria are [...] behaviorally uninteresting*’. If perfect foresight (or rational expectations) is not selective enough, our experimental data suggest that the selected equilibrium will be robust to the use of simple behavioral rules. Furthermore, once removing the assumption of the conditional optimality of savings decisions, our experiment reports an even smaller set of empirically relevant outcomes, of which the steady state becomes the most likely. While we cannot yet claim that indeterminacy is only a theoretical issue, our experiment provides an empirical example where the problem of equilibrium selection is far less salient than theory would suggest.

Our experimental environment is, of course, a stylized one. We may consider it too rudimentary to draw general conclusions about economic dynamics or to validate the claim that the long-run behavior of a competitive monetary economy must be a deterministic steady state. Instead, we provide suggestive evidence that self-fulfilling beliefs alone cannot sustain irregular and unpredictable fluctuations, even in a complex environment with (initially) heterogeneous beliefs and decentralized information. Empirically observed economic fluctuations likely require amplification mechanisms that cannot be accounted for by beliefs alone.

Appendix A. Experimental economies – LtFE (baseline sessions) (Figs. 12–16)

Fig. 12. $\rho_2 = 3$ (steady state).

Fig. 13. $\rho_2 = 5$ (2-cycle). Fig. 14. $\rho_2 = 8$ (noisy 32-cycle). Note: For Group 4, note the change in the scale of y-axis.

Fig. 15. $\rho_2 = 12$ (chaotic dynamics). Fig. 16. $\rho_2 = 13.5$ (3-cycle).

Appendix B. Experimental economies – LtOE (baseline sessions) (Figs. 17–19)

Fig. 17. $\rho_2 = 3$ (steady state).

Fig. 18. $\rho_2 = 5$ (2-cycle).

Fig. 19. $\rho_2 = 8$ (noisy 32-cycle).

Appendix C. Experimental economies – LtOE with training (robustness sessions)

C.1. LtOE with a 10-period initial training phase on the 2-cycle (Figs. 20, 21)

Fig. 20. $\rho_2 = 5$ (2-cycle) **with training**. Note: a computer crashed at period 58 in Group 1, ending prematurely the experiment in this group.

Fig. 21. $\rho_2 = 8$ (noisy 32-cycle) **with training**.

C.2. LtOE with a 10-period initial training phase on the 2-cycle and a non-linear transformation of savings ($\rho_2 = 8$) (Fig. 22)

Fig. 22. $\rho_2 = 8$ (noisy 32-cycle) **with training and a non-linear transformation of savings** $\in [0, 100]$.

Appendix D. Experimental economies – Two-population design (robustness sessions)

D.1. LtFE with the two-population design (Figs. 23, 24)

Fig. 23. $\rho_2 = 5$ (2-cycle).

Fig. 24. $\rho_2 = 13.5$ (3-cycle).

D.2. LtOE with the two-population design (Figs. 25, 26)

Fig. 25. $\rho_2 = 5$ (2-cycle) LtOE in the two-population design.

Fig. 26. $\rho_2 = 5$ (2-cycle) LtOE in the two-population design with a flattened payoff function.

Appendix E. Analysis of individual forecast time series (Fig. 28)

Fig. 27. Sample average and first-order sample autocorrelation of individual price forecasts in the 20 baseline LtFE, 120 observations. *Notes:* The first 10 periods are discarded. In red: convergence to the steady state, in blue: 2-cycle dynamics (see classification in Table 3).

Fig. 28. Cumulative distribution of the estimated coefficients in Equation (16) for the 14 baseline LtF economies that converge to a 2-cycle, 84 observations).

Fig. 29. Sample average and first-order autocorrelation of individual savings decisions in LtOEs, 168 observations. *Notes:* See Fig. 27.

