

Evans, George W.; Hommes, Cars H.; McGough, Bruce; Salle, Isabelle

Working Paper

Are long-horizon expectations (de-)stabilizing? Theory and experiments

Bank of Canada Staff Working Paper, No. 2019-27

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Evans, George W.; Hommes, Cars H.; McGough, Bruce; Salle, Isabelle (2019) : Are long-horizon expectations (de-)stabilizing? Theory and experiments, Bank of Canada Staff Working Paper, No. 2019-27, Bank of Canada, Ottawa, <https://doi.org/10.34989/swp-2019-27>

This Version is available at:

<https://hdl.handle.net/10419/210767>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Staff Working Paper/Document de travail du personnel 2019-27

Are Long-Horizon Expectations (De-)Stabilizing? Theory and Experiments



by George Evans, Cars Hommes, Bruce McGough and Isabelle Salle

Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

Bank of Canada Staff Working Paper 2019-27

August 2019

Are Long-Horizon Expectations (De-)Stabilizing? Theory and Experiments

by

George Evans¹, Cars Hommes², Bruce McGough³ and Isabelle Salle⁴

¹Department of Economics,
University of Oregon & School of Economics and Finance,
University of St. Andrews

²Amsterdam School of Economics,
University of Amsterdam,
Tinbergen Institute
and
Bank of Canada, Ottawa

³Department of Economics,
University of Oregon

⁴Market Risks and Vulnerabilities
Financial Markets
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
ISalle@bank-banque-canada.ca

Acknowledgements

We thank the participants of the CREED seminar at the University of Amsterdam on June 6, 2016; the internal seminar at Utrecht University on February 9, 2017; the “Expectations in Dynamic Macroeconomic Models” workshop on August 28–30, 2017, at the Federal Reserve Bank of St. Louis; the internal seminars at the University of New South Wales on March 6 and the University of Technology Sydney on March 9, 2018; the workshop on Theoretical and Experimental Macroeconomics at the Bank of Canada on June 25–26, 2019; and the Economic Science meeting on July 7–11 at Simon Fraser University, Vancouver, for helpful discussions. This research has been partly financed by the EU FP7 project MACFINROBODS, grant agreement No. 612796. Isabelle Salle is grateful to the International Network on Expectational Coordination, funded by the Institute for New Economic Thinking, for financing her stay at the University of Oregon during the spring of 2016. None of the above are responsible for potential errors in this paper. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of Canada.

Abstract

We consider boundedly rational agents who do not plan over the infinite future but make trading plans at a finite, arbitrary horizon. We investigate the role of that horizon in the price dynamics of an asset in a Lucas tree model. We then design a laboratory experiment to test our theoretical predictions against the behaviors of human subjects. Short-horizon markets are prone to substantial and prolonged deviations from rational expectations (RE). By contrast, markets populated by even a modest share of long-horizon forecasters exhibit convergence towards the fundamental price. Longer-horizon forecasts do display more heterogeneity and thus some departure from RE; however this same heterogeneity also prevents the coordination of subjects on wrong anchors—a pattern of behavior that leads to mispricing in short-horizon markets. Long-horizon forecasts are well-described by adaptive learning, which delivers convergence to RE equilibrium, while short-horizon forecasts exhibit destabilizing trend-chasing.

Bank topics: Asset pricing; Central bank research; Economic models and financial markets.

JEL codes: C92, D84, E70

1 Introduction

Most macroeconomic and finance models involve long-lived agents making dynamic decisions in the presence of uncertainty. The benchmark modeling paradigm is the rational expectations (RE) hypothesis, which, in a stationary environment, can be captured by a one-step-ahead formulation of the model dynamics together with boundary conditions;¹ the impact of future plans at all horizons are fully summarized by one-step-ahead forecasts. Thus, under RE the issue of the decision horizon is hidden. When agents are more plausibly modeled as boundedly rational, one-step-ahead forecasts may be inadequate, and thus a stand must be taken on the decision horizon employed. In this paper, using a simple asset-pricing model, we investigate the importance of horizon length, both theoretically and in a laboratory experiment.

Under bounded rationality, theoretical and applied aspects of the decision horizon in general equilibrium settings have been considered by a variety of authors. At the short-horizon extreme, one-step-ahead “Euler equation” learning was employed, for example, in Evans & Honkapohja (2006); more recently, Evans & McGough (2018) show that one-step-ahead “shadow-price” learning provides a broad formal framework for agent-level bounded rationality in dynamic stochastic general equilibrium models. At the other extreme, an infinite-horizon approach was developed by Preston (2005). The work most closely related to this paper is Branch et al. (2012), in which general finite decision-horizon approaches were developed, including the one used in the current paper. The details, as well as the adaptation of this approach to our experimental framework, are discussed in Section 2.

The role of the forecast horizon is more than a theoretical question, as it may be relevant to many macroeconomic or financial issues: for example, forward guidance, fiscal policy or trading strategies. In fact, whether households or firms, when planning investment and savings, think only about the next period or instead think ahead several or many periods, and also how accurately they can forecast over those horizons, affects their economic and financial decisions and their reaction to policies. For example, how agents react to news about the future is critical to monetary policy, particularly now that publicly disclosing projections of future interest rates and forecasts of future macroeconomic variables has become a major part of central banking. The effects of fiscal transfers also depend on the ability of agents to evaluate the dynamic consequences of a given policy on future taxes and expenditures. Another important example, of particular relevance to our experiment, is that in financial markets it has long been thought that whether participants seek long-term returns or short-term profit opportunities is likely to affect their trade pattern and hence the volume of transactions as well as the level and volatility of asset prices.

The primary goal of this paper is to design an asset pricing model populated by boundedly rational agents with finite planning horizons that yields experimentally testable theoretical predictions. By tuning the horizon of the expectations, we can test in the lab our theoretical predictions about the market behavior under different configurations of expectations and horizons.

We choose the framework of a consumption-based asset pricing model *à la* Lucas (1978). In this model, heterogeneous expectations about future prices constitute a motive for trade between otherwise identical agents. To implement boundedly rational decision-making in our model, we impose that agents use “ T -step optimal learning” for a given planning horizon T , one of the approaches developed in Branch et al. (2012). Expectations about future asset prices constitute a

¹These boundary conditions include initial conditions on the state, as well as no-Ponzi scheme and transversality conditions. Typically, a non-explosiveness condition ensures these latter two.

central element of the price determination and impart positive feedback into the price dynamics: higher price forecasts translate into higher prices. Our implementation of bounded rationality leads to a particularly simple connection between individual decisions and expectations about future asset prices: an individual agent’s conditional demand schedule reduces to a linear function of their endowment, the market clearing price and the agent’s expectation of the *average* asset price over any given horizon. This latter feature facilitates elicitation of forecasts from the human subjects in the lab.

Our contribution stands at the crossroad between two strands of the literature: the learning literature, as implemented, e.g. in dynamic general equilibrium models (Evans & Honkapohja 2001), and the experimental literature concerned with behavioral finance. While our primary focus lies in the former, we borrow from the latter the laboratory implementation that allows us to design an original group experiment whose main features remain as close as possible to the theoretical learning setup (see Section 3).

Our experiment belongs to the class of “learning-to-forecast” experiments (LtFEs),² which focuses on the study of expectation-driven dynamics. An LtFE isolates the effects of expectations in a controlled environment. In these experiments, participants’ beliefs are elicited and their resulting boundedly optimal economic decisions, conditional on those beliefs, are computerized. The model’s expectational feedback permits expectation-driven fluctuations and (nearly) self-fulfilling price dynamics.

Our asset-pricing model is easily summarized: there is a fixed quantity of a single durable asset, yielding a constant, perishable dividend that comprises the model’s single consumption good. The initial allocation of assets is uniform across agents (referred to, in the experiment, as participants). Each period, each agent forms forecasts of future asset prices and, based on these forecasts and their current asset holdings, forms their asset demand schedules. These schedules are then coordinated by a competitive market-clearing mechanism, yielding equilibrium price and trades. Participants’ payoffs reflect both forecast accuracy and utility maximization. To accommodate the laboratory setting, the usual infinite-lifetime assumption is modified: at the end of each period, with a constant, known probability, the asset becomes worthless and the entire economy ends. This device emulates an infinite-horizon setting within an experimental economy with a fixed time frame and also yields a constant effective discount rate induced by the probability of termination. This economy has a unique perfect-foresight equilibrium price, which we refer to as the “fundamental price,” and which is determined by the dividend and the discount factor.

We consider four experimental treatments, based on horizon length, T : one with a short horizon ($T = 1$), one with a long horizon ($T = 10$), and two treatments with mixtures of short and long horizons. We have two main questions in mind: (i) Does the horizon of expectations matter for the aggregate behavior of the market? (ii) If so, how do the horizon and the heterogeneity of those horizons affect this behavior? In particular, are long-horizon expectations (de)stabilizing?

We find that markets populated only by short-horizon forecasters are prone to significant and frequently prolonged deviations from the fundamental price. By contrast, if all traders are long-horizon forecasters, the price path is consistent with convergence to the fundamental price. A detailed analysis of individual forecasts reveals that the failure of convergence in short-horizon markets reflects the coordination of participants’ forecasts on patterns derived from price histories, e.g. “trend-chasing” behavior. By contrast, the coordination of subjects’ forecasts appears more

²See the earlier contribution of Marimon et al. (1993). A more recent study within an experimental asset-pricing model is Hommes et al. (2005). This literature is surveyed in Duffy (2016).

challenging in longer horizon treatments: long-horizon forecasters display more disagreement. The resulting heterogeneity of long-horizon expectations impedes coordination on trend-chasing behavior and favors instead adaptive learning of the fundamental price. From those two polar cases, a natural follow-up question is what share of long-horizon forecasters would be large enough to stabilize the market price. Our findings suggest that even a modest share of them is enough.

A relatively large body of the literature has investigated financial markets in a laboratory setting. We are not reviewing these studies here, as our research questions are distinct. To the best of our knowledge, all the existing LtFEs operate within reduced-form models where only one-step-ahead expectations matter for the resulting price dynamics. Several experimental studies have been concerned with belief elicitation at longer horizons; see, for instance, Haruvy et al. (2007). However, in those studies, the participants' elicited forecasts do not themselves affect the price dynamics, which is a fundamental difference relative to our framework. In a pair of experimental contributions, Hirota & Sunder (2007) and Hirota et al. (2015) studied the influence of trading horizons in a double auction stock market, but their setting differs substantially from ours. Among other differences, their studies are not LtFEs, as their participants trade in every period a security that matures at a known date.³

The rest of the paper is organized as follows. Section 2 gives a general theoretical framework, while Section 3 details both the experimental design and our hypotheses based on predictions from the learning model, within a specific calibration of the framework. We conclude with a discussion in Section 5.

2 Theoretical framework: an asset-pricing model

The underlying framework of our experiment is a consumption-based asset-pricing model *à la* Lucas (1978). This model can be interpreted as a pure exchange economy with a single type of productive asset; at time t , each unit of the asset costlessly produces y_t units of consumption. The textbook model refers to this asset as a “tree” that produces “fruit.” In the experiment, we use the framing of a “chicken” producing “eggs.” We adopt this terminology mainly for two reasons: (i) to reduce the likelihood that participants with a background in economics or finance, which was the case for two-thirds of our participants, would recognize the textbook asset-pricing model; and (ii) to simplify the emulation of the infinite-horizon environment in the lab by considering an asset with a finite life (see below).

2.1 The infinite-horizon model

Because the consumption good is non-storable, holding assets is the only way to transfer wealth into the future. There are many identical agents, each initially endowed with a finite number $q_0 > 0$ of chickens. Each chicken renders a per-period quantity y_t of eggs. In each period, there

³The authors further define the horizon as the number of periods a trader participates in the market relative to the maturity of the security. Hence, for short-horizon traders, the dividend is payoff-irrelevant and they seek capital gains. As a result, speculation motives may drive the price away from fundamentals and it loses the dividend anchor. The authors conclude that longer maturities lead to more price bubbles due to more trading opportunities. By contrast, the horizon in our setup corresponds to a forecast horizon and the forecasts determine the realized market price. We find that long-horizon planning decisions stabilize the market.

is a market for chickens. At the beginning of the period, each agent collects the eggs from her chickens and then decides how many of the eggs to consume and how many to sell for additional chickens. The agent could instead decide to sell some chickens in order to increase her current egg consumption. This decision depends on her forecast of future chicken prices, and she aims to smooth her consumption over time.

To formalize the model, we consider the representative agent's problem:

$$\begin{aligned} \max \quad & E \sum_{t \geq 0} \beta^t u(c_t) \\ & c_t + p_t q_t = (p_t + y_t) q_{t-1}, \text{ with } q_{-1} = q \text{ given,} \end{aligned}$$

where $u(\cdot)$ is a CRRA utility function with usual properties, q_{t-1} is the quantity of chickens held at the beginning of period t , c_t is the quantity of eggs consumed, p_t is the goods-price of a chicken (the price of a chicken in terms of eggs) and y_t is the dividend process to be specified. Here E denotes the subjective expectation of the agent.

For the lab implementation, we assume a deterministic and constant dividend equal to an integer y in every period. In this case, rational expectations reduce to perfect foresight, and the corresponding first-order condition (Euler equation) is given by

$$u'(c_t) = \left(\frac{p_{t+1} + y}{p_t} \right) u'(c_{t+1}).$$

Assuming identical agents, there is no trade in equilibrium; thus, $c_t = q_t y$ in every period. It follows that in the non-stochastic case, the perfect-foresight steady state satisfies:

$$c = qy \text{ and } p = \frac{\beta y}{1 - \beta}.$$

Throughout the paper we will refer to $p = (1 - \beta)^{-1} \beta y$ as the fundamental price, or the fundamental value, of the asset. We note here that in the finite-horizon case the steady state corresponds to the fundamental price and associated consumption $c = qy$. Since perfect foresight is a special case of rational expectations, we also refer to this as the REE.

2.2 The model with finite-horizon agents

We relax the assumption of perfect foresight over an infinite horizon and consider the behavior of a boundedly rational agent with a finite planning horizon $T \geq 1$. In the non-stochastic case, this agent solves the following problem at time t :

$$\begin{aligned} \max \quad & E_t \sum_{k=0}^T \beta^k u(c_{t+k}) \\ & c_{t+k} + p_{t+k} q_{t+k} = (p_{t+k} + y_{t+k}) q_{t+k-1}, \end{aligned} \tag{1}$$

with q_{t-1} , q_{t+T}^e given.

The infinite-horizon version presented in Section 2.1 can be viewed as a limit of this model as $T \rightarrow \infty$. Letting dx be the deviation of an arbitrary variable x from its steady-state value, the following proposition provides the first-order approximation to the time t asset demand dq_t in terms of current and expected future variables:

Proposition 1. Let $\sigma = -cu''(c)/u'(c)$. If dividends are constant at $dy_k = 0$ for all k then

$$\begin{aligned} dq_t = & \left(\frac{1 - \beta^T}{1 - \beta^{T+1}} \right) dq_{t-1} - \left(\frac{(1 - \beta)(1 - \beta^T)}{(1 - \beta^{T+1})} \frac{q}{p\sigma} \right) dp_t \\ & + \left(\frac{(1 - \beta)\beta^T}{1 - \beta^{T+1}} \right) dq_{t+T}^e + T \left(\frac{(1 - \beta)^2 \beta^T}{1 - \beta^{T+1}} \frac{q}{p\sigma} \right) \frac{1}{T} \sum_{k=1}^T dp_{t+k}^e. \end{aligned} \quad (2)$$

Proof. See Appendix A.1. ■

Hence, the individual demand schedule of an agent depends negatively on the market price dp_t , positively on her past and future endowments dq_{t-1} and dq_{t+T}^e , and, perhaps unexpectedly, on her forecast of the *average price* $\frac{1}{T} \sum_{k=1}^T dp_{t+k}^e$ over her planning horizon. This final dependency merits further discussion, which we attend to now.

To gain intuition for the dependence of asset demand on average price forecasts, it is helpful to rewrite the model using the more common language of *returns*. The details for the computations provided in this section may be found in Appendix A.2. For simplicity, assume perfect foresight, take $t = 0$, and assume $dq_{-1} = dp_{-1} = 0$.

First we linearize the flow constraint (1) to get

$$\begin{aligned} dc_t + pdq_t + qdp_t &= qdp_t + (p + y)dq_{t-1}, \text{ or} \\ dc_t + pdq_t &= (p + y)dq_{t-1}. \end{aligned}$$

Since $dq_{-1} = 0$, it follows that $dc_0 + pdq_0 = 0$. Thus to determine asset demand dq_0 it is sufficient to examine the demand for time-zero consumption dc_0 .

Let Q_t be the goods-value of current asset holdings: $Q_t = p_t q_t$. The agent's problem becomes

$$\max \sum_{k=0}^T \beta^k u(c_k) \text{ subject to } c_k + Q_k = R_k Q_{k-1},$$

with $Q_{-1} = pq$ and Q_T given. Here, as is standard, $R_k = p_{k-1}^{-1}(p_k + y)$. The associated Euler equation is $u'(c_k) = \beta R_{k+1} u'(c_{k+1})$, which linearizes as $dc_k = dc_{k+1} - ((1 - \beta)Q/\sigma) dR_{k+1}$.

The linearized Euler equations and flow constraints may be iterated, combined and simplified (see Appendix A.2) to yield the time-zero consumption demand,

$$dc_0 = \frac{1 - \beta}{1 - \beta^{T+1}} \left(qdp_0 - \beta^T pdq_T - \beta^T qdp_T + \frac{Q(\sigma - 1)}{\sigma} \sum_{k=1}^T \beta^k dR_k + \frac{Q}{\sigma} \beta^{T+1} \sum_{k=1}^T dR_k \right), \quad (3)$$

which, in the infinite-horizon case, may be written as

$$dc_0 = (1 - \beta) \left(qdp_0 + \frac{Q(\sigma - 1)}{\sigma} \sum_{k=1}^{\infty} \beta^k dR_k \right). \quad (4)$$

Equation (4) has a familiar form: consumption today depends on price today and on the present value of excess future returns (and the nature of the dependence on these discounted returns is governed by the size of σ relative to one).

Now notice that $dR_k = (\beta p)^{-1}(\beta dp_k - dp_{k-1})$. In words, a rise (say) in p_k raises the return in period k , but it lowers the return in period $k+1$ by an amount which, when discounted, exactly offsets the previously noted rise. This leads to the somewhat surprising result that a change in future prices does not, to first order, impact the present value of excess future returns. Specifically, $\partial/\partial dp_m \sum_{k=1}^{\infty} \beta^k dR_k = 0$. This further implies that in the infinite-horizon case, dc_0 , and hence asset demand dq_0 , does not depend on dp_m for $m \geq 1$.

However, in the finite-horizon case it is crucial to correct for the terminal condition. For $1 \leq m < T$, changes in prices still do not impact the discounted term, i.e. $\partial/\partial dp_m \sum_{k=1}^T \beta^k dR_k = 0$. Taking into account the last term on the right-hand side of Equation (3), we compute

$$\frac{\partial dc_0}{\partial dp_m} = \frac{\partial}{\partial dp_m} \frac{Q}{\sigma} \beta^{T+1} \sum_{k=1}^T dR_k = \beta^T \frac{Q}{\sigma p} (\beta - 1).$$

Importantly, while non-zero, $\partial dc_0/\partial dp_m$ is *independent* of m for $1 \leq m < T$. In Appendix A.2, we show that the same computation holds for changes in p_T , i.e. $\partial dc_0/\partial dp_T = \beta^T \frac{Q}{\sigma p} (\beta - 1)$. Since $\partial dc_0/\partial dp_m$ is independent of m for $1 \leq m \leq T$, it follows that average expected price path is a sufficient statistic for the determination of dc_0 , and hence for asset demand dq_0 .

2.3 Price dynamics with heterogeneous agents

We consider a market populated by heterogeneous agents, indexed by i , possibly forecasting over heterogeneous horizons. Formally, let α_i be the proportion of agents of type i , for $i = 1, \dots, I$, and let

$$\alpha = \{\alpha_1, \dots, \alpha_I\} \text{ and } \mathcal{T} = \{T_1, \dots, T_I\}.$$

For convenience we assume that agents of the same type hold the same forecasts; however, since we allow agents of different types to have planning horizons of the same length, this assumption is innocuous. In such a framework, heterogeneous wealth and expectations lead to different individual demand schedules, which induce out-of-equilibrium dynamics and provide a motive for trade.

At this stage, one needs to specify agents' beliefs about the future wealth at the horizon T_i , $dq_{i,t+T_i}^e$. Bearing in mind the need for simplicity for the purpose of the lab implementation, we assume that agents hold random-walk forecasts of future wealth, so that $dq_{i,t+T_i}^e = dq_{i,t-1}$, $\forall i, t$. Note that in a REE, the wealth of each agent is constant over time and hence satisfies this assumption. The random-walk assumption $dq_{i,t+T_i}^e = dq_{i,t-1}$ is a natural generalization under bounded rationality. The detailed implications of our assumption are discussed in Appendix A.3.1.

Under this assumption, the coefficients associated to the past and the future endowments in Proposition 1 sum up to one, and the individual demand schedule of an agent of type i reduces to

$$dq_i(i) = dq_{i-1}(i) - \phi g(T_i) dp_t + \phi h(T_i) d\bar{p}_t^e(i, T_i), \quad (5)$$

where $d\bar{p}_t^e(i, T_i)$ denotes the expected average price of an agent of type i with planning horizon T_i , and where

$$\phi = \frac{(1-\beta)q}{p\sigma}, \quad g(T) = \frac{1-\beta^T}{1-\beta^{T+1}}, \text{ and } h(T) = \frac{(1-\beta)T\beta^T}{1-\beta^{T+1}}.$$

Market clearing in each period implies $\sum_i \alpha_i dq_t(i) = \sum_i \alpha_i dq_{t-1}(i) = 0$, $\forall i, t$, which uniquely determines the price p_t :

$$dp_t = \sum_i \xi(\alpha, \mathcal{T}, i) d\bar{p}_t^e(i, T_i), \text{ where} \quad (6)$$

$$\xi(\alpha, \mathcal{T}, i) = \frac{\alpha_i h(T_i)}{\sum_j \alpha_j g(T_j)}.$$

Equation (6) states that the time t price only depends on the agents' forecasts of the *average* price of chickens over their planning horizon, i.e. $\{d\bar{p}_t^e(i, T_i)\}_{i=1}^I$. The asset-pricing model with heterogeneous agents is therefore an *expectational feedback* system, in which the perfect foresight steady-state price is exactly self-fulfilling and is unique.

Two special cases are of interest. If expectations are homogeneous across types, $d\bar{p}_t^e(i, T_i) = dp_t^e$, $\forall i$, the following holds:

$$dp_t = \xi(\alpha, \mathcal{T}) dp_t^e, \text{ where} \quad (7)$$

$$\xi(\alpha, \mathcal{T}) = \sum_i \xi(\alpha, \mathcal{T}, i) = \frac{\sum_i \alpha_i h(T_i)}{\sum_j \alpha_j g(T_j)}.$$

If, in addition, all agents have the same planning horizon T , i.e. $\mathcal{T} = \{T\}$, we have

$$dp_t = \xi(T) dp_t^e, \text{ where } \xi(T) = \frac{(1-\beta)T\beta^T}{1-\beta^T}. \quad (8)$$

The following proposition characterizes the expectational feedback parameter $\xi(\alpha, T)$ when expectations are homogeneous.

Proposition 2. *Let $I \geq 1$, $\alpha_i \geq 0$, $\sum \alpha_i = 1$, $T_i \geq 1$. Then:*

1. $0 < \xi(\alpha, \mathcal{T}) < 1$.
2. *If agents have the same planning horizon and if $T, T' \in \mathbb{N}$ and $1 \leq T < T'$, then $\xi(T) > \xi(T')$.*
3. *If $\mathcal{T} = \{T_1, T_2\}$ and if $T_1 < T_2$, then an increase in α_1 leads to an increase in $\xi(\alpha, \mathcal{T})$.*

Proof. See Appendix A.3.2. ■

The above results are easily interpreted: Item 1 of Proposition 2 says that the feedback in this system is positive but always less than one: a rise in the expected average future price leads to a rise in the current realized price. Item 2 considers the case in which there is a single planning horizon and states that increasing the planning horizon reduces the expectational feedback (for a given discount factor β). In particular, the strongest feedback occurs with planning horizon equal to one, in which case $\xi(1) = \beta$. Item 3 considers the case in which there are two agent types with different planning horizons, and states that increasing the proportion of agents using the shorter planning horizon increases the expectational feedback.

It is natural to examine how price expectations might evolve over time as new information becomes available. Here we consider whether agents using simple learning rules would eventually coordinate their forecasts on the REE. Put differently, is the REE stable under adaptive learning?

To assess stability we assume each of the N agents uses the following adaptive learning rule:

$$d\bar{p}_t^e(i, T_i) = d\bar{p}_{t-1}^e(i, T_i) + \gamma_t(dp_{t-1} - d\bar{p}_{t-1}^e(i, T_i)). \quad (9)$$

Here, $0 < \gamma_t \leq 1$ is called the “gain” sequence, which is assumed to satisfy $\sum_t \gamma_t = \infty$. There are two prominent cases in the literature: the “decreasing gain” case with $\gamma_t = t^{-1}$, which provides equal weight to all data; and the “constant gain” case with $\gamma_t = \gamma$, which discounts past data. While we assume here that all agents use the same gain, a continuity argument extends our results to allow for local heterogeneity.

Corollary 1. *Under both decreasing gain and constant gain, with $\gamma < 1$, $d\bar{p}_t^e(i, T_i) \rightarrow 0$ and hence $dp_t \rightarrow 0$.*

Proof. See Appendix A.3.3. ■

Corollary 1 shows that under adaptive learning of the form (9), the price dynamics converge to the REE, i.e. to the fundamental value. This result is independent of the number of agent-types and the length of their horizons.

This convergence result concerns asymptotic stability, but we are also interested in transition dynamics. For simplicity, we focus here on the constant gain learning, and show in the Appendix that the market price converges to zero at a constant rate $1 - \zeta$, where $\zeta = dp_t/dp_{t-1}$. The following result characterizes the dependence of the rate of convergence on the planning horizon.

Corollary 2. *Under constant gain learning, the rate at which the market price converges to the fundamental value is increasing in the individual planning horizons T_i .*

Proof. See Appendix A.3.3. ■

Additional comments regarding transition dynamics are warranted. Stochastic models of the form (7) have been studied in closely related setups that allow for heterogeneous expectations across agents and constant gain learning. From the adaptive learning literature it is also known that the extent and speed of convergence depends on the size of the expectational feedback parameter, which corresponds here to $\xi(\alpha, T)$.⁴ Furthermore, in short-horizon settings a number of authors have noted the possibility that near-random-walk behavior leads to significant departures from REE and excess volatility in asset prices.⁵ In our model this phenomenon arises most forcefully when $T = 1$ and β is near one so that ξ is near one. The central idea is as follows. For simplicity, consider the case of homogeneous horizons and expectations, as in Equation (8), except that we include a small i.i.d. random shock so that $dp_t = \beta dp_t^e + u_t$. In this case, a random-walk forecast rule $dp_t^e = dp_{t-1}$ implies $dp_t = \beta dp_{t-1} + u_t$, i.e. it generates a dynamic process for prices that, while asymptotically stationary, is very close to a random walk. More generally, the key novelty of our theoretical results and associated experiments is to explicitly examine the impact of the forecast horizon on the behavior of asset prices.

A value of the expectational feedback parameter ξ near one also has implications for forecast accuracy. In particular, for some simple salient forecast models, including those based on possibly-weighted sample averages (γ small) or near random walks (γ large), expectations are nearly self-fulfilling. Thus in this case, even if price level is far from the REE, the agents’ forecast errors can be small.

⁴See, e.g. Ch. 3.2, 3.3 and 7.5 of Evans & Honkapohja (2001).

⁵See, e.g. Branch & Evans (2011), Bullard et al. (2010), Hommes & Zhu (2014), Adam et al. (2016) and Blanchard & Watson (1982). For related experimental results, see Hommes et al. (2005).

The theoretical results and discussion above point to the following implications for this model under learning, which we would expect to be reflected experimentally:

Implication 1: There will be convergence over time of the price and the individual forecasts towards the REE.

Implication 2: The extent and speed of convergence will be greater the smaller is the expectational feedback parameter $\xi(\alpha, \mathcal{T})$.

Implication 3: The extent by which forecasts deviate from REE will be smaller the smaller is $\xi(\alpha, \mathcal{T})$.

Implication 4: The level of price volatility will be lower the smaller is $\xi(\alpha, \mathcal{T})$.

These implications are reflected in the hypotheses we develop and test below.

3 The experimental design

In this section we outline the design of our group experiment and list specific hypotheses used to test the theoretical results obtained in the previous section. The experiment is couched in terms of a metaphorical asset market in which the assets are chickens (and thus finite-lived), and the dividends are eggs (and thus perishable), which comprise the experiment's unique consumption good. Participants play the role of traders who make saving decisions based on forecasts of future chicken prices. In the experiment, participants submit price forecasts that are then coupled with the decision rules derived in Section 2 to determine their demand-for-saving schedules. Equilibrium prices and saving decisions are then determined each period via market clearing.

3.1 Environment and procedures

Each group in the experiment is composed of $J = 10$ participants. At the opening of a market, each forecaster/trader is endowed with a given number of chickens. This number is the same across all forecasters/traders, but participants can only observe their own endowment and do not know the total number of chickens in the market.

Upon entering the lab, each participant is assigned the single task of forecasting the average market price of a chicken in terms of eggs over a given horizon, and this horizon remains the same throughout the experiment. Trading is computerized: based on participants' forecasts, market clearing determines the actual amount of trade of each participant, given by Equation (5), together with the current market price as specified by Equation (6). The timing of the events within one experimental period is illustrated in Figure 1.

The dividend is common knowledge, and participants operate under no-short-selling and no-debt constraints. In each period, they have to consume at least one egg. Eggs are both the consumption good and the medium of exchange, but only chickens are transferable between periods in the experiment (see Crockett et al. (2019) for a similar setup).

Transposing this type of model to a laboratory environment requires resolving a number of issues, as discussed for instance in Asparouhova et al. (2016). Two major concerns are the emulation of stationarity and infinitely lived agents. Stationarity is an essential feature as it rules out rational motives to deviate from fundamentals, hence allowing us to get cleaner data on potential behav-




Figure 1: Timing of events within one period of an experimental market

ioral biases. An infinite-lifetime setting, together with exponential discounting and the dividend process, determines the fundamental value of the asset. This may play an important role in the belief formation process of the participants.

We use the standard random termination method originally proposed by Roth & Murnighan (1978) to deal with infinite lifetime in the laboratory. If each experimental market has a constant and common-knowledge probability of ending in each period, the probability of continuation is known to theoretically coincide with the discount factor. In the instructions of our experiment, the metaphor of the chickens allows us to tell the participants the story of an avian flu outbreak that may occur with a 5% probability in each period (corresponding to a discount factor $\beta = 0.95$). If this is the case, the market terminates: all chickens die and become worthless.

As for the stationarity issue, we choose a constant dividend process. The fundamental value associated with this dividend value and discount factor was not given to the participants. However, we think it likely that the experimental environment, including in particular the constant dividend process, is concrete enough to induce the idea of a fundamental value for a chicken in terms of eggs to the participants.

However, as discussed in Asparouhova et al. (2016), a major difficulty lies in the constant termination probability (or, equivalently, discount factor). Participants should perceive the probability of a market to end to be the same at the beginning of the experimental session as towards the end of the time span for which they have been recruited. We therefore use the “repetition” design of Asparouhova et al. (2016): we recruited the participants for a given time and ran as many markets as possible within this time frame. Furthermore, we recruited them for 2 hours and 30 minutes but completed most of the sessions within 2 hours so as to keep the participants’ perception of the session’s end in the distant future throughout the experiment (see also Charness & Genicot (2009) for such an implementation). We did so by starting a new market only if not more than 1 hour and 50 minutes had elapsed since the participants entered the lab. In case a market was still running after this time constraint, the experimenter would announce that the current 20-period block (see below) was the last one.

Finally, our framework involves two additional difficulties. Most importantly, participants have to form forecasts over a given horizon, say over the next 10 periods, but the market may terminate before period 10. In this case, the average price corresponding to their elicited predictions is not realized, and participants’ tasks cannot be evaluated (see below how the payoffs are determined). In order to circumvent this issue, we use the “block” design proposed by Fréchette & Yuksel (2017): each market is repeated in blocks of a given number of periods, and the termination or continuation of the market is observed only at the end of each block. This design allows the experiment to continue at least for the number of periods specified in the block, without altering

	<i>Markets</i>				
	Market 1	Market 2	Market 3	Market 4	Market 5
Dividend y (“eggs”)	2	4	1	5	3
Fundamental price p	38	76	19	95	57
Initial endowment of “chickens” q	4100	2100	8200	1700	2700

Table 1: Calibration of the markets, all groups, all treatments

the emulation of the stationary and infinite living environment from a theoretical viewpoint.

In our experiment, the length of a block is taken to be 20 periods, which corresponds to the expected lifetime of a chicken with a 5% probability of termination. The random draws in each period are “silent,” and participants observe only every 20 periods whether the chickens have died during the previous 20 periods. If this occurred, the market terminates and they enter a new market from period 1 on. If this did not occur, the market continues for another 20-period block. In period 40, participants observe whether a termination draw has occurred between periods 20 and 40. If this is the case, the market terminates and a new one starts; if not, participants play another 20-period block till period 60, etc. Only periods during which the chickens have been alive count towards the earnings of the participants.

We need to prevent any knowledge of the fundamental to be carried over from one market to the next. If participants happened to coordinate their price expectations around the fundamental price in one market, the realized price would converge to its fundamental value, and we would not learn much from the next market if the fundamental value remains the same. For this reason, we keep the same termination probability (i.e. the same discount rate β) but vary the dividend process y between markets so as to vary the equilibrium price, and pose a new problem of coordination to the participants in every market. We further vary the initial endowment of chickens to keep liquidity and utility levels constant across markets. Those values are collected in Table 1.⁶ Upon entering each new market, participants receive the corresponding values through a pop-up message on their screen, and those values remain on the screen throughout the market (see Figure 2).

To avoid perfect predictions, we add a small noise term to the price. This term is drawn from a normal distribution with mean zero and standard deviation 0.25, and is the same for all markets.

3.2 Payoffs

We elicit price forecasts from participants, but those forecasts translate into trade decisions, and the predictions of our theoretical model partly rely on the properties of the utility function and the incentive to smooth consumption over time. For this reason, the payoff of the participants consists of two parts: at the end of each market, all participants receive experimental points based *either* on forecast accuracy *or* on their egg consumption with equal probability. This random choice of payoff is designed to avoid “hedging” and maintain equal incentives towards the two objectives (forecasting and consuming) throughout each market. Payoff tables are reported in Appendix D.

⁶We remark that only integer values of chickens and eggs are allowed to be traded/consumed. The large number of chickens renders this imposition inconsequential.

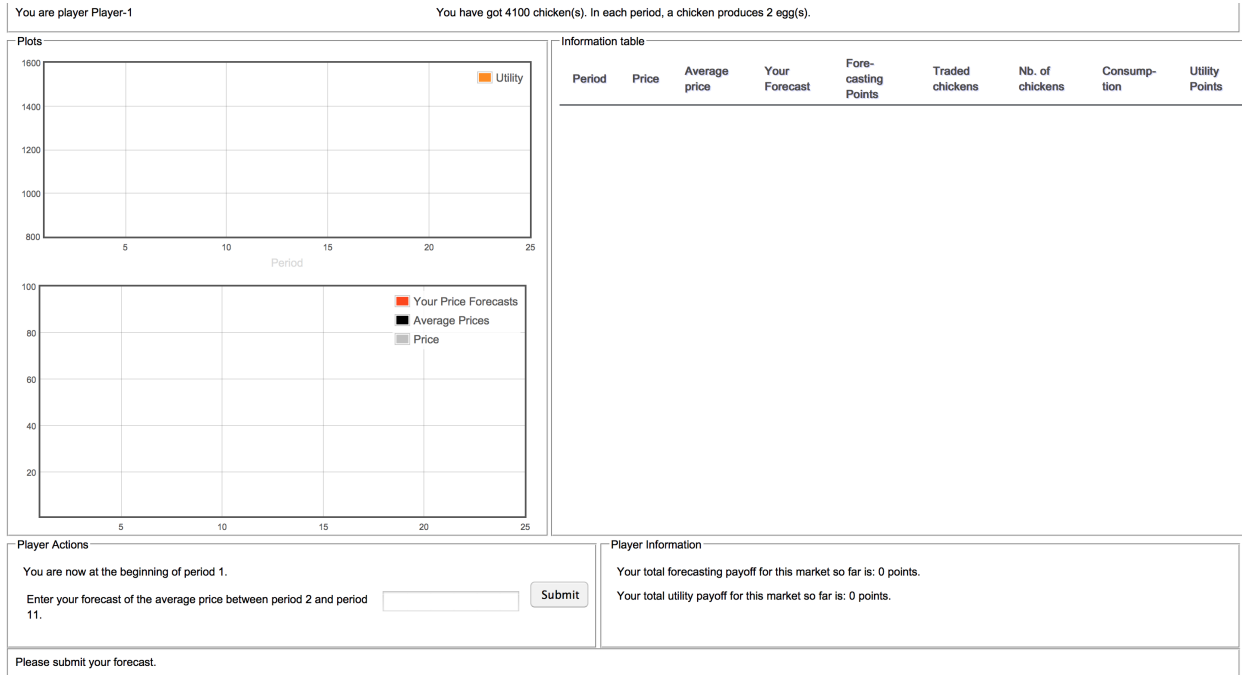


Figure 2: Screen shot of the experimental interface for a long-horizon participant

The consumption payoff is a linear transformation of a CRRA utility function $u(\cdot) = \frac{c^{1-\sigma}-1}{1-\sigma}$ with $\sigma = 1$. We use the logarithmic function $120 \cdot \ln(c)$ ($c \geq 1$) for several reasons. First, by specifying a utility function, we have a tighter control on subjects' preferences. Second, by choosing a concave reward scheme, we can induce consumption smoothing behavior that underlies the predictions from the theoretical model (see also Crockett et al. (2019) for such a payoff function). As noted above, participants are paid only for periods during which chickens continue to be alive. The payoff based on utility is a straightforward sum of their utility realized in each of those periods.

Payoffs based on predictions are made as follows. First, as usual in LtFEs, the forecasting payoff is a decreasing function of the squared forecast errors:

$$\max \left(1100 - \frac{1100}{49} (\text{error})^2, 0 \right).$$

If the error is higher than 7, the payoff is zero. Second, we must take account of the fact that there are necessarily periods before the death of the chickens for which forecast errors are not available. Consequently, the number of realized average prices over T periods, and the associated forecasting payments, is lower than the number of utility payments that take place in every period. To circumvent this discrepancy, the last rewarded forecast is paid $T + 1$ times to the participants. This also incentivizes them to submit accurate forecasts for every single period, as they are uncertain about which one will be the last and, hence, the most rewarded. In the special case in which the chickens die in the first block before $T + 1$ periods, participants were paid on utility. At the end of all the markets, the total number of points earned by each participant was converted into euros at a pre-announced exchange rate, and paid privately.