Appendix F. Analysis of individual savings time series (LtOE)

Fig. 29 summarizes the descriptive statistics of the individual savings time series in the baseline 12 experimental groups as well as the 16 robustness groups that feature an initial 10-period training phase (referred to as Tr. T, and S for sessions using also the non-linear transformation of savings values). We have to include those robustness sessions here to compare the behavior of the subjects around the steady state and along the two-cycle, as the baseline LtOEs all converge to the steady state.

Using the same econometric procedure as in the LtFE (see Section 5), we estimate behavioral rules from the individual savings decisions in the LtOE. Specifically, we estimate the following behavioral rule for each participant³⁹:

$$s_{i,t} = \alpha + \beta_{s_{t-1}} s_{i,t-1} + \beta_{s_{t-2}} s_{i,t-2} + \beta_{R_{t-1}} R_{t-1} + \varepsilon_{i,t} \quad (18)$$

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 (18) 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 $|\hat{\beta}_{s_{t-1}}| \in (0, 1)$, as well as a stable AR(2) rule if $\hat{\beta}_{R_{t-1}} = 0$ and $|\hat{\beta}_{s_{t-2}} + \hat{\beta}_{s_{t-1}}| \in (0, 1)$.

Following those nested cases, Table 7 classifies the 168 subjects, 150 in LtOEs that converge or oscillate around the steady state, and 18 in LtOEs that display regular up-and-down oscil-

³⁹ For instance, Bao et al. (2017) estimate a similar rule with an AR(1) structure for quantity decisions in an LtOE in an asset pricing model.

Fig. 30. Cumulative distribution of the estimated coefficients in Equation (18), 168 observations.

lations in phase with the 2-cycle (see the discussion in Section 6.1, those three groups are in bold in Table 7).⁴⁰ Fig. 30 reports the distributions of the estimates of the four parameters of the general rule for all the 168 subjects. The main insights from this exercise may be presented as follows.

First, we notice that the intercept is always significant, and for all but one subject is positive (see the first plot in Fig. 30). A high concentration of the estimated values is observed between 0.5 and 0.6, which broadly corresponds to the steady state value of individual savings.

Next, we focus on the 150 estimates in economies that converge towards the steady state (i.e. we exclude Gps. 1 and 5 of Tr. T with $\rho_2 = 5$ and Gp. 5 of Tr. S): 94 subjects (63%) are characterized by an AR rule with significant intercept ($\hat{\beta}_{R_{i,t-1}} = 0$), 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 (18) (where $\hat{\beta}_{R_{i,t-1}} = 0$) as $s^{*SS} \equiv \frac{\alpha}{1 - \hat{\beta}_{S_{i,t-1}} - \hat{\beta}_{S_{i,t-2}}}$. Fig. 31a below displays the frequency distribution of those relative estimated distances to the steady state. The average relative distance is -0.01 , and we cannot reject the null hypothesis that it is equal to zero.⁴¹

⁴⁰ Since savings decisions are more variable than price predictions in economies that converge to the steady state, the estimation of (18) is less problematic in the LTOEs than in the LTFEs. In the LTOEs, only 3 subjects have strictly constant savings decisions after the first 10 periods.

⁴¹ A two-sided Wilcoxon rank sum test gives a p-value of 0.6481.

Table 7

Distribution of savings rules among the 168 subjects in the LtOE in the 12 baseline LtO sessions and the 16 robustness sessions with the initial training phase.