3.3 Instructions and information

Participants were given instructions that they could read privately at their own pace (these instructions are reported in Appendix D). The instructions contain a general description of the markets for chickens, explanations about the forecasting task and how it translates into computerized trading decisions, information about the payoffs, and payoff tables, as well as an example. The instructions convey a qualitative statement of the expectations feedback mechanism that characterizes the underlying asset pricing model.

At the end of the instructions, the participants had to answer a quiz on paper. Two experimenters were in charge of checking the accuracy of their answers, discussing their potential mistakes and answering privately any question. The first market opens only after all participants had answered accurately all questions of the quiz. This procedure allows us to be confident that all participants start with a reasonable understanding of the experimental environment and their task. Of the participants, 90% (218) reported that the instructions were understandable, clear or very clear.

3.4 Hypotheses and treatments

In the experiment, participants are either short-horizon forecasters, in which case they forecast the next period's price ($T = 1$) or they are long-horizon forecasters, in which case they forecast the average price over the next ten periods ($T = 10$). We design four treatments to test the implications of the model discussed in Section 2.3 as distinguished by the proportion of long-horizon forecasters. These treatments, labeled L, M50, M70 and S respectively, are summarized in Table 2. This two-type setup leads to a special case of the pricing equation (6) developed in Section 2.2 that allows for a dependence on the average forecasts of both the short-horizon and the long-horizon forecasters. From the implications established in Section 2.3, we now derive the first three main hypotheses to be tested through the experimental treatments.

First, based on Corollary 1, one might expect convergence in all treatments since the feedback parameter is always lower than one. However, the implications at the end of Section 2.3 suggest that convergence to the REE can be tenuous if ξ is close to one. Item 3 of Proposition 2 shows that the feedback parameter $\xi \in (0, 1)$ increases with the share of short-horizon forecasters α . Figure 3 illustrates the effect of that share α on the feedback parameter ξ within the specific calibration of the model that we implement in the laboratory. In particular, Figure 3a shows how a rise in α increases ξ . As is clear from that figure, the impact on ξ of the share of short-horizon forecasters is nonlinear, magnifying the stabilization power of even a small share of long-horizon agents. Figure 3b illustrates the same phenomenon by plotting the expectational feedback equation (7) for the four different treatments. Note that the convergence challenge posed by a near-unity feedback parameter is particularly acute in Tr. S since in that case $\xi = \beta$.

These considerations, together with Implications 1, 2 and 3, suggest the first hypotheses:

Hypothesis 1a (Price convergence). *Under each treatment, participants' average forecasts and the resulting price level converge towards the REE.*

Hypothesis 1b (Price deviation). *The higher the share of short-horizon forecasters, the more likely failures to converge towards the REE to occur.*

	<i>Treatments</i>			
	Tr. L	Tr. M50	Tr. M70	Tr. S
Share α (and number of forecasters) with horizon $T = 1$	0 (0 subject)	0.5 (5 subjects)	0.7 (7 subjects)	1 (10 subjects)
Share $1 - \alpha$ (and number of forecasters) with horizon $T = 10$	1 (10 subjects)	0.5 (5 subjects)	0.3 (3 subjects)	0 (0 subject)

Table 2: Summary of the four experimental treatments

Furthermore, Implication 4 together with Proposition 2 provides the following hypothesis:

Hypothesis 2 (Price volatility). *Increasing the share of short-horizon planners increases the level of price volatility.*

Our theoretical results suggest that coordination of agents' expectations will increase over time as agents learn the REE. Since heterogeneous expectations provide a motive for trade in our experiment, we test the following hypothesis in all experimental treatments:

Hypothesis 3 (Eventual coordination). *The price predictions of the participants become more homogeneous over time. As a consequence, trade decreases over time.*

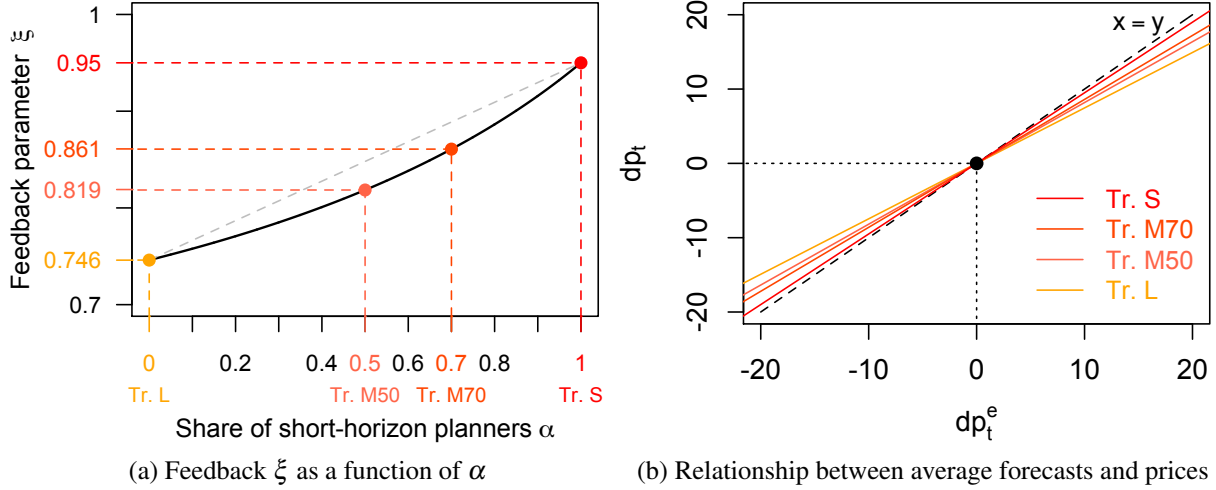
Besides providing an empirical test of the theoretical implications of the model, one further advantage of learning-to-forecast experiments is that they make it possible to collect “clean” data on individual expectations because the information, underlying fundamentals and incentives are under the full control of the experimenter. We can then use this rich dataset to test additional hypotheses regarding participants' forecasting behavior. In the current context, it is of interest to compare how short-horizon participants form their forecasts, and how they differ from long-horizon ones. A variety of factors suggest that long-horizon forecasting is more challenging than short-horizon forecasting. For example, long-horizon forecasting involves a sequence of shocks and thus more noise. Also, short-horizon forecasting involves contemplation of only a single data point, and hence a lighter cognitive load.

This discussion suggests that there may be more variation of price forecasts for long-horizon forecasters than for short-horizon forecasters. As a measure of variation of price forecasts, we use *cross-sectional dispersion*, defined in terms of the relative standard deviation of subjects' forecasts within each period. We have the following two companion hypotheses:

Hypothesis 4 (Coordination and forecast horizons). *Long-horizon forecasters exhibit more heterogeneity, i.e. greater cross-sectional dispersion of price forecasts, than do short-horizon forecasters.*

Hypothesis 5 (Trade volume and forecast horizons). *The higher the share of long-horizon forecasters in a market, the greater the cross-sectional dispersion of price forecasts and, hence, the higher the trade volume.*

After a few words on the implementation of the experiment, Section 4 discusses our findings in the context of these hypotheses.



Note: Left panel: Feedback ξ in price equation (6) as a function of the share α of short-horizon forecasters, $\beta = 0.95$, $T_1 = 1$ and $T_2 = 10$; Right panel: relationship between realized prices and forecasts (in deviation from fundamentals).

Figure 3: Price equation in the four experimental treatments

3.5 Implementation

The experiment was programmed using the Java-based PET software.⁷ The experimental sessions were run in the CREED lab at the University of Amsterdam between October 14 and December 16, 2016. Most subjects (124 out of 240) had participated in experiments on economic decision making in the past, but no person participated more than once in this experiment. With the four treatments, each involving six groups of ten subjects, we recruited a total of 240 subjects, who participated in a total of 63 markets, ranging from 20 to 60 periods. The average earnings per participant amount to 22.9 euros (ranging from 10.8 to 36.6 euros).

4 The experimental results

We first provide, in Section 4.1, a graphical overview of the price data from the experimental markets. In Section 4.2 we then proceed to conduct cross-treatment statistical comparisons in the context of our five hypotheses. Section 4.3 conducts an empirical assessment of convergence to REE using price data. Finally, in Section 4.4 we use individual data to empirically assess the types of forecasting models used by short- and long-horizon agents.⁸

⁷The PET software was developed by AITIA, Budapest under the FP7 EU project CRISIS, Grant Agreement No. 288501.

⁸Throughout our study, we adopt a 5% confidence threshold to assess statistical significance. When carrying out econometric analysis, we use OLS estimates, autocorrelation in error terms is detected by Breusch-Godfrey tests, and heteroskedasticity using Breusch-Pagan tests. When needed, we use the consistent estimators described in Newey & West (1994). Significant differences between distributions are established using K-S tests and Wilcoxon rank sum tests to address non-normality issues.

4.1 A first look at the data

Figure 4 displays an overview of the realized prices in the experimental markets for each of the four treatments. Each line represents a market, with the reported levels corresponding to the deviations from the market's fundamental value, expressed in percentage points.⁹ Plots with individual forecast data for each single market are given in Appendix B: see figures 10 to 12. In those figures, blue corresponds to long-horizon ($T = 10$) forecasts, red to short-horizon ($T = 1$) forecasts, dots to rewarded forecasts and crosses to non-rewarded forecasts. Finally, the solid line is the realized price and the dashed horizontal line is the fundamental price.

A first visual inspection of the market price data leads us to identify three different emerging patterns: (i) *convergence* to the fundamental price (see, for instance, in Figure 4d, Tr. L, Gp. 2 in purple or Gp. 6 in orange); (ii) *mispricing*, that we characterize by mild or dampening oscillations around a price value that is different from the fundamental value; either above the fundamental price, i.e. *overpricing*, or below the fundamental price, i.e. *underpricing* (see, for an example of each type of mispricing, the two markets played by Gp. 1 in Tr. M70 on Figure 4b, red lines); and (iii) *bubbles and crashes*, described by large and amplifying oscillations (where the top of the “bubble” reached several times the fundamental value); see, for instance, the markets of the first group in Tr. S (Figure 4a, red lines).

This first glance at the data already leads us to question Hypothesis 1a, as it is clear that not every market exhibits price convergence towards the fundamental value. On the other hand, we see patterns in the data that are in line with Hypothesis 1b: while large deviations from fundamentals are observed in the short-horizon treatments (Tr. S and Tr. M70), they are absent from the long-horizon treatments (Tr. M50 and Tr. L), where the price appears to evolve closer to the REE. Moreover, the problem of mispricing seems particularly acute in the short-horizon markets.


Interestingly, though, the observed bubbles break endogenously, which is *not* usual in LtFEs.¹⁰ There are several specific features of our environment that may be behind this phenomenon: (i) the framing as, at some point, participants may wonder whether a chicken that can die and produces only few eggs in each period is really worth so many eggs, or (ii) incentives related to the payoff-relevant utility in our setting. In particular, some participants reported in the end-of-experiment questionnaire that they attempted to bring back the price to lower values because they experienced low payoff along a bubble.¹¹

In the rest of this section, we explore the differences between treatments and confront these with our theoretical implications and experimental hypotheses. We now formulate five main results in the context of our five hypotheses.


⁹The apparent asymmetry around zero in the proportional deviations from fundamental values reflects the fact that the price cannot be negative, hence deviations cannot be lower than -100%, but there is no upper bound, except for the artificial one of 1000 that is unknown to the subjects until they hit it.

¹⁰The only exception is Market 2 of Group 2, in Tr. S, where one participant hits the upper-bound of 1000 and receives the message that his predictions have to be lower than this number. Note that this bound has been implemented for technical reasons, and none of the participants were aware of this bound, unless they reach it. This bound was reached 25 times out of the 18,170 forecasts elicited across all markets and subjects (which is about 0.1% of all forecasts).


¹¹We also note that a high price provides incentives to sell – and therefore to submit a lower prediction than the average of the group – a strategy that was also reported a few times.




(a) Treatment S: 100% of short-horizon forecasters



(b) Treatment M70: 70% of short-horizon forecasters, 30% of long-horizon forecasters



(c) Treatment M50: 50% of short-horizon forecasters, 50% of long-horizon forecasters



(d) Treatment L: 100% of long-horizon forecasters

Note: for instance, +100% means that the price is twice the fundamental.

Figure 4: Overview of the realized price levels in all experimental markets: the plots report deviations in percentage points from the fundamental value

4.2 Cross-treatment comparison

Table 3 reports cross-treatment comparisons of aggregate data. The first rows in particular show that there are significant cross-treatment differences regarding the price deviation (from fundamental), price volatility and, to a lesser extent, forecast dispersion: see Table 3 for definitions of these terms. These differences lead us to confirm the visual impression that the horizon of the forecasters matters for price dynamics and convergence towards the REE. The discrepancy between the realized price and the fundamental is strikingly lower in Tr. L than in Tr. S. Moreover, while the discrepancy from the REE is not statistically different between Tr. L and Tr. M50, prices are significantly closer to the fundamental price in those two treatments than in Tr. M70. These difference lead us to reject Hypothesis 1a in favor of Hypothesis 1b and report the following finding:

Finding 1 (Price convergence). *Increasing the share of long-horizon forecasters from zero percent to 30% and also from 30% to 50% significantly reduces price deviation from the REE.*

Turning now to Hypothesis 2, we find, as anticipated, that long-horizon forecasters have a stabilizing influence on prices. The price in Tr. S is significantly more volatile than in all other treatments, while price volatility is not significantly different between Tr. M50 and Tr. L. Those observations yield the following finding, consistent with Hypothesis 2:

Finding 2 (Price volatility). *Increasing the share of long-horizon forecasters from zero percent to 30% and also from 30% to 50% significantly reduces price volatility.*

Put conversely, increasing the share of short-horizon forecasters pulls the price level away from its fundamental value and increases its volatility. Our results suggest a *threshold effect* in the share of short-horizon forecasters on price convergence and volatility. Quite a large share of short-horizon forecasters (more than half of the market) is necessary to hinder stabilization and convergence.

To consider the issue of coordination between participants, we first look at the overall dynamics of the markets over time.¹² The trade volume significantly decreases in all treatments except Tr. S, and similar dynamics are observed for the within-participants forecast dispersion over time. Therefore, in the context of Hypothesis 3, we obtain the following result:

Finding 3 (Eventual coordination). *Participants' forecasts become more homogeneous over time and, consequently, the trade volume decreases over time, except in Tr. S.*

Regarding the differences of participants' degree of coordination between treatments, Table 3 shows that there is some evidence that forecasts are more homogeneous in treatments with short-horizon forecasters than in treatments with long-horizon forecasters. The forecast dispersion is higher in Tr. L than in Tr. S, and higher in Tr. M70 than in Tr. S. Within-horizon coordination can also be compared in the mixed treatments, especially in Tr. M50, where the groups of short-horizon and long-horizon forecasters are of the same size. For instance, in this treatment, looking at the first market of Gp. 4, or at all markets in Gp. 5 and 6, it is clear that short-horizon forecasts are closer

¹²A regression per treatment of the trade volume on the period leads to the coefficients -0.433, -0.348, -0.699 and 0.021 for, respectively, Tr. L, M50, M70 and S, with the associated p-values $< 2e-13$ except for Tr. S with 0.493. Similarly, with the forecast dispersion as a dependent variable, the same estimated coefficients are -0.004, -0.004, -0.005 and 6.185e-05 with the associated p-values of 0.020, 5.4e-06, 0.002 and 0.935.

	Diff-diff treatments					
	<i>L-S</i>	<i>L-M70</i>	<i>L-M50</i>	<i>M70-S</i>	<i>M50-S</i>	<i>M50-M70</i>
Price deviation^a (p-value)	-0.564 (0.000)	-0.111 (0.000)	0.012 (0.205)	-0.453 (0.000)	-0.576 (0.000)	-0.123 (0.000)
Price volatility^b (p-value)	-2.12 (0.000)	-0.111 (0.000)	-0.029 (0.315)	-2.013 (0.000)	-2.094 (0.000)	-0.082 (0.000)
Trade volume^c (p-value)	0.088 (0.000)	0.061 (0.000)	0.14 (0.000)	0.027 (0.000)	-0.052 (0.000)	-0.079 (0.000)
Forecast dispersion^d (p-value)	0.161 (0.015)	0.08 (0.53)	0.115 (0.047)	0.081 (0.005)	0.046 (0.532)	-0.035 (0.025)
EER (forecasts)^e (p-value)	-0.071 (0.231)	-0.026 (0.924)	-0.083 (0.452)	-0.045 (0.304)	0.012 (0.5)	0.057 (0.622)
EER (utility)^e (p-value)	0.01 (0.984)	-0.003 (0.492)	0.002 (0.614)	0.013 (0.663)	0.008 (0.754)	-0.01 (0.414)

Note: The table reports the differences between treatments, and the associated p-values of the one-sided Wilcoxon rank sum tests. In bold are the significant differences between treatments. K-S tests give the same predictions, except between treatments M70 and S regarding the volatility of the price, in which case the pair-difference becomes insignificant.

^a Average of the absolute price deviation from its fundamental value p_m , over all periods $t \geq 1$ of each market m , computed as $\frac{|p_{m,t} - p_m|}{p_m}$.

^b Relative price standard deviation computed over all periods t of each market m as $\frac{\sqrt{\text{Var}(p_{m,t})}}{\text{mean}(p_{m,t})}$.

^c Sum over all periods t and all markets m of exchanged assets among subjects in proportion of the steady-state endowment q_m , i.e. $\sum_{j=1}^{10} \left| \frac{q_{j,t} - q_{j,t-1}}{q_m} \right|$.

^d Relative standard deviation between subjects' forecasts $\frac{\sqrt{\text{Var}(p_{j,t}^e)_{j \in J}}}{\text{mean}(p_{j,t}^e)_{j \in J}}$, $t \geq 1$, averaged over all periods of each market.

^e Earnings Efficiency Ratio (EER) computed over all periods of each market, averaged over the 10 subjects as follows: (i) for the forecasting task, it is the average number of forecasting points earned in each market over the total amount of points possible in the market (1100 per period in case of perfect prediction); (ii) for the consumption task, it is the average number of utility points earned in each market over the total utility points earned at equilibrium (1081 per period).