		<i>intercept only</i> $\hat{\beta}_{s_{t-1}} = \hat{\beta}_{s_{t-2}} = 0$	<i>stable AR(1) rule</i> $\hat{\beta}_{s_{t-2}} = 0, \hat{\beta}_{s_{t-1}} \in]0, 1[$ $\hat{\beta}_{R_{t-1}} = 0$	<i>stable AR(2) rule</i> $\hat{\beta}_{s_{t-2}} \neq 0, \hat{\beta}_{s_{t-1}} \in [0, 1[$	<i>mixed rule</i> $ \hat{\beta}_{s_{t-2}} , \hat{\beta}_{s_{t-1}} \in [0, 1[$ $\hat{\beta}_{R_{t-1}} > 0 \quad \hat{\beta}_{R_{t-1}} < 0$	
$\rho_2 = 3$	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	2		2	1	1
	Gp 2	2	1	1	1	1
	Gp 3	3	1		1	1
	Gp 4	1	1			4
$\rho_2 = 8$	Gp 1	1			2	3
	Gp 2	2	2		1	1
	Gp 3	2	1		1	2
	Gp 4	3		1		2
<i>Sessions with a 10-period initial training phase (Tr. T)</i>						
$\rho_2 = 5$	Gp 1		4	2		
	Gp 2	1	2	2		1
	Gp 3	1	2	2	1	
	Gp 4	1	3	1		1
	Gp 5		1	5		
$\rho_2 = 8$	Gp 1		3	2		1
	Gp 2	1	3	2		
	Gp 3	2	1			3
	Gp 4	1	2	2		1
<i>Sessions with a 10-period initial training phase and a non-linear transformation of savings (Tr. S)</i>						
$\rho_2 = 8$	Gp 1	1	2	1	1	1
	Gp 2		1		2	3
	Gp 3			2	1	3
	Gp 4	1	2			3
	Gp 5			6		
	Gp 6	1	1	2	1	1
	Gp 7		1	2	1	2
TOTAL		34 (20%)	38 (23%)	40 (24%)	17 (10%)	39 (23%)

Notes: The groups that display up-and-down oscillations in phase with the 2-cycle are highlighted in bold.

Note that a comparison between Fig. 31a and Fig. 8b displaying the distribution of the same distances obtained from the estimated forecasting rules in the LtFE reveals a better coordination on the steady state in the LtFEs than in the LtOEs – i.e. the estimated distances are strikingly smaller and closer to zero in the LtFEs. This is fully in tune with the discussion in Section 4.2.

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

Fig. 31. Estimated long-run savings values with respect to equilibrium. *Notes:* Left panel: Frequency distribution of the relative distances between the steady state value of savings and the long-run estimated saving values s^{*SS} for the 94 subjects that are characterized by an AR(1) or AR(2) rule in groups converging towards the steady state. Right panel: Frequency distribution of the distances between the average of the 2-cycle values of savings and the estimated long-run saving values s^{*2} in (18) for the economies that display up-and-down oscillations in phase with the 2-cycle, 18 subjects.

negative.⁴² Those 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 optimal savings decision (bottom left corner of the payoff table), as long as the expected return is not too small.

In the three economies in which the dynamics resemble an ‘attenuated’ 2-cycle, all 18 participants 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.⁴³ These signs are fully consistent with a two-cycle dynamics – see Hommes et al. (2013) for a discussion of AR rules associated with 2-cycle dynamics.

Finally, we compute the average relative distances of the long-run estimated savings equilibrium s^{*2} to the average value of savings along the 2-cycle for each of the three economies (see Fig. 31b). This distance is very small for two of three groups, i.e. -0.0188 in Group 1 with $\rho_2 = 5/\text{Tr. T}$ and -0.018 in Group 5 with $\rho_2 = 8/\text{Tr. S}$. It is larger for Group 5 with $\rho_2 = 5/\text{Tr. T}$ (-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 equilibrium 2-cycle, and we observe the wider oscillations in Group 1 with $\rho_2 = 5/\text{Tr. T}$ and Group 5 with $\rho_2 = 8/\text{Tr. S}$.

Therefore, we conclude that *the individual estimates of savings rules reveal the wide use of simple heuristics that are consistent with the experimental aggregate savings patterns.*

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

⁴³ The average estimates of $\hat{\beta}_{s_{t-1}}$ and $\hat{\beta}_{s_{t-2}}$, respectively in Group 1 with $\rho_2 = 5/\text{Tr. T}$, in Group 5 with $\rho_2 = 5/\text{Tr. T}$ and in Group 5 with $\rho_2 = 8/\text{Tr. S}$ are -0.7172 , -0.3753 and -0.3646 , and 0.1646 , 0.4777 , and 0.5135 .