Table 3: Cross-treatment statistical comparisons

to each other than the long-horizon ones (see again Figure 11 in Appendix B). This impression is confirmed by a statistical analysis: in this treatment, the average dispersion between short-horizon forecasters is 0.057, versus 0.163 among the long-horizon forecasters, and the difference is significant (p-value $< 2.2e - 16$). Therefore, in line with the two companion hypotheses 4 and 5, we find the following:

Finding 4 (Coordination and forecast horizons). *Long-horizon forecasters exhibit greater cross-section dispersion of price forecasts than do short-horizon forecasters.*

In addition, looking jointly at the trade volume and forecast dispersion rows in Table 3, we find the following:

Finding 5 (Trade volume and forecast horizons). *The higher the share of long-horizon forecasters in a market, the greater the cross-sectional dispersion of price forecasts and the higher the trade volume.*

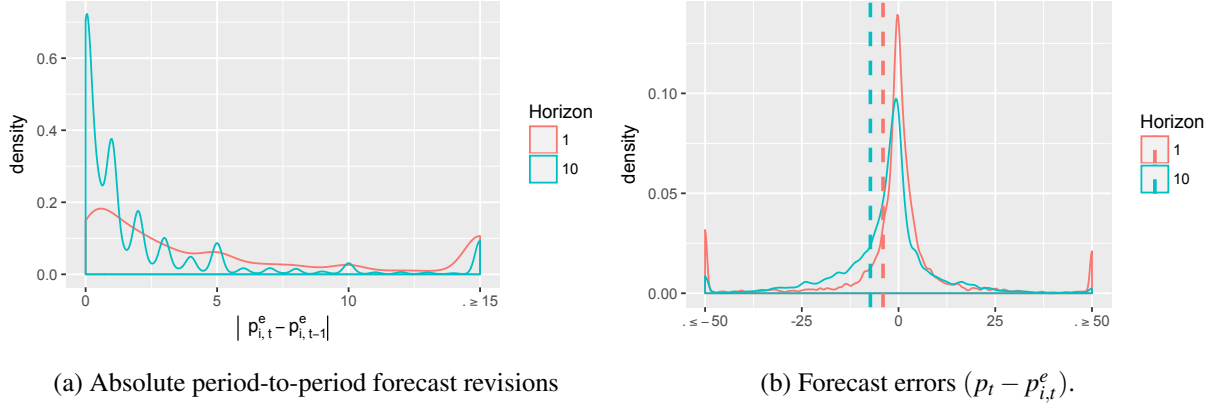
The fact that predictions at a longer horizon involve more disagreement between forecasters than short-run forecasts has been documented both in survey data (Bundick & Hakkio 2015) and in experimental work (Haruvy et al. 2007), albeit in non-self-referential environments. This corresponds also to what we observe in our experiment. One reason is the greater uncertainty involved in long-horizon forecasting. Taking a detour to look at individual forecast data reveals that long-horizon forecasts display a substantially higher amount of inertia and bias than short-horizon forecasts. To see that, the distribution of the period-to-period forecast revisions in Figure 5a reveals that the long-horizon forecasts remain unchanged more frequently (in one-quarter of the cases vs. 10%), and the forecast revisions are of smaller amplitude (with a median change of 6.828 vs. 15.962) than in the case of short-horizon forecasts (those differences being significant, with $p\text{-values} < 2.2e - 16$).¹³ Most interestingly, updates of long-horizon forecasts are strikingly biased towards round numbers (look at the spikes at each integer in the density plots in Figure 5a). This feature, absent from the short-horizon forecasts, might translate a greater uncertainty at long horizons than at short horizons, and hence the submission by participants of less accurate forecasts.

By studying questionnaire data, we find that the observed heterogeneity in forecasts may have two additional causes. First, some individuals were experimenting by occasionally trying to deviate from the average prediction of the group in order to trade (either submitting a very low prediction to sell, or a very high prediction to buy). This form of strategic behavior was reported by 16 participants, for instance in the first market of Tr. S, Group 3; in Tr. L, Group 4; in Tr. M70, Groups 2 and 5; and in Tr. M50, Groups 5 and 6. In those later groups, the resulting dispersion of individual forecasts is particularly salient as strategic deviations for the purpose of trade were reported by two participants to those markets as a frequent, and not occasional, strategy. Of course, such attempts turned out to be detrimental to the participants' payoff, as both the forecasting payoff (due to larger forecast errors) and the utility payoff (due to an irregular consumption pattern) drastically decreased. Second, some participants were confused despite the care that we devoted to explaining the instructions, and simply submitted random forecasts. Fortunately, this was the case for only very few participants, namely one subject in each of the following markets: Tr. L, Groups 1 and 6; Tr. M50, Groups 3 and 6; and M70, Group 3. However, strategic forecasts or confusion remain marginal phenomena in our experiment, and affect all treatments and groups indifferently.

Additional results There are two additional considerations of interest that are less directly connected to our hypotheses: first, we consider possible learning effects resulting from repetition; and second, we discuss the implications of performance metrics based on received utility versus forecast accuracy.

The repetition design of our experiment allows us to look for the presence of a *learning effect* in sequential markets with the same group of subjects. Repetitions of the seminal Smith et al. (1988) bubble experiment have shown that large deviations from fundamentals disappear if the market is repeated several times with the same participants (Dufwenberg et al. 2005).

¹³Those differences persist if we only consider the first periods of the markets, which indicates that the smaller and less frequent adjustments of long-run expectations are not a consequence of the relative stabilization of the price in those markets compared with those populated by short-horizon forecasters. They are rather due to the mechanical feature of the long-horizon forecasts as we elicit averages over future time spans that partly overlap.



Note: Left panel: the dashed lines represent the means of the distributions. The first 10 periods of each market are discarded to allow for a learning phase, especially for the long-horizon participants who were able to assess the accuracy of their first forecast and adjust upon their error from period 11 on only: 6,269 short-horizon observations and 4,971 long-horizon observations.

Figure 5: Statistics of individual forecasts, all markets pooled together

In our experiment, most figures in Appendix B instead convey the impression that price fluctuations do not decrease with participants' experience. On the contrary, a bubble can take several markets to arise, and price deviations from fundamental tend to amplify with market repetitions. This is especially the case in Groups 1, 2 and 4 of Tr. S. Deviations from fundamental tend also to increase with market repetition in Gp. 5 of Tr. L, for instance. Two linear regressions of the absolute deviations of the price and the forecasts from the REE on the order of the market confirms the absence of learning effect.¹⁴ Those observations lead us to conclude that there is no clear learning effect in repeated markets in terms of price deviations and convergence of forecasts towards the REE. However, we remark that by design, repeated markets had different fundamental prices, which makes it difficult to carry over knowledge from one market to the next.

Turning now to the role of performance metrics, we take a final look at Table 3 and consider the earnings of the participants in the experiment in the different treatments. While not directly connected to our hypotheses, incentives are an essential ingredient of theory testing using laboratory experiments. The data from the last two rows of Table 3 reveal that there is no noticeable difference in participants' earnings across treatments, whether based on utility or forecasting.

To explain this absence of payoff difference across treatments, it should be first recalled that Tr. S contains most of the large price oscillations, along which participants' forecasting earnings are very low, as their predictions cannot keep up with the realized prices. Forecasting over the next 10 periods, on the other hand, naturally involves greater uncertainty and forecasting errors, hence lower forecasting payoff. Quantitatively, Figure 5b shows that when compared with long-horizon forecasts, the distribution of short-horizon forecasts has a mean closer to zero but also places more weight on extreme values.

Hence, *average* forecasting efficiency losses are of comparable orders of magnitude across treatments, but arise from different causes. A similar argument goes for utility losses: as all participants initially receive the symmetric steady-state endowment (even if they do not know

¹⁴The coefficients are 4.984 and 6.6376 for, respectively, the price and the forecast distances as dependent variables.

it), trade leads to utility losses. This is the case along the large fluctuations observed in Tr. S (and to a lesser extent in Tr. M70) that have a detrimental distributional effect. In Tr. L and Tr. M50, the higher dispersion of individual forecasts generates trade, which leads to utility losses as well, as discussed above.

We next take a detailed look at the price data to shed more light on the convergence or divergence patterns reported in the different experimental treatments.

4.3 Assessing convergence to the REE

We now test formally whether convergence to the fundamental value occurs in the experimental markets. We follow the method presented in Noussair et al. (1995) which consists in estimating the value to which the price would converge asymptotically if a market were extrapolated into the infinite future. As the lengths of our markets differ and most are short due to the stochastic termination rule, this approach appears particularly suited to our experiment.


To test for convergence of market price towards its fundamental value, we estimate the following equation for each of the four treatments separately:

$$\frac{p_{g,m,t} - p_{g,m}}{p_{g,m}} = \frac{1}{t} \sum_{g=1}^6 \sum_{m \in \Omega_{M_g}} D_{g,m} b_{1,g,m} + \frac{t-1}{t} \sum_{g=1}^6 \sum_{m \in \Omega_{M_g}} D_{g,m} b_{2,g,m}, \quad (10)$$

with $p_{g,m,t}$ the realized market price in period t in Group $g \in \{1, \dots, 6\}$ and market m ; Ω_{M_g} the number of markets played by Group g ; $D_{g,m}$ a dummy taking the value 1 if the price comes from Group g and market m and 0 otherwise. $p_{g,m}$ is the fundamental value of the price in Group g and market m .

The estimated coefficients of these regressions provide summary statistics per market of the initial (coefficients $\hat{b}_{1,g,m}$) and asymptotic (coefficients $\hat{b}_{2,g,m}$) price values. If $\hat{b}_{2,g,m}$ is not significantly different from zero for a given market, we cannot reject the hypothesis of convergence in a strong sense towards the fundamental value of market m of Group g . If the inequality $|\hat{b}_{1,g,m}| > |\hat{b}_{2,g,m}|$ holds significantly, market m of Group g converges in a weak sense. The results of such an exercise are collected in Table 4. Figure 6 further reports the distributions of the estimated coefficients per treatment, while the details of the estimations are deferred to Appendix C.

Figure 6 reveals a net decrease in the estimated distances of the price to fundamental in Tr. M70, M50 and L (compare the paired box plots per treatment). However, such a decrease is not observed in Tr. S. The estimated final distances are particularly concentrated around zero in Tr. L, and even more strikingly in Tr. M50 (to see that, look at the height and the position of the second boxplot with respect to the zero line for each treatment). The econometric analysis indeed shows that weak convergence is obtained in all but one markets with long-horizon forecasters only (Tr. L), and the vast majority of them in the mixed treatment with 50% of each type of forecasters (Tr. M50). By contrast, less than two-thirds of the markets with a majority of short-horizon forecasters (Tr. M70) exhibit weak convergence, and even less than one-half of those with short-horizon forecasters only (Tr. S). Results on strong convergence show a similar pattern, as the share of markets exhibiting equilibration increases with the share of long-horizon forecasters.



Note: A box plot illustrates a distribution by reporting the four quartiles, with the thick line being the median, and the two whiskers being respectively $Q1$ and $Q4$ within the lower limit of $Q1 - 1.5(Q3 - Q1)$ and the upper limit of $Q3 + 1.5(Q3 - Q1)$. Outside that range, data points, if any, are outliers and represented by the dots. In the figure, each pair of box plots represents a treatment. The first box plot of each pair gives the distribution of the estimated initial values $\hat{b}_{1,g,m}$, the second one the estimated asymptotic values $\hat{b}_{2,g,m}$ in Eq. (10). The zero line represents zero deviation, i.e. convergence to fundamental.

Figure 6: Distribution of estimated initial ($\hat{b}_{1,g,m}$) and final ($\hat{b}_{2,g,m}$) price values in relative deviation from fundamental per treatment

	<i>Market level</i>	
	Weak convergence: $ \hat{b}_{1,g,m} > \hat{b}_{2,g,m} $	Strong convergence: $ \hat{b}_{2,g,m} = 0$
Tr. S	7/18 \simeq 39%	3/18 \simeq 17%
Tr M70	11/18 \simeq 61%	2/18 \simeq 11%
Tr. M50	10/13 \simeq 77%	3/13 \simeq 23%
Tr. L	13/14 \simeq 93%	4/14 \simeq 29%

Note: Number of markets exhibiting weak and strong convergence, as defined in the main text, over the total number of markets in each treatment, and corresponding fractions of converging markets.

Table 4: Fractions of markets exhibiting significant convergence

As a complement to Finding 1, we draw from this exercise the following insight:

Finding 6 (Statistical convergence). *Convergence to the REE is more frequently observed when the share of long-horizon forecasters is increased with respect to the share of short-horizon forecasters.*

This finding conforms with Hypothesis 1b; and further, based on Table 4, Hypothesis 1a is rejected.

We now examine factors that contribute to the convergence failures observed in Tr. M70 and Tr. S. From the individual plots reported in Appendix B, we make the observation that initial conditions in a given market may be correlated with terminal conditions in the previous market. Price patterns sometimes reproduce themselves from one market to another, without clear connection

to the fundamental, even though the information from the previous market is not displayed to the participants entering a new market. Looking at Tr. S, for instance, reveals that mispricing is carried out between the four markets of Gp. 3 or between the first two markets of Gp. 4. Price oscillations are also repeated between the second and the third markets in Gp. 5.

To formally assess this observation, we compute the correlation between the estimated initial price values $\{\hat{b}_{1,g,m}\}$ and the price levels prevailing in the preceding market. This correlation is 0.6644 (p-value 0.0000) when the previous prevailing prices is measured as the average price over the last 10 periods of the previous market, and is 0.3444 (p-value: 0.0057) when measured as simply the last observed price in the preceding market.¹⁵

We may also use the econometric analysis from Equation (10) to assess the role of the previously observed price levels in convergence failures. We do so by conducting an analysis of the variance of the estimated asymptotic coefficients $\{\hat{b}_{2,g,m}\}$ in terms of three factors: the fundamental value; the price in period one; and the last price in the previous market.¹⁶ Results are reported in Figure 7 and reveal a striking feature of the experiment: the asymptotic price values are almost entirely driven by the fundamental values in Tr. L and M50, while initial price levels and previously realized prices explain a considerable amount of the asymptotic price values in Tr. M70, and an even larger amount in Tr. S. This analysis confirms the dynamics reported in Table 4, and sheds further light on Hypotheses 1a and 1b: coordination of subjects' forecasts on a wrong anchor, namely past observed prices, is responsible for the lack of convergence observed in Tr. M70 and Tr. S and, hence, the rejection of Hypothesis 1a.

Finding 7 (Fundamental and non-fundamental factors).

- (i) *When the share of long-horizon forecasters is large enough, the asymptotic market price is driven by fundamentals only.*
- (ii) *By contrast, if short-horizon forecasters dominate, the asymptotic market price becomes partly driven by non-fundamental factors, in particular past observed price levels.*

To shed some light on the causal mechanisms behind those results, we now seek to understand how the participants formed their price forecasts and how those individual behaviors connect to the observed market prices in the experiment.

4.4 Participants' forecasts and aggregate outcomes

At the end of the experiment, participants were asked to describe in a few words their strategies. Analysis of the answers makes clear that the vast majority of participants, aside from strategic deviations for trading purposes, made use of past prices. The observation that expectations about future market prices depend on past trends has also found wide support in the experimental literature – see the early evidence reported in Smith et al. (1988) and Andreassen & Kraus (1990), and more recent evidence found in Hommes et al. (2005) and Haruvy et al. (2007); see also the empirical literature, starting from early contributions such as Shiller (1990).

¹⁵For the first markets, we took 50 as the previous value because it corresponds to the middle point of the empty price plot that the participants observe before entering their first forecast; see the screen shots in Figure 2. Removing the first markets results in fewer data points, but the correlation pattern persists.

¹⁶The variance decomposition was done using the Fourier amplitude sensitivity test.

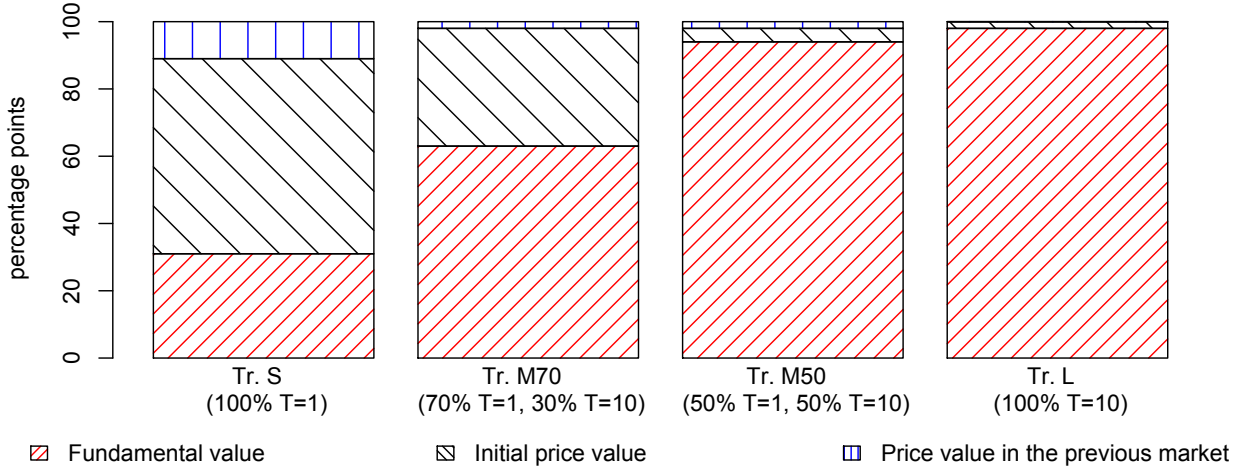


Figure 7: Contribution to the variance of the estimated final values $\hat{b}_{2,g,m}$

To estimate the dependence of participants' forecasts on past data, we begin with the following class of simple, yet flexible, agent-specific forecasting models:

$$p_{j,t}^e = \beta_0 + \beta_1 p_{t-1} + \beta_2 p_{t-2} + \delta_1 p_{j,t-1}^e. \quad (11)$$

Clearly, participants could have paid attention to more lags of the observable variables – and actually a few reported to have done so – but most referred to at most the last two of prices in their strategy. It should also be noted that including lagged expectations is an indirect way of accounting for the influence of additional lags of prices.¹⁷

We focus towards the following three special cases of the forecasting model (11):

- Naive expectations:** $\beta_0 = \beta_2 = \delta_1 = 0$ and $\beta_1 = 1$
- Adaptive expectations:** $\beta_0 = \beta_2 = 0$, $\beta_1 \in (0, 1)$, and $\beta_1 + \delta_1 = 1$
- Trend-chasing expectations:** $\beta_0 = \delta_1 = 0$, $\beta_1 > 1$, and $\beta_1 + \beta_2 = 1$.

Under naive expectations, agents forecast as $p_{j,t}^e = p_{t-1}$. Although we label this “naive,” these are the optimal forecasts if the price process follows a random walk, and naive expectations are therefore “nearly rational” when prices follow a near-unit root process. Under adaptive expectations, agents forecast as

$$p_{j,t}^e = p_{j,t-1}^e + \beta_1 (p_{t-1} - p_{j,t-1}^e). \quad (12)$$

This rule is known to be optimal if the price process is the sum of a random walk component and a white noise component, i.e. a mixture of permanent and transitory shocks: see Muth (1961). From the viewpoint of adaptive learning, (12) comprises the constant-gain algorithm for estimating the mean of the price process. We further note that, in this case, $p_{j,t}^e$ is the sum of past prices geometrically discounted at rate $1 - \beta_1$. Finally, we observe that if all agents use the same adaptive forecasting rule then, by combining (12) with the implications of market clearing given by equation (8), we observe that the resulting price process is always stable.

¹⁷In principle, this forecasting model could generate negative price forecasts, in which case it would be natural for agents to impose a non-negativity condition.