Appendix G. 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:

consumption of chips when young = 200 – 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:

savings in money = number of chips saved $\times$ 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:

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		

Appendix H. 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 131.2. 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 $\rho_2 = 5$

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)

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
0.005	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	15	11	10
0.02	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	11	11	11	11	15	11
0.03	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	11	15	10
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

Pay-off table $\rho_2 = 8$

Appendix I. Computer interfaces (Figs. 32–33)

Fig. 32. Subjects' computer interface (LtFE).

Fig. 33. Subjects' computer interface (LtOE).

Appendix J. LtF instructions for even generation [odd generation]

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 a consumption good at a given price in each period. We will refer to this consumption good as chips. **You are a Professional Forecaster, and you have to make predictions of the price of a chip.**

In every period, two generations of individuals – the young and the old – trade chips. 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, you advise an individual of the generations that are young in even [odd] periods, namely in periods 2, 4, 6, etc. At the beginning of each even [odd] period, you have to predict the price of the chips in the next odd [even] period, and your

prediction is then used by the young individual for **making a savings decision in the current even [odd] period.**

There are six forecasters who make price predictions in even [odd] periods, and six other forecasters who make predictions for the generations that are young in odd [even] periods. Each forecaster is played by a participant like you. Once the individuals you are advising become old, they do not need your price forecasts, as they just consume the number of chips their savings can buy. Therefore, in odd [even] periods, you do not make predictions and you wait for the forecasters to the young generation in odd [even] periods to submit their price forecasts for the next even [odd] period.

In general, **the price predictions of participants for the next period** determine the number of chips young individuals are selling to the old ones *in the current period*, and therefore **the price of the chips in the current period**: the **higher the price forecasts 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 **price predictions for the next period only influences the price in the current period, not the price in the next period.**

In economies similar to this one, the price of chips has historically been between 1 and 100.

Information about your prediction task

The experiment lasts for 100 periods. **At the beginning of each even [odd] period, you have to submit a prediction of the price of the chips in the next odd [even] 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 pay-off depends on your forecast error, that is the difference between your price forecast and its realized value (we explain below how your pay-off is exactly computed). **You will then observe your forecast error and your corresponding pay-off for a forecast made at the beginning of any period at the end of the next period.**

The experiment starts in period 1, you have to wait for the other 6 forecasters to submit a prediction for the price in period 2. Once all 6 forecasters have submitted their price forecast, the young individuals they are advising 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. [The experiment starts in period 1, you are asked to submit a forecast for the next period (period 2). Once all 6 forecasters have submitted their price forecast, 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 are then entering period 2.]

You are then entering period 2. You are now asked to submit a price forecast for the next period (period 3). Once all 6 forecasters have done so, you will observe the realized price in period 2. You are then entering period 3, and have to wait for the other 6 participants to submit their forecasts for the price in period 4. Once they all have done so, you will observe the price in period 3, your forecast error based on the forecast that you made in period 2 for period 3, and your corresponding payoff for period 3. You are then entering period 4, and have to submit a price forecast for period 5, etc. This sequence of events takes place in each of the 100 periods of the experiment. [In period 2, you wait for the other 6 forecasters to submit a prediction for the price in period 3. Once they all have done so, you will observe the realized price in period 2, your forecast error based on the forecast that you made in period 1 for period 2, and your corresponding payoff for period 2. You are then entering period 3, and have to submit a price

forecast for period 4, etc. This sequence of events takes place in each of the 100 periods of the experiment.]

The computer interface is mainly self-explanatory. 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 past price forecasts from the beginning of the experiment on;
- Your pay-off from the beginning of the experiment on.

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.

Information about your pay-off

Your pay-off 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 pay-off will be displayed on the computer screen in terms of points, and is computed as follows:

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

There is a pay-off table with the instructions. It shows your pay-off 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 pay-off that you can get in any period. The more accurate your forecast, the lower your squared forecast error, and the higher your pay-off. If your forecast error is higher than 7, you get 0 point, and this is the minimum pay-off 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 pay-off 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 pay-off 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 pay-off in the experiment is computed, and converted into euro. **For each 1300 points you make, you earn 0.70 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 the economy enters period 7, are you waiting or submitting a price forecast?
...
 2. If you enter period 12 [13], for which period are you asked to submit a price forecast?
...
 3. If you enter a price prediction for period 11 [10], which period's price will be influenced by your prediction?
...
 4. 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?
...
 5. 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?
...
 6. 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?
...
 7. Suppose all forecasters like you are predicting at the beginning of period 12[13] a "high" price for period 13[14], would you say that:
 - (a) The price in period 13[14] is likely to be high;
 - (b) The price in period 13[14] is likely to be low;
 - (c) The price in period 12[13] is likely to be high;
 - (d) The price in period 12[13] is likely to be low;
 - (e) Forecasts of the price for period 13[14] do not influence the price in period 13[14];
 - (f) Forecasts of the price for period 13[14] do not influence the price in period 12[13].
- N.B.: **multiple answers are possible.**

Appendix K. LtO instructions for even generation [odd generation]

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 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 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. 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 than their 50 chips in the next period, when they will be old. In each period, a young individual then consumes:

consumption of chips when young = 200 – quantity of chips saved

You work for a Professional Saving Advisor Bureau, and you have to give saving advice.

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:

savings in money = number of chips saved $\times$ 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 the young individuals save, the lower the realized price of chips, and the more chips the old individuals can buy with their savings and consume.**

As a professional saving advisor, you make savings decisions for an individual of the generations who are young in even[odd] periods, namely in periods 2, 4, 6, [1, 3, 5] etc. You then have to make a saving decision at the beginning of each even[odd] period. There are six advisors like you who give saving advice in even [odd] periods, and six others who advise the generations that are young in odd [even] periods. Each advisor is played by a participant like you.

Once the individuals you advise have become old, they just consume the number of chips their savings can buy from the new young individuals. The consumption of chips of an old individual then equals:

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

Hence, old individuals do not need your savings advises, and in odd [even] periods, you do not give advice and have to wait for the other six advisors to the young generations in odd [even] periods to submit their savings advice.

In general, **your saving decision influences what the individual consumes both when young in a given even [odd] period, and when old in the next odd [even] 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**, 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 odd [even] period to guide your savings decision in the current even [odd] period.**

Information about your task as an advisor

The savings advisor bureau exists for 100 periods. Each individual lives for two periods, consumes and saves when young, and consumes when old. At the beginning of each even [odd] period, you have to submit a savings decision for a young individual. **Your pay-off depends on the consumption of chips of this individual both when young and when old** (we explain below how your pay-off 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 pay-off of your savings decision made in even [odd] periods only at the end of the next (odd) [(even)]period, when he will have become old.**

The experiment starts in period 1, you have to wait for the other 6 advisors to make savings decisions for the individuals in the generations that are young in odd periods. Once all 6 advisors have done so, the young individuals they are advising consume and save chips in period 1, 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. You are then entering period 2. You are now asked to submit a saving decision for this period (period 2). Once all six advisors have done so, the young individuals in period 2 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 are then entering period 3 and have to wait for the other six participants to submit their savings decisions, which will determine the price in period 3 and the return on savings between periods 2 and 3. You then observe the consumption of the young person you advised in period 2 both in period 2 (when young) and 3 (when old), and therefore the corresponding pay-off of your savings decision made in period 2. You then enter period 4 and are asked to submit a savings advise for a new young individual in this period, etc. [The experiment starts in period 1, and you are asked to submit a saving decision for this period. Once all six participants like you have entered their savings decision, all young people consume and save chips in period 1, 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. You are then entering period 2. You have to wait for the other six advisors to make savings decisions for the individuals of the young generation in period 2. Once they all have done so, those young individuals consume and save chips in period 2, 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 pay-off of your savings decision made in period 1. You are then entering period 3 and are asked to enter a savings advise for a new young individual in this period, etc.]