Under trend-chasing expectations, agents forecast as

$$p_{j,t}^e = p_{t-1} + \phi(p_{t-1} - p_{t-2}), \text{ where } \phi = \beta_1 - 1 > 0. \quad (13)$$

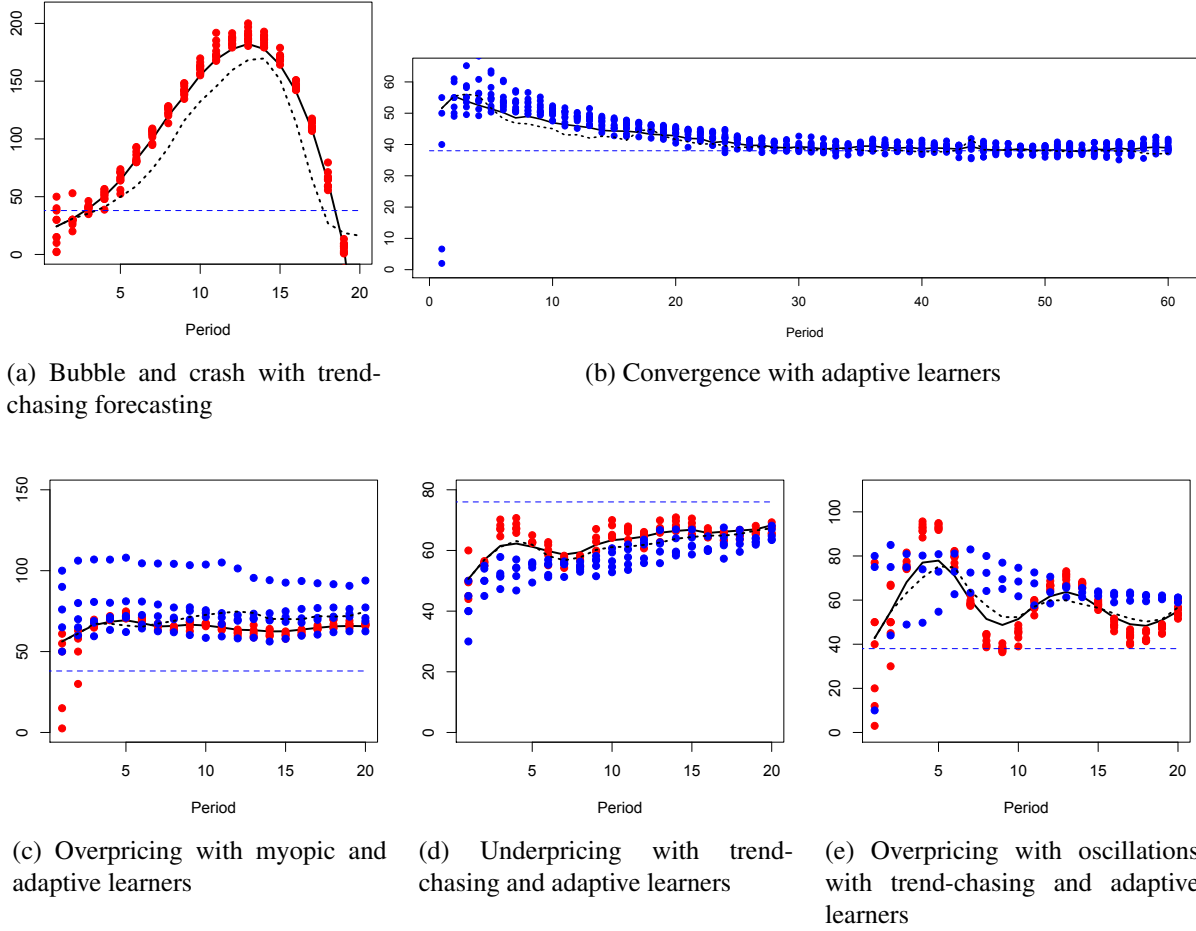
This rule performs well in bubble-like environments in which price changes are persistent. In fact, this forecasting rule would be optimal if the first difference in prices were to follow an AR(1) process (which would be stationary if $\phi < 1$). Intuitively, agents are forecasting based on the assumption that the proportion ϕ of last period's price change will continue into the future.¹⁸ Finally, assuming homogeneity and combining (13) with the implications of market clearing given by Equation (8), we observe that the resulting price process is stable provided that $\phi\xi < 1$; additionally, given any ξ , there is value $\bar{\phi} < 1$ so that $\phi \in (\bar{\phi}, 1)$ implies that the prices dynamics are stable and cyclical, i.e. governed by complex eigenvalues of modulus less than one.

Figure 8 illustrates the potential for these simple forecasting rules to explain the price data in five different experimental markets: see graphs (a) to (e). In each case, the dashed horizontal line is the fundamental price and the dotted line corresponds to the realized price in the experimental market. The dots correspond to simulated price forecasts and the solid line provides the implied, simulated market prices. To construct the simulated price forecasts, a parametric specification of a particular forecasting model is chosen, and, for each agent, is initialized using their forecasts in the first two periods of the experiment. In each subsequent period, the agents' forecasts are determined using the parametric forecasting model and previously determined simulated prices, to which a small, idiosyncratic white noise shock is added. We note that, in all cases, the lines identifying the simulated and experimental price time series are close to each other. Figure 8 also highlights the systematic differences between treatments and horizons in belief formation and links them to the observed price patterns.

Graph (a) provides an example of trend-chasing behavior that emerged from treatment S. The simulated data are based on the parameterization $\phi = \beta_1 - 1 = 0.3$. This market strikingly illustrates the possibility of a bubble and crash being generated by trend-chasing forecast rules. Graph (b) gives an example of adaptive expectation associated to treatment L. The parameterization is $\beta_1 = 0.7$ and $\delta_1 = 0.3$, and it illustrates apparent convergence to the fundamental price, as predicted by our theoretical results. Graphs (c) to (e) provide outcomes from markets with mixtures of horizons and corresponding forecast models. Graphs (c) and (d) correspond to results from treatment M50, in which the short-horizon forecasters are naive and trend-following, respectively, and the long-horizon forecasters form expectations adaptively. The simulated price paths depend on initial conditions, i.e. on the individual-specific initial forecasts in each market, and this is one significant factor in the different dynamics exhibited in these graphs. The presence of short-horizon forecasters accentuates the model's positive feedback and allows for a variety of near-rational market equilibrium outcomes. In graph (c) this takes the form of a persistent departure from fundamentals, while in graph (d) the short-horizon trend-chasers generate cyclic dynamics as well as apparent convergence. Finally, graph (e) corresponds to M70 with short-horizon trend-chasing forecasters and with long-horizon forecasters forming expectations adaptively. Here the cyclicity arising from the trend-followers is even more pronounced, while the presence of only 30% long-horizon types appears insufficient to impart convergence.

Using a step-by-step elimination process, we examined individual participant-level forecast data, pooled across markets, and looked for simplifications of the model (11) in an attempt to

¹⁸Under adaptive learning, agents might actually update over time their estimates of this proportion ϕ as new data become available.



Note: The blue dashed line is the fundamental price, the dotted lines represent the prices in the experimental markets, the dots and the solid lines are the simulated forecasts and prices. The forecasts in the first two periods are taken from the experiment. An idiosyncratic shock distributed as $\mathcal{N}(0,2)$ is added then in each subsequent period to the forecasts. Fig. (a): Tr. S, Gp. 1, Market 1, trend-chasing forecasting model with $\beta_1 = 1.3$ (see Eq. (11) below); Fig. (b): Tr. L, Gp. 2, Market 1, convergence with adaptive learning, $\delta_1 = 0.3$; Fig. (c): Tr. M50, Gp. 6, Market 1, overpricing with static short-horizon forecasters ($\beta_1 = 1$) and adaptive long-horizon forecasters ($\delta_1 = 0.1$); Fig. (d): Tr. M50, Gp. 1, Market 2, trend-chasing short-horizon forecasters ($\beta_1 = 1.3$) and adaptive long-horizon forecasters ($\delta_1 = 0.1$); Fig. (e): Tr. M70, Gp. 6, Market 1, trend-chasing short-horizon forecasters ($\beta_1 = 1.75$), adaptive long-horizon forecasters ($\delta_1 = 0.1$).

Figure 8: Simulated versus experimental time series for selected price patterns

determine if, and to what extent, participants used one of the three simple rules listed above. We found, considering all 240 participant forecast series,¹⁹ that more than one-half of the short-horizon participants had forecasts that were consistent with trend-chasing rules, and more than one-third of the long-horizon participants had forecasts that were consistent with adaptive expectations.²⁰

¹⁹The experiments included 18 treatment S, 14 treatment L, 18 treatment M70, and 13 treatment M50 markets, with 10 participants in each market, for a total of 630 market-participant forecast series.

²⁰For the vast majority of the participants (212 of 240), the step-by-step elimination process leads to a forecasting model in which at least one variable other than the intercept is significant. Also, the average R^2 is high for each treatment (ranging from an average of 0.884 in Tr. L to 0.962 in Tr. M50), which confirms the ability of the simple

The estimated coefficients $\hat{\beta}_0, \hat{\beta}_1, \hat{\beta}_2$ and $\hat{\delta}_1$ from (11) for each participant are illustrated in Figure 9: smaller, solid triangles identify long-horizon forecasters and larger triangles identify short-horizon forecasters.²¹ Panel 9a shows a scatterplot of the components $\hat{\beta}_1$ and $\hat{\beta}_2$ for each participant. Under the restrictions that $\hat{\beta}_0 = \hat{\delta}_1 = 0$ and $\hat{\beta}_1 > 1$, the trend-chasing model aligns with the constellation of points on the part of the downward-sloping dashed line that lies within the shaded region. In line with the previous discussion, it can be seen that a substantial number of the short-horizon points in panel 9a lie on, or close to, the trend-chasing constellation. The trend-chasing restrictions cannot be rejected for 56% of the short-horizon forecasters.

Panel 9b shows the corresponding scatterplot of the components $\hat{\beta}_1$ and $\hat{\delta}_1$. Under the assumptions that $\hat{\beta}_0 = \hat{\beta}_2 = 0$ and $0 < \hat{\beta}_1 < 1$, the adaptive-expectations model aligns with the constellation of points on the part of the downward-sloping dashed line that lies within the shaded region in panel 9b.²² Analogous to behavior exhibited by short-horizon forecasters, a substantial number of the long-horizon points in panel 9b lie on, or close to, the adaptive-expectations constellation. The adaptive-expectations restrictions cannot be rejected for more than one-third of the participants in long-horizon treatments. We summarize these findings as follows:

Finding 8 (Individual forecast behaviors). *Short-horizon and long-horizon forecasters display different forecasting behaviors:*

- (i) *More than one-half of the short-horizon forecasters form forecasts consistent with trend-chasing behavior.*
- (ii) *More than one-third of the long-horizon forecasters form forecasts consistent with adaptive expectations.*

These results are consistent with the earlier hypotheses 1b and 2 in the sense that trend-chasing behavior tends to preclude, and adaptive learning tends to impart convergence to REE.

As noted above, adaptive learning is more general than the special cases that we have emphasized. In this sense, most forecasters across treatments make predictions consistent with some form of adaptive learning. Yet, Finding 8 suggests that long-horizon forecasters use a greater variety of forecasting models than do short-horizon forecasters, which supports Hypothesis 4 and the related discussion in Section 4.2.

The plots in Figure 9 include estimates that do not appear, even after accounting for statistical significance, to align with any of the special cases identified above. For example, there are points in panel 9b with $\hat{\delta}_1 > 0$, so that the corresponding subject is not a trend-chaser, and $\hat{\delta}_1 + \hat{\beta}_1 > 1$, so that the corresponding subject does not have adaptive expectations. There are several possible explanations. First, it is possible that some subjects use less parsimonious forecasting rules than are captured by the class (11). Second, given that most subjects participated in multiple markets, it is quite possible that some of these participants used different rules in different markets. Our pooling estimation strategy does not account for this. Third, in general, under the adaptive learning approach, in addition to the intercept, the other coefficients in the subjects' forecasting rules may evolve over time to reflect recent patterns of the data. Finally, we note that if ξ is near one,

class of rules (11) to capture the main features of participants' behaviors.

²¹A few of the participants' estimated coefficients lie outside the ranges chosen for Figure 9.

²²We note that naive expectations corresponds to limiting cases (i.e. $\hat{\beta}_1 \rightarrow 1$) of both trend-chasing and adaptive-expectations forecasting models.

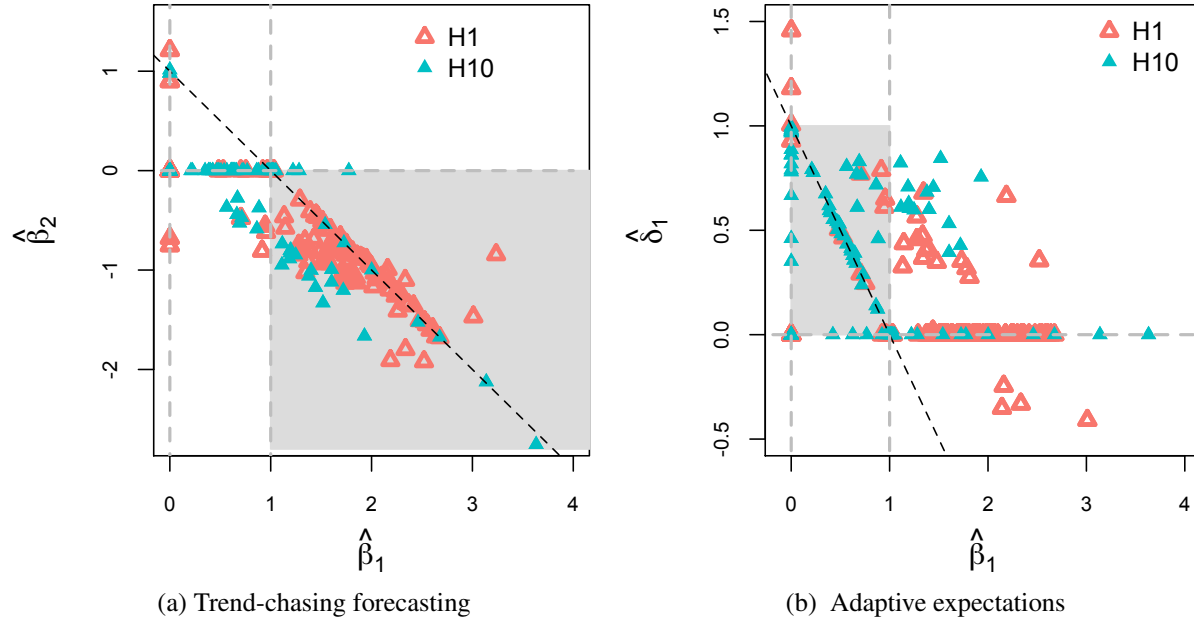


Figure 9: Distribution of the estimated coefficients of Equation (11) for the 240 subjects

then *any* forecast of the deviation of price from fundamentals is nearly self-fulfilling; this point is particularly relevant for treatment S.

Findings 6 to 8 shed further light on the impact of the feedback parameter ξ on price dynamics. In long-horizon treatments, this feedback is smaller and empirically is associated with adaptive expectations leading to more price stability and more frequent convergence to the fundamental price. By contrast, short-horizon treatments with ξ near one are empirically associated with trend-chasing behavior and are consistent with persistence departures from the fundamental price.

5 Conclusions

We investigate how the expectation horizons affect the price dynamics in a self-referential asset market, that is, in a market in which price expectations of the traders in turn determine the asset price. To do so, we proceed in two stages. Firstly, we develop a model with boundedly rational heterogeneous agents and learning to derive theoretical predictions of the effects of the planning horizon on the market price. Secondly, we provide an empirical test of our predictions by implementing this asset market in a lab experiment, in which we elicit price forecasts at different horizons from human subjects, who then trade accordingly.

Our main results may be summarized as follows. Markets populated by only short-horizon forecasters fail to converge to the REE and exhibit mispricing, sometimes with large and prolonged deviations from fundamentals, which stands in contradiction to rational expectations and most theories of learning. By contrast, in line with our established theoretical predictions, we show that even a relatively modest share of long-horizon forecasters operating on the market is enough to stabilize the market and induce convergence to the REE.

These results bring up further elements of discussion. In our design, the incentive to smooth out consumption and the constant dividend flows discourage capital-gain-seeking behaviors. The absence of incentives to large capital gains limits the possibility that the observed price fluctuations may be due to rational or informed traders betting on future crowd behaviors (De Long et al. 1990). Those large price oscillations are hard to reconcile with rational bubbles because the observed price growth rates appear disconnected from the value of the discount factor (Blanchard & Watson 1982). Our design also rules out the possibility of heterogeneous beliefs about future dividends, which can constitute a cause of deviations from fundamentals. Nor do those fluctuations arise from confusion about how the market works, as the vast majority of participants reported to understand their experimental task. Finally, we can exclude the role of liquidity in mispricing, as this is kept constant across all treatments and markets.

Thus, speculation and confusion are not responsible for departures from REE. Instead we find that the coordination of forecasts on unstable belief formation processes and the anchorage of individual expectations on non-fundamental factors are responsible for mispricing in short-horizon markets. Instability of this type has been noted in the adaptive learning literature. Our experiment shows that this theoretical outcome constitutes an empirical concern as well.

Admittedly, several previous experimental studies eliciting only one-step-ahead forecasts have reported similar coordination failures, albeit in simpler designs, but our results demonstrate the robustness of such failures to a number of important dimensions along which our study clearly differs: (i) a different framing that does not use the vocabulary of speculative asset markets; (ii) the emulation of a stationary and infinite environment that induces discounting with a stochastic ending after which the value of the asset drops to zero; (iii) the repetition of series of markets with each group of participants, with different fundamentals but the same structure; (iv) a payoff scheme that incentivizes participants to smooth their consumption of the dividend; and (v) a higher cognitive load in the theoretically more stable treatments, i.e. the treatments that elicit long-horizon forecasts.

These robustness results are particularly interesting for two main reasons. First, those dimensions are major ingredients of the underlying theoretical framework, but were absent from previous experimental work. Second, any of those dimensions taken individually could have been thought *a priori* to be able to dampen coordination failures. Even if those dimensions are present all together, we show that this is not the case.

Furthermore, the stabilizing role of long-horizon planners is not obvious, because long-horizon forecasting is more challenging than short-horizon forecasting: participants have to think in terms of average over a given number of future periods, instead of single data points, and the observability of the forecast errors and the resulting feedback from the experimental environment is delayed and postponed to the end of the forecast horizon, once the average price is realized. Long-horizon forecasters also display more disagreements between each other, as a result of a greater uncertainty. Despite these obstacles, they contribute to stabilize the market.

One reason is that, within our model, increasing the share of long-horizon planners decreases the risk of near-unit root price dynamics, which simplifies the coordination problem of the agents despite the more challenging forecasting task. Another stabilizing force relates to the nature of the forecasting task over a long horizon: long-horizon forecasters do revise their forecasts less often and less dramatically than the short-horizon participants. Such a conservative behavior mitigates the fluctuations of the price, especially at the beginning of the markets, when participants are still learning how to form accurate predictions and coordinate. Given the dependence on past

price patterns observed in the unstable markets with short-horizon forecasters only, the forecast inertia displayed by long-horizon participants undoubtedly contributes to eliminating large price fluctuations and helping them learn adaptively the fundamental price.

One interesting insight from this result is that heterogeneity in behavior is not necessarily detrimental to market stabilization. On the contrary, in our setup, the greater heterogeneity observed in the long-horizon forecasts contributes to breaking the coordination of participants' beliefs on non-fundamental factors, such as past price trends, or incorrect anchors, such as initial price levels. Heterogeneity in behavior is therefore a crucial component of learning and convergence to equilibrium.

A second interesting insight from our study is that the type of behavioral rules that agents use seems to depend on the environment and the nature of the task that they are confronted with. Trend-chasing behavior is widely observed among short-horizon forecasters while adaptive behavior well describes long-run predictions. This suggests that boundedly rational agents may not be easily classified into different invariant behavioral types.

Finally, our study conveys theoretical implications for macro-finance models with learning (and possibly heterogeneous) agents. We show that the forecasting and decision horizon of agents plays a central role in the determination of asset prices. This indicates that the forecasting horizon of agents must be taken into account when assessing economic models and designing policy.

References

- Adam, K., Marcet, A. & Nicolini, J. P. (2016), 'Stock Market Volatility and Learning', *Journal of Finance* **71**(1), 33–82.
- Andreassen, P. B. & Kraus, S. J. (1990), 'Judgmental extrapolation and the salience of change', *Journal of Forecasting* **9**(4), 347–372.
- Asparouhova, E., Bossaerts, P., Roy, N. & Zame, W. (2016), 'Lucas' In The Laboratory', *Journal of Finance* **71**(6), 2727–2780.
- Blanchard, O. J. & Watson, M. W. (1982), Bubbles, Rational Expectations and Financial Markets, in P. Wachtel, ed., 'Crisis in the Economic and Financial Structure', Lexington, MA.
- Branch, W. & Evans, G. W. (2011), 'Learning about Risk and Return: A Simple Model of Bubbles and Crashes', *American Economic Journal: Macroeconomics* **3**(3), 159–191.
- Branch, W., Evans, G. W. & McGough, B. (2012), Finite Horizon Learning , in T. Sargent & J. Vilmunen, eds, 'Macroeconomics at the Service of Public Policy', Oxford University Press.
- Bullard, J., Evans, G. W. & Honkapohja, S. (2010), 'A model of near-rational exuberance', *Macroeconomic Dynamics* (2), 166–188.
- Bundick, B. & Hakkio, C. S. (2015), 'Are longer-term inflation expectations stable?', *Macro Bulletin* pp. 1–3. Federal Reserve Bank of Kansas City, March 9.
- Charness, G. & Genicot, G. (2009), 'Informal Risk Sharing in an Infinite-Horizon Experiment', *Economic Journal* **119**(537), 796–825.