This sequence of events takes place in each of the 100 periods of the experiment.

The computer interface is mainly self-explanatory. 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 2 up to the previous period;
- The average savings decisions among the 6 advisors to your generations from the beginning of the experiment (period 2) [(period 1)] on;
- Your savings decisions from the beginning of the experiment on;
- The corresponding consumption of chips of the individuals you advised when young from the beginning of the experiment on;
- The consumption of chips when old of the individuals you advised when young from the beginning of the experiment on;
- Your pay-off from the beginning of the experiment on.

The two plots (left panel) indicate your savings decisions together with the average decisions of the advisors to your generations 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 even [odd] period, **the return on savings you need to forecast for the next odd [even] period and the savings decision you need to make for the current even [odd] 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 saving decision and when you have to wait.

Information about your pay-off

Your pay-off 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 pay-off. You do not need to calculate his utility, and hence your pay-off yourself. **There is a pay-off table on your desk.** 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 even [odd] period (horizontal axis) according to your forecast of the return on savings in the next odd [even] 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 131.2. 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 $10/20 = 0.5$, this person consumes $200 - 90 = 110$ when young, and $50 + 0.5 \times 90 = 95$ when old, and your payoff is 422 {661} points. For the same savings decision and current price, if the next period's price turns out to be 5, the return on savings is $10/5 = 2$ and this person consumes $50 + 2 \times 90 = 230$ when old, and your payoff is 630 {707} points.

The sum of your payoff from your savings advice over the different periods is shown in the bottom right of the screen. At the end of the experiment, your cumulative payoff over the experiment is computed, and converted into euro. **For each 1300 points you make, you earn 1 euro. 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 fill 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 eight questions before starting the experiment.

If you have any questions, please ask them now!

Questionnaire

1. If you enter period 16, are you waiting or submitting a savings advice?

...

2. If you enter period 6 [7], for which period do you need to forecast the return on savings to make your savings decision?

...

3. If you make a saving decision at the beginning of period 10, in which period will you observe your corresponding pay-off?

...

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

...

5. Suppose that in period 10 [9], you advised to save 4 chips, the price of the chips was 30 in this period, and 10 in the next period (period 11 [10]). What is the return on savings between period 10 [9] and period 11 [10]?

...

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

...

7. 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?

...

8. Suppose you have decided for a young individual to save 100 chips in a given period, would you say that:
- (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 periods, 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.

$\rho_2 = 5$, L/O-two population, flattened pay-off ($K = 154$, $C = 2.5$, $\alpha = 1$)