- Crockett, S., Duffy, J. & Yehuda, I. (2019), ‘An Experimental Test of the Lucas Asset Pricing Model’, *Review of Economic Studies* **86**, 627–667.
- De Long, J. B., Shleifer, A., Summers, L. H. & Waldmann, R. J. (1990), ‘Positive Feedback Investment Strategies and Destabilizing Rational Speculation’, *The Journal of Finance* **45**(2), 379–395.
- Duffy, J. (2016), Macroeconomics: A Survey of Laboratory Research, in J. H. Kagel & A. E. Roth, eds, *The Handbook of Experimental Economics, Volume 2*, Introductory Chapters, Princeton University Press.
- Dufwenberg, M., Lindqvist, T. & Moore, E. (2005), ‘Bubbles and Experience: An Experiment’, *American Economic Review* **95**(5), 1731–1737.
- Evans, G. W. & Honkapohja, S. (2001), *Learning and Expectations in Macroeconomics*, Princeton University Press.
- Evans, G. W. & Honkapohja, S. (2006), ‘Monetary Policy, Expectations and Commitment’, *Scandinavian Journal of Economics* **108**(1), 15–38.
- Evans, G. W. & McGough, B. (2018), ‘Learning to Optimize’. mimeo, University of Oregon.
- Fréchette, G. & Yuksel, S. (2017), ‘Infinitely Repeated Games in the Laboratory: Four Perspectives on Discounting and Random Termination’, *Experimental Economics* **20**(2), 279–308.
- Haruvy, E., Lahav, Y. & Noussair, C. (2007), ‘Traders’ expectations in asset markets: Experimental evidence’, *American Economic Review* **97**(5), 1901–1920.
- Hirota, S., Huber, J., Stock, T. & Sunder, S. (2015), Investment Horizons and Price Indeterminacy in Financial Markets, Cowles Foundation Discussion Papers 2001, Cowles Foundation for Research in Economics, Yale University.
- Hirota, S. & Sunder, S. (2007), ‘Price bubbles sans dividend anchors: Evidence from laboratory stock markets’, *Journal of Economic Dynamics and Control* **31**(6), 1875–1909.
- Hommes, C., Sonnemans, J., Tuinstra, J. & van den Velden, H. (2005), ‘Coordination of Expectations in Asset Pricing Experiments’, *Review of Financial Studies* **33**, 955–908.
- Hommes, C. & Zhu, M. (2014), ‘Behavioral learning equilibria’, *Journal of Economic Theory* **150**(C), 778–814.
- Lucas, R. E. (1978), ‘Asset Prices in an Exchange Economy’, *Econometrica* **46**, 1429–1445.
- Marimon, R., Spear, S. E. & Sunder, S. (1993), ‘Expectationally Driven Market Volatility: An Experimental Study’, *Journal of Economic Theory* **61**(1), 74–103.
- Muth, J. F. (1961), ‘Rational Expectations and the Theory of Price Movements’, *Econometrica* **29**(3), 315–335.
- Newey, W. K. & West, K. D. (1994), ‘Automatic Lag Selection in Covariance Matrix Estimation’, *Review of Economic Studies* **61**, 631–653.

- Noussair, C., Plott, C. & Riezman, R. (1995), 'An Experimental Investigation of the Patterns of International Trade', *American Economic Review* **85**(3), 462–491.
- Preston, B. (2005), 'Learning about Monetary Policy Rules when Long-Horizon Forecasts Matter', *International Journal of Central Banking* **1**(2), 81–126.
- Roth, A. E. & Murnighan, J. K. (1978), 'Equilibrium behavior and repeated play of the prisoner's dilemma', *Journal of Mathematical Psychology* **17**, 189–198.
- Shiller, R. J. (1990), 'Speculative Prices and Popular Models', *Journal of Economic Perspectives* **4**(2), 55–65.
- Smith, V. L., Suchanek, G. L. & Williams, A. W. (1988), 'Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets', *Econometrica* **56**(5), 1119–1151.

A Finite-horizon learning in the Lucas model

A.1 Proof of Proposition 1

Without loss of generality, let $t = 0$. Let $Q_k = p_k q_k$, and $R_k = p_{k-1}^{-1}(p_k + y_k)$, so that $c_k + Q_k = R_k Q_{k-1}$. The associated first-order condition (FOC) is $u'(c_k) = \beta R_{k+1} u'(c_{k+1})$. Linearizing the FOC and iterating gives

$$\begin{aligned} dc_k &= dc_{k-1} + \frac{(1-\beta)Q}{\sigma} dR_k, \text{ or} \\ dc_k &= dc_0 + \frac{(1-\beta)Q}{\sigma} \sum_{m=1}^k dR_m. \end{aligned} \quad (14)$$

Linearizing $c_k + Q_k = R_k Q_{k-1}$ and iterating gives

$$\begin{aligned} dc_k &= R dQ_{k-1} - dQ_k + Q dR_k, \text{ or} \\ \sum_{k=0}^T \beta^k dc_k &= R dQ_{-1} - \beta^T dQ_T + Q \sum_{k=0}^T \beta^k dR_k, \end{aligned} \quad (15)$$

where $R = \beta^{-1}$. Combining (14) and (15), we get

$$\sum_{k=0}^T \beta^k \left(dc_0 + \frac{(1-\beta)Q}{\sigma} \sum_{m=1}^k dR_m \right) = R dQ_{-1} - \beta^T dQ_T + Q \sum_{k=0}^T \beta^k dR_k,$$

or

$$\left(\frac{1-\beta^{T+1}}{1-\beta} \right) dc_0 = R dQ_{-1} - \beta^T dQ_T + Q \sum_{k=0}^T \beta^k dR_k - \frac{(1-\beta)Q}{\sigma} \sum_{k=0}^T \beta^k \sum_{m=1}^k dR_m.$$

Now notice

$$\sum_{k=0}^T \beta^k \sum_{m=1}^k dR_m = \sum_{k=1}^T \left(\frac{\beta^k - \beta^{T+1}}{1-\beta} \right) dR_k.$$

It follows that

$$dc_0 = \frac{1-\beta}{1-\beta^{T+1}} \left(R dQ_{-1} - \beta^T dQ_T + Q dR_0 + \frac{Q}{\sigma} \sum_{k=1}^T \psi(k, T) dR_k \right), \quad (16)$$

where $\psi(k, T) = \beta^k(\sigma - 1) + \beta^{T+1}$.

The linearized flow constraint provides

$$dQ_0 = R dQ_{-1} + Q dR_0 - dc_0.$$

Combine with 16 to get

$$\begin{aligned} dQ_0 &= R \left(\frac{\beta(1-\beta^T)}{1-\beta^{T+1}} \right) dQ_{-1} + Q \left(\frac{\beta(1-\beta^T)}{1-\beta^{T+1}} \right) dR_0 \\ &\quad + \left(\frac{\beta^T(1-\beta)}{1-\beta^{T+1}} \right) dQ_T - \left(\frac{1-\beta}{1-\beta^{T+1}} \right) \left(\frac{Q}{\sigma} \right) \sum_{k=1}^T \psi(k, T) dR_k, \end{aligned}$$

or

$$dQ_0 = \phi_0(T) dQ_{-1} + \phi_1(T) dR_0 + \phi_2(T) dQ_T + \phi_3(T) \sum_{k=1}^T \psi(k, T) dR_k.$$

Next, linearize the relationship between prices, dividends and returns:

$$dR_k = \frac{1}{p}(dp_k + dy_k - Rdp_{k-1}).$$

Since $\beta R = 1$, we may compute

$$\begin{aligned} \sum_{k=1}^T \beta^k (dp_k - Rdp_{k-1}) &= \beta^T dp_T - dp_0 \\ \sum_{k=1}^T (dp_k - Rdp_{k-1}) &= dp_T - Rdp_0 - R(1 - \beta) \sum_{k=1}^{T-1} dp_k. \end{aligned}$$

It follows that $\sum_{k=1}^T \psi(k, T) dR_k$

$$\begin{aligned} &= \frac{1}{p} \sum_{k=1}^T \psi(k, T) dy_k + \frac{\sigma - 1}{p} \sum_{k=1}^T \beta^k (dp_k - Rdp_{k-1}) + \frac{\beta^{T+1}}{p} \sum_{k=1}^T (dp_k - Rdp_{k-1}) \\ &= \frac{1}{p} \sum_{k=1}^T \psi(k, T) dy_k + \frac{\sigma - 1}{p} (\beta^T dp_T - dp_0) + \frac{\beta^{T+1}}{p} \left(dp_T - Rdp_0 - R(1 - \beta) \sum_{k=1}^{T-1} dp_k \right) \\ &= \frac{1}{p} \sum_{k=1}^T \psi(k, T) dy_k + \frac{\beta^T}{p} (\sigma - 1 + \beta) dp_T - \frac{1}{p} (\sigma - 1 + \beta^T) dp_0 - \frac{\beta^t (1 - \beta)}{p} \sum_{k=1}^{T-1} dp_k. \end{aligned}$$

Finally, assuming dividends are constant, and using these computations, together with $dQ_k = pdq_k + qdp_k$, we may write the demand for trees as

$$dq_0 = \theta_0(T) dq_{-1} + \theta_*(T) dp_{-1} + \theta_1(T) dp_0 + \theta_2(T) dq_T + \theta_3(T) \sum_{k=1}^{T-1} dp_k + \theta_4(T) dp_T,$$

where

$$\begin{aligned} \theta_0(T) &= \phi_0(T) &= R \left(\frac{\beta(1-\beta^T)}{1-\beta^{T+1}} \right) \\ \theta_*(T) &= \frac{\phi_0(T)q}{p} - \frac{\phi_1(T)}{\beta p^2} &= 0 \\ \theta_1(T) &= -\frac{q}{p} + \frac{\phi_1(T)}{p^2} - \frac{\phi_3(T)}{p^2} (\sigma - 1 + \beta^T) &= -\frac{(1-\beta)q}{(1-\beta^{T+1})p\sigma} (1 - \beta^T) \\ \theta_2(T) &= \phi_2(T) &= \frac{(1-\beta)\beta^T}{1-\beta^{T+1}} \\ \theta_3(T) &= -\frac{(1-\beta)\beta^T}{p^2} \phi_3(T) &= \frac{(1-\beta)^2 \beta^T}{1-\beta^{T+1}} \frac{q}{p\sigma} \\ \theta_4(T) &= \phi_2(T) \frac{q}{p} + \frac{\phi_3(T)}{p^2} ((\sigma - 1)\beta^T + \beta^{T+1}) &= \theta_3(T) \end{aligned}$$

The result follows. ■

A.2 Computations for Section 2.2

We begin with a distinct short proof that when $dp_0 = 0$, time zero consumption demand, dc_0 , depends only on the sum of future prices. To this end, set $dq_{-1} = dp_0 = 0$, and let dq_t be given. The linearized budget constraints yield

$$\begin{aligned} dc_0 + pdq_0 + qdp_0 &= (p+y)dq_{-1} + qdp_0, & \text{or } dc_0 &= -pdq_0 \\ dc_1 + pdq_1 + qdp_1 &= (p+y)dq_0 + qdp_1, & \text{or } \beta dc_1 &= pdq_0 - \beta pdq_1 \\ dc_2 + pdq_2 + qdp_2 &= (p+y)dq_1 + qdp_2, & \text{or } \beta^2 dc_2 &= pdq_1 - \beta^2 pdq_2 \\ &\vdots & &\vdots \\ dc_t + pdq_t + qdp_t &= (p+y)dq_{t-1} + qdp_t, & \text{or } \beta^t dc_t &= pdq_{t-1} - \beta^t pdq_t. \end{aligned}$$

Summing, we obtain

$$\sum_{n=0}^t \beta^n dc_n = -\beta^t pdq_t. \quad (17)$$

The agent's FOC may be written $p_n u'(c_n) = \beta(p_{n+1} + y)u'(c_{n+1})$, which linearizes as

$$dc_{n+1} = dc_n + \psi(\beta dp_{n+1} - dp_n) \equiv dc_n + \psi \Delta p_{n+1},$$

where $\psi = (\sigma\beta)^{-1}q(1-\beta)$ and $\Delta p_{n+1} \equiv \beta dp_{n+1} - dp_n$. Backward iteration yields $dc_n = dc_0 + \psi \sum_{m=1}^n \Delta p_m$, which may be imposed into (17) to obtain

$$\sum_{n=0}^t \beta^n dc_0 + \psi \sum_{n=1}^t \beta^n \sum_{m=1}^n \Delta p_m = -\beta^t pdq_t. \quad (18)$$

Now a simple claim:

Claim. $\sum_{n=1}^t \beta^n \sum_{m=1}^n \Delta p_m = \beta^{t+1} \sum_{n=1}^t dp_n$.

The argument is by induction. For $t = 1$, use $dp_0 = 0$ to get the equality. Now assume it holds for $t - 1$. Then

$$\begin{aligned} \sum_{n=1}^t \beta^n \sum_{m=1}^n \Delta p_m &= \sum_{n=1}^{t-1} \beta^n \sum_{m=1}^n \Delta p_m + \beta^t \sum_{m=1}^t \Delta p_m \\ &= \beta^t \sum_{n=1}^{t-1} dp_n + \beta^t \sum_{m=1}^t \Delta p_m \\ &= \beta^t \sum_{n=1}^{t-1} dp_n + \beta^t \sum_{m=1}^t \beta dp_m - \beta^t \sum_{m=1}^t dp_{m-1} = \beta^{t+1} \sum_{m=1}^t \beta dp_m, \end{aligned}$$

where the second equality applies the induction hypothesis.

Combining this claim with equation (18) demonstrates that when $dp_0 = 0$, time zero consumption demand, dc_0 , depends only on $\sum_{n=1}^t dp_n$.

We turn now to the computations used to provide intuition for Proposition 1. Recall, it was assumed that $dq_{-1} = dp_{-1} = 0$. It follows that $R_0 Q_{-1}$ linearizes as qdp_0 . Thus we may write equation (15) as

$$\sum_{k=0}^T \beta^k dc_k = qdp_0 - \beta^T pdq_T - \beta^T qdp_T + Q \sum_{k=0}^T \beta^k dR_k. \quad (19)$$

Next, we derive equation (3) from the body, which we reproduce here for convenience:

$$dc_0 = \frac{1-\beta}{1-\beta^{T+1}} \left(qdp_0 - \beta^T p dq_T - \beta^T q dp_T + \frac{Q(\sigma-1)}{\sigma} \sum_{k=1}^T \beta^k dR_k + \frac{Q}{\sigma} \beta^{T+1} \sum_{k=1}^T dR_k \right). \quad (20)$$

Combining (14) and (19), we get

$$\sum_{k=0}^T \beta^k \left(dc_0 + \frac{(1-\beta)Q}{\sigma} \sum_{m=1}^k dR_m \right) = q_{-1} dp_0 - \beta^T dQ_T + Q \sum_{k=1}^T \beta^k dR_k,$$

or

$$\left(\frac{1-\beta^{T+1}}{1-\beta} \right) dc_0 = qdp_0 - \beta^T dQ_T + Q \sum_{k=1}^T \beta^k dR_k - \frac{(1-\beta)Q}{\sigma} \sum_{k=1}^T \beta^k \sum_{m=1}^k dR_m.$$

Now notice

$$\sum_{k=1}^T \beta^k \sum_{m=1}^k dR_m = \sum_{k=1}^T \left(\frac{\beta^k - \beta^{T+1}}{1-\beta} \right) dR_k,$$

from which equation (20) follows. Finally, using (20), we compute $\partial dc_0 / \partial dp_T$, and find

$$\frac{\partial dc_0}{\partial dp_T} = -\beta^T q + \frac{Q(\sigma-1)}{\sigma p} \beta^T + \frac{Q}{\sigma p} \beta^{T+1} = \beta^T \frac{Q}{\sigma p} (\beta - 1).$$

A.3 Random-walk forecasts of future wealth

In this section, we couple Lemma 1 with the assumption that all agents forecast future chicken holdings dq_{t+T}^e as if they were to follow a random walk, i.e. $dq_{t+T}^e = dq_{t-1}$. We analyze the behavior of demand, planned consumption, chicken holdings and equilibrium prices under this additional assumption.

A.3.1 Individual demand schedule

Without loss of generality, assume homogeneous planning horizons and omit index i . Denote the expected average price over the next T periods by $d\bar{p}_t^e(T)$:

$$d\bar{p}_t^e(T) \equiv \frac{1}{T} \sum_{k=1}^T dp_{t+k}^e.$$

Then demand of the agent may be written

$$dq_t = dq_{t-1} - \phi g(T) dp_t + \phi h(T) d\bar{p}_t^e(T)$$

$$dc_t = y dq_{t-1} + p \phi g(T) dp_t - p \phi h(T) d\bar{p}_t^e(T),$$

where

$$\phi = \frac{(1-\beta)q}{p\sigma}, \quad g(T) = \frac{1-\beta^T}{1-\beta^{T+1}}, \quad \text{and} \quad h(T) = \frac{(1-\beta)T\beta^T}{1-\beta^{T+1}}.$$

It follows that the agent's demand for chickens is decreasing in current price and increasing in expected average price.

We now consider the agent's time t plan for holding chickens over the planning period t to $t + T$. Assuming that, at time t , the agent believes that her expected average price over the time period $t + k$ to $t + T$ will be $d\bar{p}_t^e(T)$ for each $k \in \{1, \dots, T - 1\}$, we may compute the plans for chicken holdings as

$$dq_{t+k} = dq_{t+k-1} - \phi(g(T-k) - h(T-k))d\bar{p}_t^e(T).$$

Letting $\Delta_T(j) = g(T-j) - h(T-j)$, it follows that

$$dq_{t+k} = dq_{t-1} - \phi g(T)dp_t + \phi h(T)d\bar{p}_t^e(T) - \phi \left(\sum_{j=1}^k \Delta_T(j) \right) d\bar{p}_t^e(T) \quad (21)$$

$$\begin{aligned} dc_{t+k} &= ydq_{t-1} - y\phi g(T)dp_t + y\phi h(T)d\bar{p}_t^e(T) \\ &\quad - \phi y \left(\sum_{j=1}^{k-1} \Delta_T(j) \right) d\bar{p}_t^e(T) + p\phi \Delta_T(k)d\bar{p}_t^e(T). \end{aligned} \quad (22)$$

Written differently, we have

$$\Delta dq_t = -\phi(g(T)dp_t - h(T)d\bar{p}_t^e(T)) \quad (23)$$

$$\Delta dq_{t+k} = -\phi \Delta_T(k)d\bar{p}_t^e(T). \quad (24)$$

The formulae identifying the changes in consumption are less revealing.

Now observe that $\Delta_T(k) > 0$. Indeed, letting $n = T - k$, we have

$$\begin{aligned} \Delta_T(k) &= \frac{1 - \beta^n - (1 - \beta)n\beta^n}{1 - \beta^{n+1}} = (1 - \beta) \left(\frac{\frac{1 - \beta^n}{1 - \beta} - n\beta^n}{1 - \beta^{n+1}} \right) \\ &= (1 - \beta) \left(\frac{\sum_{i=0}^{n-1} \beta^i - n\beta^n}{1 - \beta^{n+1}} \right) = (1 - \beta) \left(\frac{\sum_{i=0}^{n-1} (\beta^i - \beta^n)}{1 - \beta^{n+1}} \right) > 0. \end{aligned}$$

We may now consider the following thought experiments. Here we assume all variables are at steady state (i.e. zero in differential form) unless otherwise stated. All references to periods $t + k$ concern "plans," not realizations, and it is assumed that $k \in \{1, \dots, T - 1\}$.