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	206	211	218	229	239	248	255	262	266	270	272	272	271	269	265	259	251	240	225	205	175	88
0.075	207	217	229	251	270	287	301	313	323	331	337	341	343	343	341	336	329	319	305	286	256	169
0.1	208	223	240	272	299	322	341	358	371	382	390	396	399	400	398	394	387	377	363	343	313	225
0.2	213	246	283	345	395	434	465	489	507	519	528	533	534	532	527	520	509	495	476	452	418	326
0.3	218	268	322	407	469	514	547	569	585	594	599	599	596	590	581	570	556	538	517	490	453	359
0.4	223	289	357	458	526	572	602	621	632	637	638	634	628	619	607	593	576	557	534	505	467	372
0.5	228	308	389	501	571	614	641	655	662	664	661	655	646	634	621	605	587	566	542	513	474	377
0.6	232	327	418	538	607	646	668	679	682	681	675	667	656	643	628	612	593	571	546	517	477	380
0.7	237	345	445	569	635	671	688	696	696	692	684	674	662	649	633	616	596	574	549	519	479	382
0.8	241	362	469	595	658	689	704	708	706	700	691	680	667	652	636	618	598	576	550	520	480	383
0.9	246	379	492	618	677	704	715	717	713	705	695	683	669	654	638	620	600	577	551	521	481	384
1	250	394	512	638	693	716	724	723	718	709	698	685	671	656	639	621	600	578	552	521	481	384
1.1	255	409	531	655	705	725	731	728	721	712	700	687	673	657	640	621	601	578	552	522	482	384
1.2	259	424	549	670	716	733	736	732	724	714	702	688	674	658	641	622	601	579	553	522	482	384
1.3	264	437	565	683	725	739	740	735	726	716	703	689	674	658	641	622	602	579	553	522	482	385
1.4	268	450	580	694	733	744	744	737	728	717	704	690	675	659	641	623	602	579	553	522	482	385
1.5	272	463	594	704	739	748	746	739	730	718	705	691	675	659	642	623	602	579	553	522	482	385
1.6	276	475	607	713	745	752	749	741	731	719	705	691	676	659	642	623	602	579	553	522	482	385
1.7	281	486	618	721	750	755	751	742	731	719	706	691	676	660	642	623	602	579	553	523	482	385
1.8	285	497	629	728	754	757	752	743	732	720	706	692	676	660	642	623	602	580	553	523	482	385
1.9	289	508	640	734	757	759	754	744	733	720	706	692	676	660	642	623	602	580	553	523	482	385
2	293	518	649	740	760	761	755	745	733	720	707	692	676	660	642	623	603	580	554	523	482	385
3	332	599	716	772	777	770	760	748	735	722	708	693	677	660	643	623	603	580	554	523	482	385
4	367	655	751	785	782	773	761	749	736	722	708	693	677	660	643	624	603	580	554	523	482	385
5	399	693	771	791	784	774	762	749	736	722	708	693	677	660	643	624	603	580	554	523	482	385
6	428	721	783	794	785	774	762	749	736	722	708	693	677	660	643	624	603	580	554	523	482	385
7	455	741	791	795	786	774	762	749	736	722	708	693	677	660	643	624	603	580	554	523	482	385
8	479	756	796	796	786	774	762	749	736	722	708	693	677	660	643	624	603	580	554	523	482	385
9	502	768	800	797	786	774	762	749	736	722	708	693	677	660	643	624	603	580	554	523	482	385
10	522	777	802	797	786	774	762	749	736	722	708	693	677	660	643	624	603	580	554	523	482	385
20	659	807	809	798	787	775	762	749	736	722	708	693	677	660	643	624	603	580	554	523	482	385