1. **A rise in price.** If $dp_t > 0$, then by equations (23) and (24) chicken holdings are reduced in time t by $-\phi g(T)dp_t$ and the reduction is maintained throughout the period. Consumption rises in period t by $p\phi g(T)dp_t$ and is reduced in subsequent periods by $y\phi g(T)dp_t$. Intuitively, the rise in price today, together with change in expected future prices, lowers the return to holding chickens between today and tomorrow, causing the agent to substitute toward consumption today. After one period, the new, lower steady-state levels of consumption and chicken holdings are reached and maintained through the planning period.
2. **A rise in expected price.** If $d\bar{p}_t^e(T) > 0$, then by Equations (23) and (24), current chicken holdings rise by $\phi h(T)d\bar{p}_t^e(T)$ and then decline over time. Consumption falls in time t , rises in time $t + 1$, and is otherwise more complicated. Notice that our assumption of random-walk expectations of future chicken holdings forces holdings back to the original steady state.

A.3.2 Proof of Proposition 2

First item 1. That $\xi(\alpha, T) > 0$ follows from construction. The argument is completed by observing

$$\begin{aligned}\xi(\alpha, T) &= \frac{(1-\beta) \sum_i \left(\frac{\alpha_i T_i \beta^{T_i}}{1-\beta^{T_i+1}} \right)}{\sum_j \left(\frac{\alpha_j (1-\beta^{T_j})}{1-\beta^{T_j+1}} \right)} = \frac{\sum_i \left(\frac{\alpha_i T_i \beta^{T_i}}{1-\beta^{T_i+1}} \right)}{\sum_j \left(\frac{\alpha_j \left(\frac{1-\beta^{T_j}}{1-\beta} \right)}{1-\beta^{T_j+1}} \right)} \\ &= \frac{\sum_i \left(\frac{\alpha_i T_i \beta^{T_i}}{1-\beta^{T_i+1}} \right)}{\sum_j \left(\frac{\alpha_j \left(\sum_{k=0}^{T_j-1} \beta^k \right)}{1-\beta^{T_j+1}} \right)} < \frac{\sum_i \left(\frac{\alpha_i T_i \beta^{T_i}}{1-\beta^{T_i+1}} \right)}{\sum_j \left(\frac{\alpha_j \left(\sum_{k=0}^{T_j-1} \beta^{T_j} \right)}{1-\beta^{T_j+1}} \right)} = 1.\end{aligned}$$

Next, item 2. For convenience, we allow T to take any positive real value. It suffices to show that

$$f(x) = \log \xi(x) - \log(1-\beta) = \log x + x \log \beta - \log(1-\beta^x)$$

is decreasing in x for $x > 0$. Notice that

$$f'(x) = \frac{1}{x} + \frac{\log \beta}{1-\beta^x},$$

hence for $x > 0$,

$$f'(x) \leq 0 \iff \frac{1}{\log \beta^{-1}} \leq \frac{x}{1-\beta^x} \equiv h(x).$$

Using L'Hopital's rule, we find that $h(0) = 1/\log \beta^{-1}$; thus it suffices to show that $h'(x) > 0$ for $x > 0$. Now compute

$$h'(x) = \frac{1 - \beta^x(1 + x \log \beta^{-1})}{(1 - \beta^x)^2}.$$

It follows that for $x > 0$,

$$h'(x) > 0 \iff h_1(x) \equiv \frac{1-\beta^x}{\beta^x} > x \log \beta^{-1} \equiv h_2(x).$$

Since $h_1(0) = h_2(0)$ and

$$h'_1(x) = \beta^{-x} \log \beta^{-1} > \log \beta^{-1} = h'_2(x),$$

the result follows.

Finally, item 3. Let $g(T_i) = (1 - \beta^{T_i+1})^{-1}(1 - \beta^{T_i})$. Assume $T_1 < T_2$, and, abusing notation somewhat, write

$$\xi(\alpha, T) = \frac{\alpha \xi(T_1) g(T_1) + (1-\alpha) \xi(T_2) g(T_2)}{\alpha g(T_1) + (1-\alpha) g(T_2)},$$

where we recall that

$$\xi(T) = (1-\beta) \frac{T \beta^T}{1-\beta^{T+1}}.$$

It suffices to show $\xi_\alpha > 0$. But notice this holds if and only if

$$\begin{aligned}
& (\alpha g(T_1) + (1 - \alpha)g(T_2))(\xi(T_1)g(T_1) - \xi(T_2)g(T_2)) \\
& > (\alpha \xi(T_1)g(T_1) + (1 - \alpha)\xi(T_2)g(T_2))(g(T_1) - g(T_2)) \\
& \iff \\
& \alpha(\xi(T_1) - \xi(T_2)) > (1 - \alpha)(\xi(T_2) - \xi(T_1)).
\end{aligned}$$

The last inequality holds from item 3 of the lemma and the fact that $T_2 > T_1$. ■

A.3.3 Proof of Corollaries 1 and 2

Proof of Corollary 1. Here we provide the argument for the constant gain case. The decreasing gain case is somewhat more involved but ultimately turns on the same computations.

Stack agents' expectations into the vector $d\bar{p}_t^e$, and let

$$\hat{\xi} = \begin{pmatrix} \xi(\alpha, \mathcal{T}, 1) & \cdots & \xi(\alpha, \mathcal{T}, N) \\ \xi(\alpha, \mathcal{T}, 1) & \cdots & \xi(\alpha, \mathcal{T}, N) \\ \vdots & \ddots & \vdots \\ \xi(\alpha, \mathcal{T}, 1) & \cdots & \xi(\alpha, \mathcal{T}, N) \end{pmatrix}.$$

Observe that $\hat{\xi}$ has an eigenvalue of zero with multiplicity $N - 1$, and the remaining eigenvalue given by $\text{tr } \hat{\xi} = \sum_i \xi(\alpha, \mathcal{T}, i)$, which, by Proposition 2, is contained in $(0, 1)$.

The recursive algorithm characterizing the beliefs dynamics of agent i may be written,

$$d\bar{p}_t^e(i, T_i) = d\bar{p}_{t-1}^e(i, T_i) + \gamma \left(\sum_i \xi(\alpha, \mathcal{T}, i) d\bar{p}_{t-1}^e(i, T_i) - d\bar{p}_{t-1}^e(i, T_i) \right),$$

or, in stacked form,

$$d\bar{p}_t^e = \left((1 - \gamma)I_N + \gamma \hat{\xi} \right) d\bar{p}_{t-1}^e. \quad (25)$$

Stability requires that the eigenvalues of $(1 - \gamma)I_N + \gamma \hat{\xi}$ be strictly less than one in modulus, and this is immediately implied by our above observation about the eigenvalues of $\hat{\xi}$. Via Eq. (6), convergence of expected price deviations to zero implies convergence of the realized price deviation to zero. ■

Proof of Corollary 2. The matrix $(1 - \gamma)I_N + \gamma \hat{\xi}$ has, as eigenvalues, $N - 1$ copies of $1 - \gamma$ and

$$\zeta = 1 - \gamma + \gamma \sum_i \xi(\alpha, \mathcal{T}, i).$$

Denote by S the corresponding matrix of eigenvectors and change coordinates: $z_t = S^{-1} d\bar{p}_t^e$. The dynamics (25) becomes the decoupled system $z_t = \Lambda z_{t-1}$. Denote by z_t^ζ the component of z_t corresponding to the eigenvalue ζ . With the aid of a computer algebra system, it is straightforward to show that

$$z_t^\zeta = \left(\sum_i \xi(\alpha, \mathcal{T}, i) \right)^{-1} \sum_i \xi(\alpha, \mathcal{T}, i) d\bar{p}_t^e(i, T_i) = \xi(\alpha, \mathcal{T})^{-1} dp_t.$$

It follows that

$$\frac{dp_t}{dp_{t-1}} = \frac{z_t^\zeta}{z_{t-1}^\zeta} = \zeta.$$

The argument is completed by noting that ζ is decreasing in T_i . ■

B Individual plots

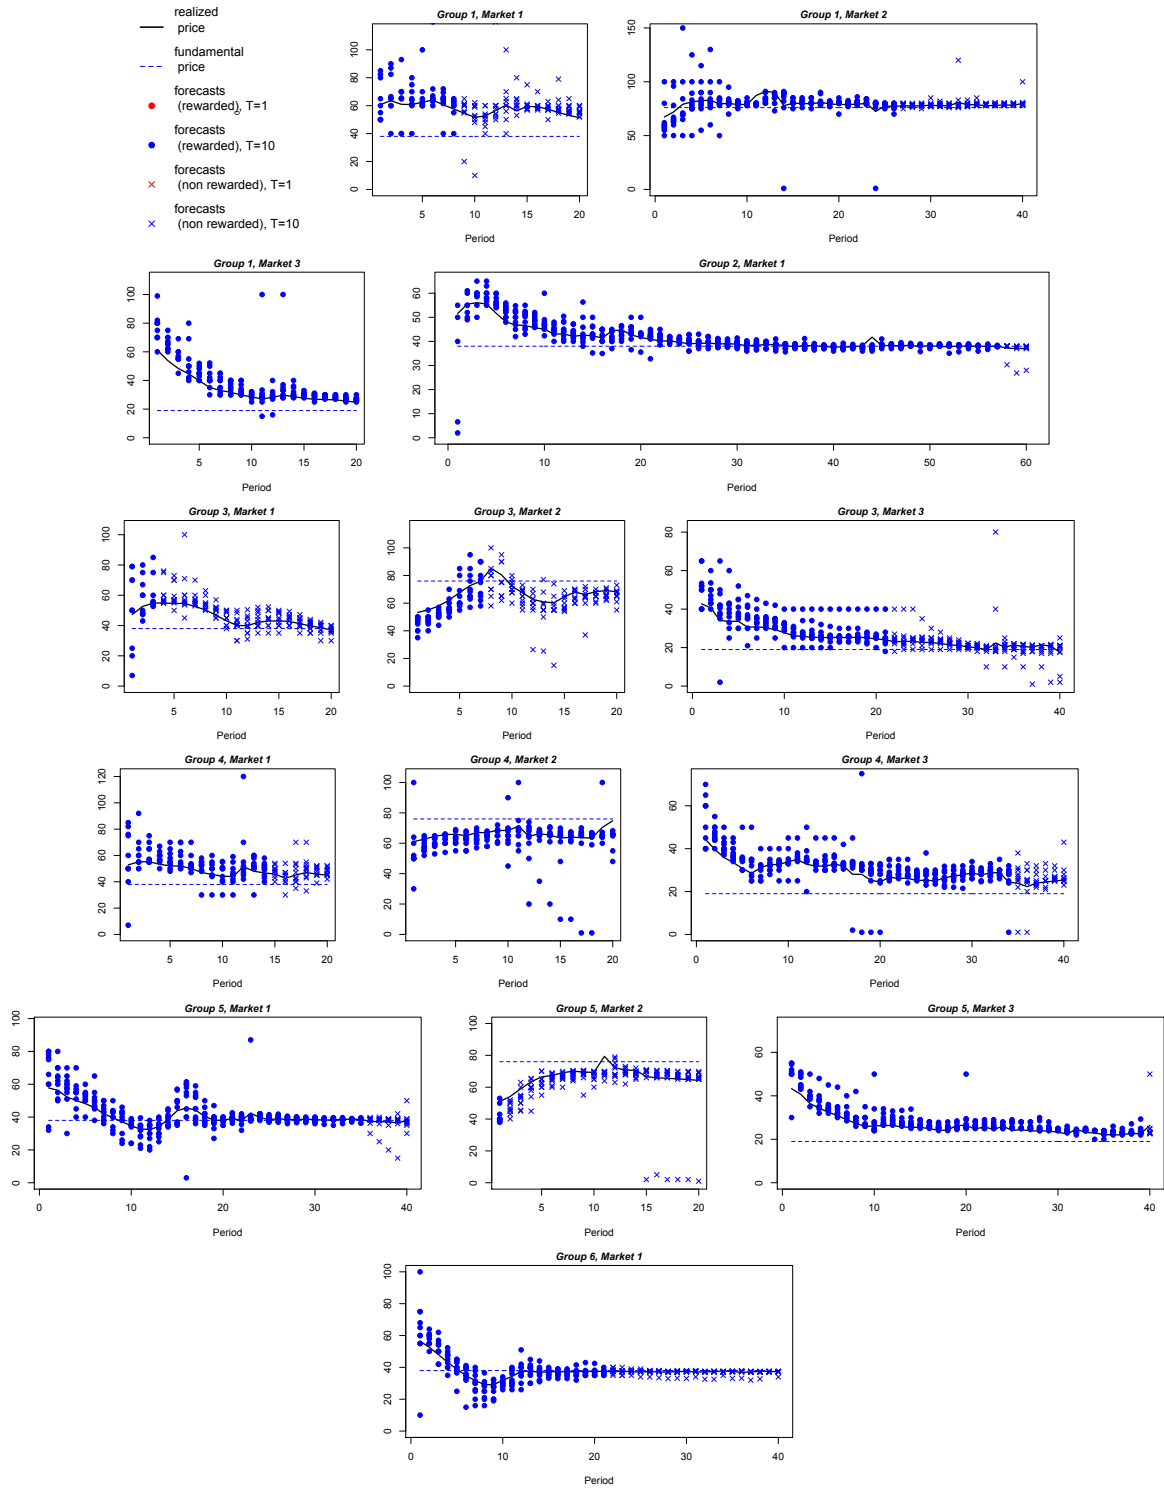


Figure 10: Prices and participants' forecasts in Treatment L: 100% $T = 10$ (6 groups)

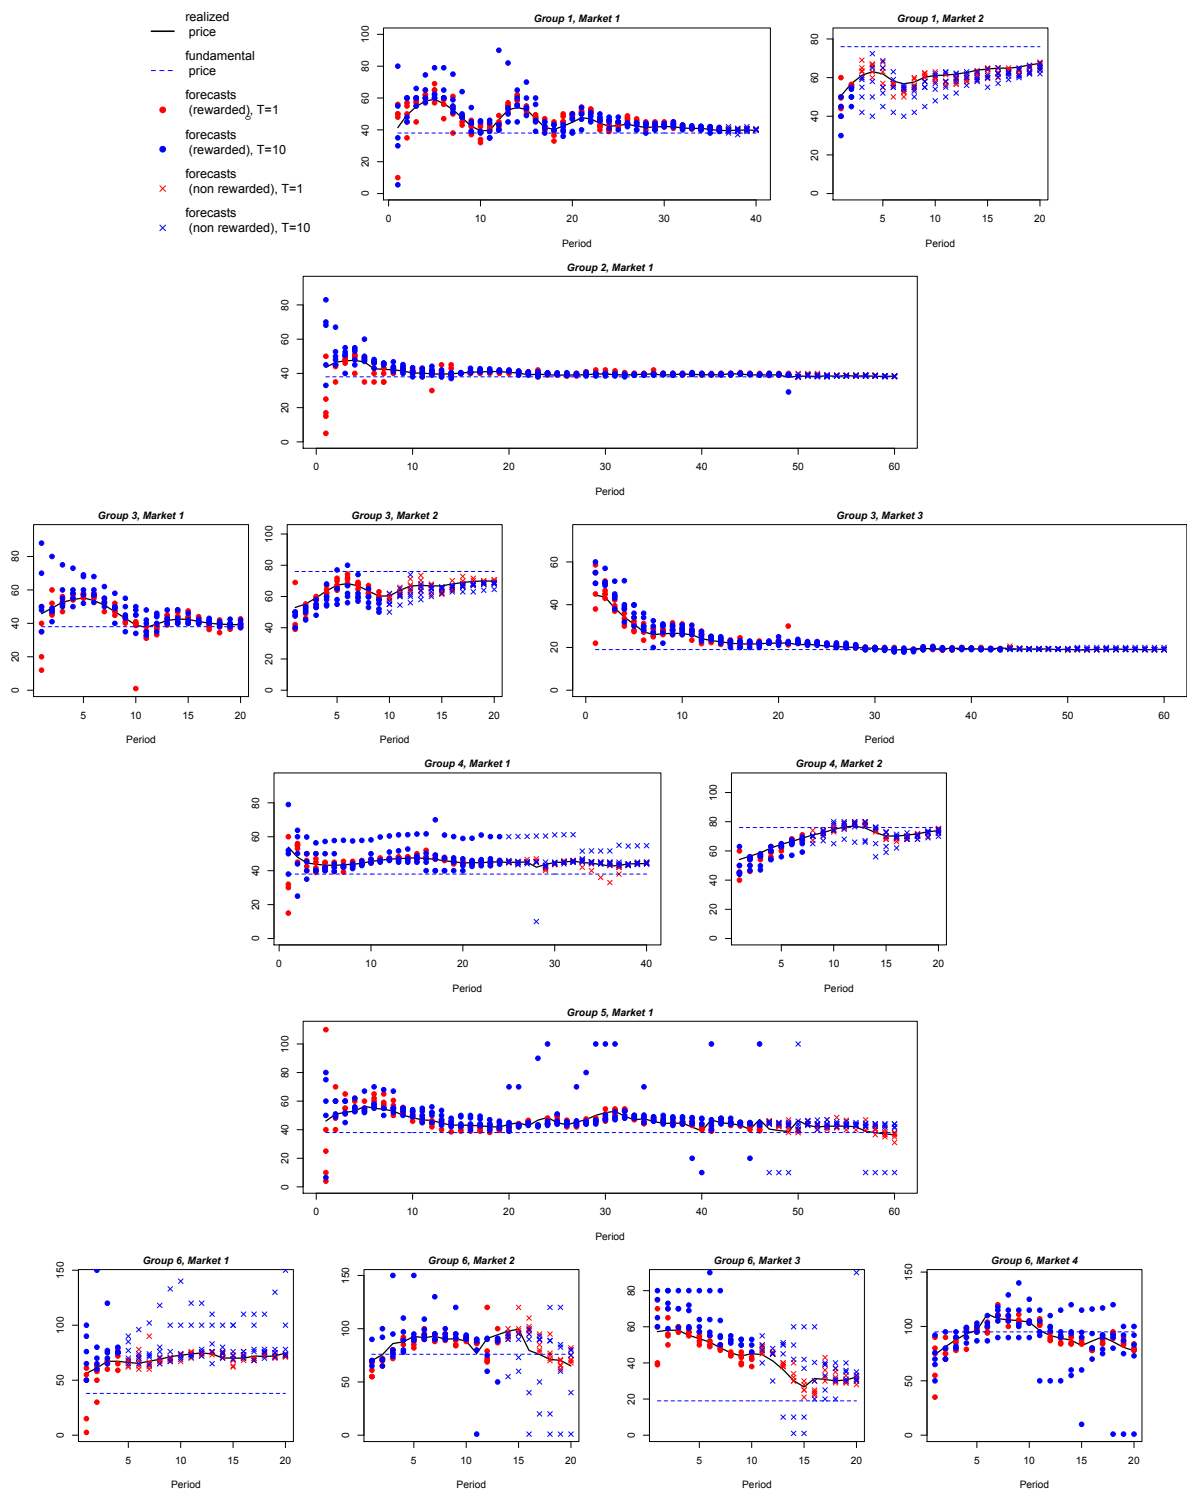


Figure 11: Prices and participants' forecasts in Treatment M50: 50% $T = 1$ /50% $T = 10$ (6 groups)