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

References

- Aliprantis, C.D., Plott, C.R., 1992. Competitive equilibria in overlapping generations experiments. *Econ. Theory* 2 (3), 389–426.
- Arifovic, J., Evans, G., Kostyshyna, O., 2014. Are Sunspots Learnable? An Experimental Investigation in a Simple General Equilibrium Model. Simon Fraser University. Manuscript.
- Azariadis, C., Guesnerie, R., 1986. Sunspots and cycles. *Rev. Econ. Stud.* 53 (5), 725–737.
- Bao, T., Duffy, J., Hommes, C., 2013. Learning, forecasting and optimizing: an experimental study. *Eur. Econ. Rev.* 61, 186–204.
- Bao, T., Hommes, C., Makarewicz, T., 2017. Bubble formation and (in)efficient markets in learning-to-forecast and -optimize experiments. *Econ. J.* 127, F581–F609.
- Benhabib, J., Farmer, R.E., 1999. Indeterminacy and sunspots in macroeconomics. In: Taylor, J.B., Woodford, M. (Eds.), *Handbook of Macroeconomics*. In: *Handbook of Macroeconomics*, vol. 1. Elsevier, pp. 387–448. chapter 6.
- Bernasconi, M., Kirchkamp, O., 2000. Why do monetary policies matter? An experimental study of saving and inflation in an overlapping generations model. *J. Monet. Econ.* 46 (2), 315–343.
- Bullard, J., Duffy, J., 1998. Learning and the stability of cycles. *Macroecon. Dyn.* 2 (01), 22–48.
- Camerer, C., 2003. *Behavioral Game Theory*. Princeton University Press.
- Duffy, J., Fisher, E.O., 2005. Sunspots in the laboratory. *Am. Econ. Rev.* 95, 510–529.
- Evans, G., Honkapohja, S., 1995. Adaptive learning and expectational stability: an introduction. In: Salmon, A.K.M. (Ed.), *Learning and Rationality in Economics*. Basil Blackwell, Oxford, pp. 102–126.
- Evans, G., McGough, B., 2018. *Learning to Optimize*. University of Oregon. Manuscript.
- Evans, G.W., Honkapohja, S., 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press.
- Gale, D., 1973. Pure exchange equilibrium of dynamic economic models. *J. Econ. Theory* 6, 12–36.
- Gardini, L., Hommes, C., Tramontana, F., de Vilder, R., 2009. Forward and backward dynamics in implicitly defined overlapping generations models. *J. Econ. Behav. Organ.* 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. *J. Econ. Dyn. Control* 33, 1052–1072.
- Hirota, S., Sunder, S., 2007. Price bubbles sans dividend anchors: evidence from laboratory stock markets. *J. Econ. Dyn. Control* 31 (6), 1875–1909.
- Hommes, C., 2011. The heterogeneous expectations hypothesis: some evidence from the lab. *J. Econ. Dyn. Control* 35, 1–24.
- Hommes, C., Sonnemans, J., Tuinstra, J., van de Velden, H., 2005. Coordination of expectations in asset pricing experiments. *Rev. Financ. Stud.* 18 (3), 955–980.
- Hommes, C., Sorger, G., Wagener, F., 2013. Consistency of linear forecasts in a nonlinear stochastic economy. In: *Global Analysis of Dynamic Models in Economics and Finance*. Springer-Verlag, Berlin Heidelberg.
- Hommes, C., Zhu, M., 2014. Behavioral learning equilibria. *J. Econ. Theory* 150 (C), 778–814.
- Lim, S.S., Prescott, E.C., Sunder, S., 1994. Stationary Solution to the Overlapping Generations Model of Fiat Money: Experimental Evidence. *Physica-Verlag HD, Heidelberg*, pp. 75–97.
- Lucas, R.E., 1972. Expectations and the neutrality of money. *J. Econ. Theory* 4 (2), 103–124.
- Lucas, R.E., 1986. Adaptive behavior and economic theory. *J. Bus.* 59 (4), S401–S426.
- Marimon, R., Spear, S.E., Sunder, S., 1993. Expectationally driven market volatility: an experimental study. *J. Econ. Theory* 61 (1), 74–103.
- Marimon, R., Sunder, S., 1993. Indeterminacy of equilibria in a hyperinflationary world: experimental evidence. *Econometrica* 61 (5), 1073–1107.
- Marimon, R., Sunder, S., 1994. Expectations and learning under alternative monetary regimes: an experimental approach. *Econ. Theory* 4 (1), 131–162.
- Marimon, R., Sunder, S., 1995. Does a constant money growth rule help stabilize inflation?: experimental evidence. *Carnegie-Rochester Conf. Ser. Public Policy* 43, 111–156.
- Spiliopoulos, L., 2012. Pattern recognition and subjective belief learning in a repeated constant-sum game. *Games Econ. Behav.* 75 (2), 921–935.

- Van Huyck, J.B., Cook, J.P., Battalio, R.C., 1994. Selection dynamics, asymptotic stability, and adaptive behavior. *J. Polit. Econ.* 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. *J. Stat. Softw.* 11 (10), 1–17.