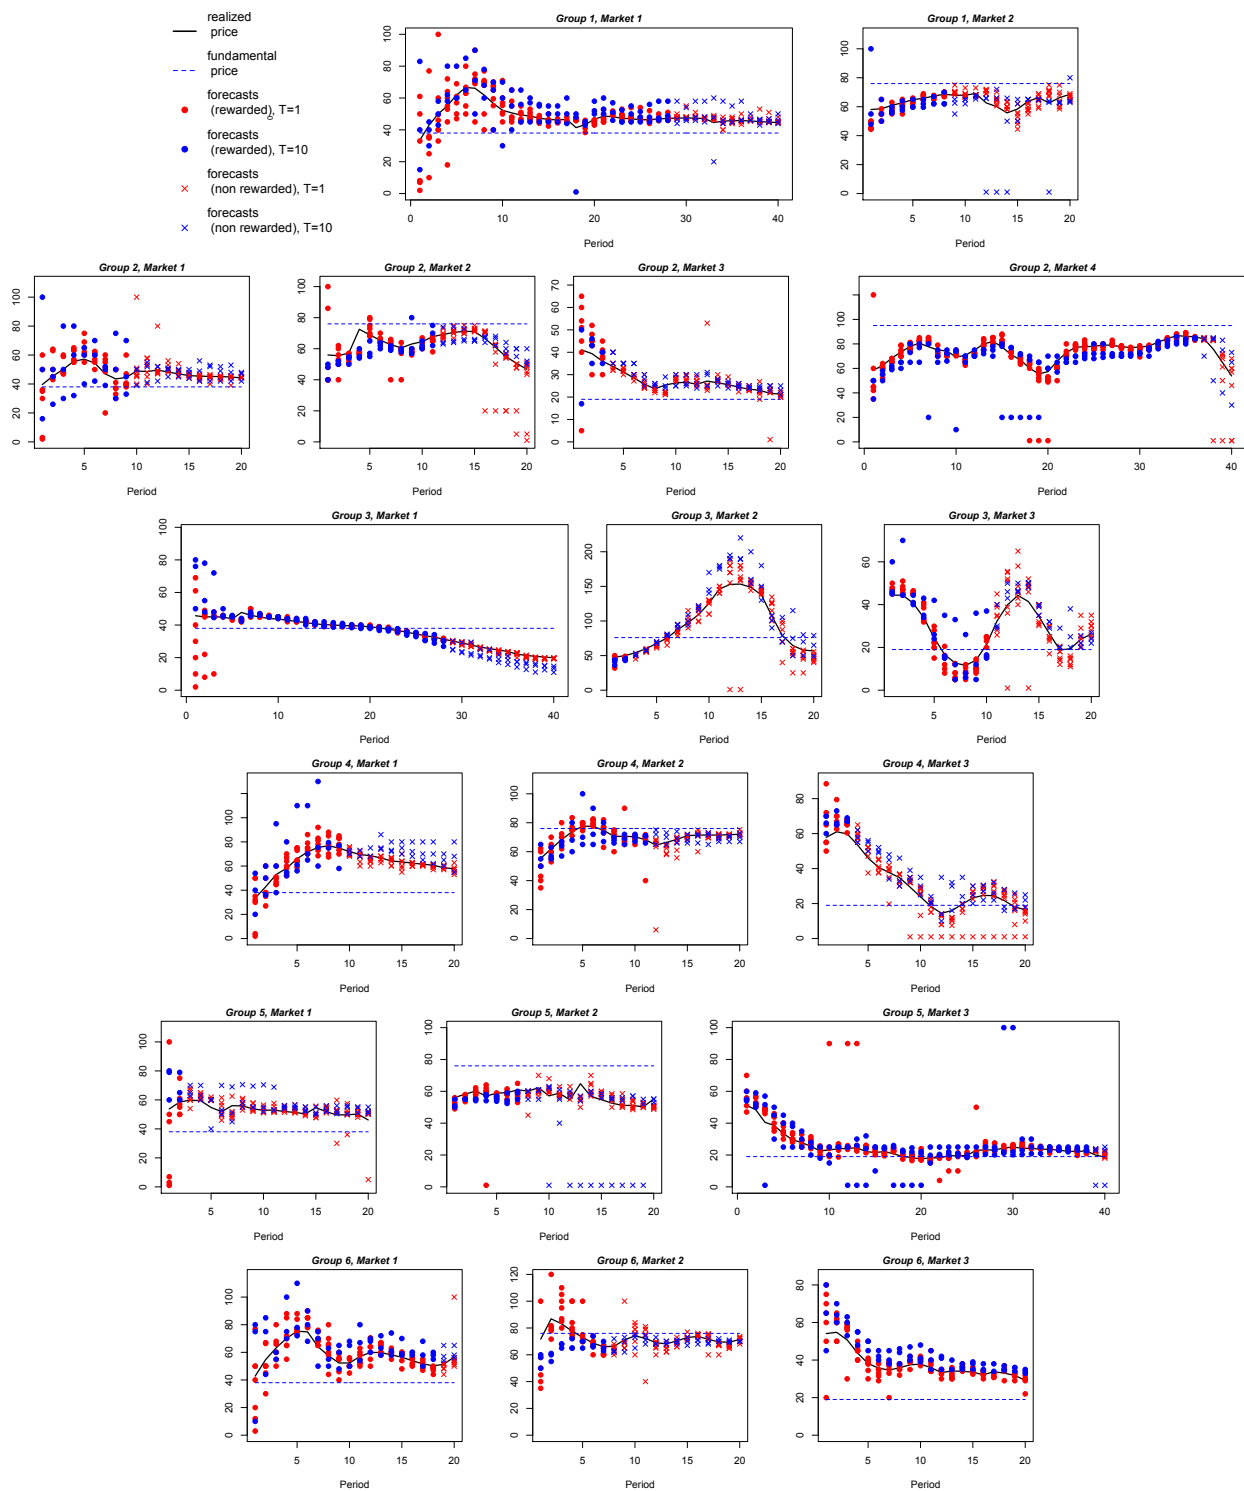


Figure 12: Prices and participants' forecasts in Treatment M70: 70% $T = 1$ /30% $T = 10$

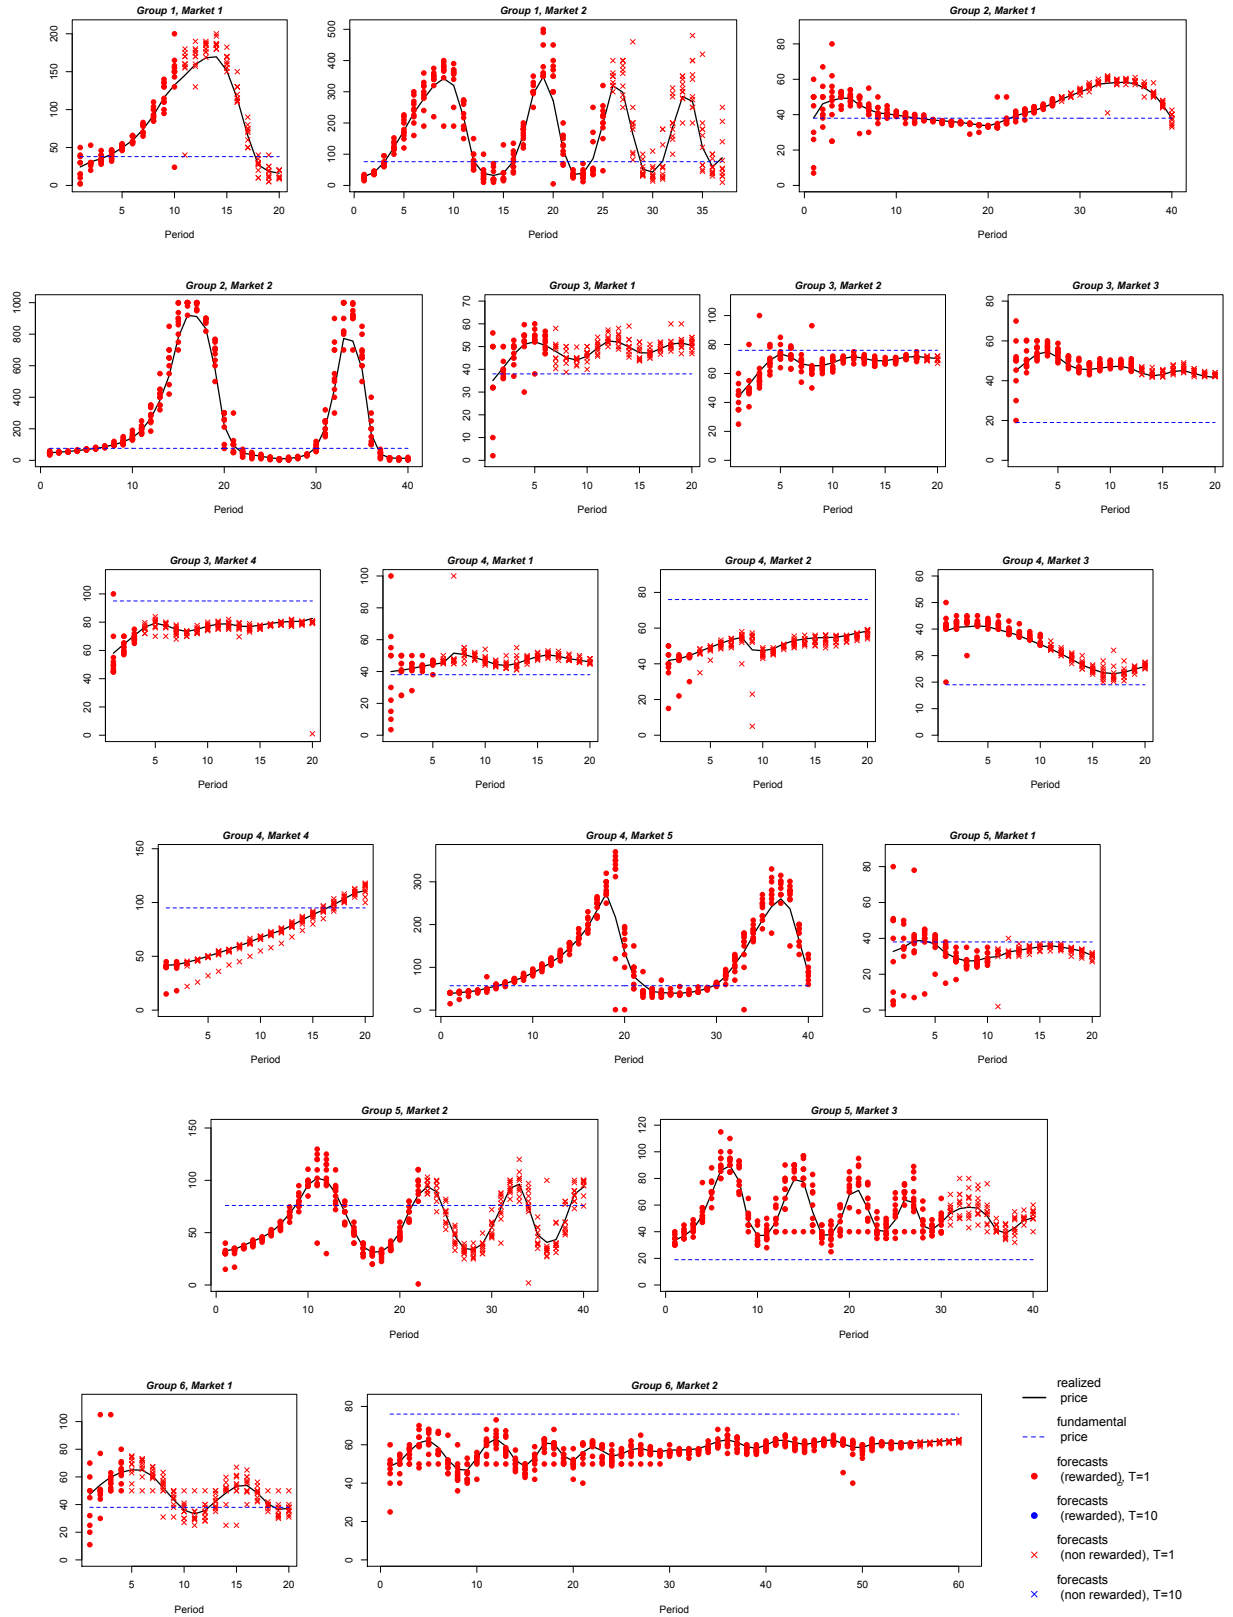


Figure 13: Prices and participants' forecasts in Treatment S: 100% $T = 1$ (6 groups)

C Test of the equilibration hypothesis: estimation outcomes of Equation (10)

For each treatment, we report below the estimated coefficients $\{\hat{b}_{1,m}, \hat{b}_{2,m}\}$ in Equation (10) for each market $m_{i,m}$ (subscript i corresponds to the group number and subscript m to the number of the market). Standard deviations are reported between brackets. We highlight in bold the markets that exhibit weak convergence, and denote with a star those exhibiting strong convergence.

$$\begin{aligned} \text{Tr. S} \quad & m_{1,1} : \left\{ \begin{matrix} -1.012, 1.715 \\ (0.464) \quad (0.449) \end{matrix} \right\}; m_{1,2} : \left\{ \begin{matrix} -0.565, 1.321 \\ (0.365) \quad (0.224) \end{matrix} \right\}; m_{2,1}^* : \left\{ \begin{matrix} 0.040, 0.183 \\ (0.104) \quad (0.103) \end{matrix} \right\}; m_{2,2} : \left\{ \begin{matrix} -1.821, 2.714 \\ (0.958) \quad (0.903) \end{matrix} \right\}; \\ & m_{3,1} : \left\{ \begin{matrix} -0.069, 0.335 \\ (0.015) \quad (0.011) \end{matrix} \right\}; \mathbf{m}_{3,2} : \left\{ \begin{matrix} -0.433, -0.048 \\ (0.009) \quad (0.005) \end{matrix} \right\}; m_{3,3} : \left\{ \begin{matrix} 1.646, 1.393 \\ (0.122) \quad (0.084) \end{matrix} \right\}; \mathbf{m}_{3,4} : \left\{ \begin{matrix} -0.408, -0.155 \\ (0.007) \quad (0.006) \end{matrix} \right\}; \\ & m_{4,1} : \left\{ \begin{matrix} -0.001, 0.265 \\ (0.018) \quad (0.011) \end{matrix} \right\}; \mathbf{m}_{4,2} : \left\{ \begin{matrix} -0.500, -0.290 \\ (0.021) \quad (0.017) \end{matrix} \right\}; \mathbf{m}_{4,3} : \left\{ \begin{matrix} 1.469, 0.554 \\ (0.217) \quad (0.180) \end{matrix} \right\}; \mathbf{m}_{4,4}^* : \left\{ \begin{matrix} -0.805, -0.121 \\ (0.130) \quad (0.111) \end{matrix} \right\}; \\ & m_{4,5} : \left\{ \begin{matrix} -1.224, 1.220 \\ (0.506) \quad (0.354) \end{matrix} \right\}; m_{5,1} : \left\{ \begin{matrix} -0.076, -0.143 \\ (0.028) \quad (0.019) \end{matrix} \right\}; \mathbf{m}_{5,2} : \left\{ \begin{matrix} -0.693, -0.126 \\ (0.064) \quad (0.047) \end{matrix} \right\}; m_{5,3} : \left\{ \begin{matrix} 1.068, 1.910 \\ (0.335) \quad (0.162) \end{matrix} \right\}; \\ & m_{6,1} : \left\{ \begin{matrix} 0.527, 0.232 \\ (0.134) \quad (0.099) \end{matrix} \right\}; \mathbf{m}_{6,2} : \left\{ \begin{matrix} -0.384, -0.223 \\ (0.026) \quad (0.014) \end{matrix} \right\}. \end{aligned}$$

Observations: 518; Adj. $R^2 = 0.381$; F Statistic: 9.849 (df = 36; 482)

$$\begin{aligned} \text{Tr. M70} \quad & m_{1,1} : \left\{ \begin{matrix} 0.137, 0.295 \\ (0.082) \quad (0.053) \end{matrix} \right\}; \mathbf{m}_{1,2} : \left\{ \begin{matrix} -0.245, -0.142 \\ (0.003) \quad (0.002) \end{matrix} \right\}; m_{2,1} : \left\{ \begin{matrix} 0.154, 0.272 \\ (0.027) \quad (0.021) \end{matrix} \right\}; \mathbf{m}_{2,2} : \left\{ \begin{matrix} -0.257, -0.156 \\ (0.016) \quad (0.013) \end{matrix} \right\}; \\ & \mathbf{m}_{2,3} : \left\{ \begin{matrix} 1.410, 0.256 \\ (0.034) \quad (0.028) \end{matrix} \right\}; \mathbf{m}_{2,4} : \left\{ \begin{matrix} -0.391, -0.202 \\ (0.007) \quad (0.004) \end{matrix} \right\}; m_{3,1}^* : \left\{ \begin{matrix} 0.493, -0.127 \\ (0.141) \quad (0.093) \end{matrix} \right\}; \mathbf{m}_{3,2} : \left\{ \begin{matrix} -0.655, 0.431 \\ (0.023) \quad (0.019) \end{matrix} \right\}; \\ & \mathbf{m}_{3,3} : \left\{ \begin{matrix} 1.497, 0.286 \\ (0.030) \quad (0.024) \end{matrix} \right\}; m_{4,1} : \left\{ \begin{matrix} -0.147, 0.831 \\ (0.040) \quad (0.033) \end{matrix} \right\}; \mathbf{m}_{4,2} : \left\{ \begin{matrix} -0.248, -0.043 \\ (0.008) \quad (0.006) \end{matrix} \right\}; \mathbf{m}_{4,3}^* : \left\{ \begin{matrix} 2.993, 0.184 \\ (0.142) \quad (0.115) \end{matrix} \right\}; \\ & \mathbf{m}_{5,1} : \left\{ \begin{matrix} 0.542, 0.371 \\ (0.023) \quad (0.019) \end{matrix} \right\}; m_{5,2} : \left\{ \begin{matrix} -0.227, -0.255 \\ (0.012) \quad (0.010) \end{matrix} \right\}; \mathbf{m}_{5,3} : \left\{ \begin{matrix} 2.140, 0.088 \\ (0.030) \quad (0.018) \end{matrix} \right\}; m_{6,1} : \left\{ \begin{matrix} 0.348, 0.569 \\ (0.054) \quad (0.043) \end{matrix} \right\}; \\ & m_{6,2} : \left\{ \begin{matrix} 0.047, -0.069 \\ (0.006) \quad (0.005) \end{matrix} \right\}; \mathbf{m}_{6,3} : \left\{ \begin{matrix} 2.210, 0.716 \\ (0.041) \quad (0.033) \end{matrix} \right\}. \end{aligned}$$

Observations: 441; Adj. $R^2 = 0.783$; F Statistic: 45.21 (df = 36; 405)

$$\begin{aligned} \text{Tr. M50} \quad & m_{1,1} : \left\{ \begin{matrix} 0.373, 0.163 \\ (0.108) \quad (0.055) \end{matrix} \right\}; \mathbf{m}_{1,2} : \left\{ \begin{matrix} -0.341, -0.157 \\ (0.012) \quad (0.012) \end{matrix} \right\}; \mathbf{m}_{2,1} : \left\{ \begin{matrix} 0.295, 0.035 \\ (0.045) \quad (0.012) \end{matrix} \right\}; \mathbf{m}_{3,1}^* : \\ & \left\{ \begin{matrix} 0.391, 0.128 \\ (0.080) \quad (0.075) \end{matrix} \right\}; \mathbf{m}_{3,2} : \left\{ \begin{matrix} -0.337, -0.106 \\ (0.012) \quad (0.010) \end{matrix} \right\}; \mathbf{m}_{3,3}^* : \left\{ \begin{matrix} 1.846, 0.029 \\ (0.140) \quad (0.043) \end{matrix} \right\}; \mathbf{m}_{4,1} : \left\{ \begin{matrix} 0.368, 0.166 \\ (0.019) \quad (0.010) \end{matrix} \right\}; \mathbf{m}_{4,2} : \end{aligned}$$

$$\left\{ \begin{matrix} -0.361, -0.033 \\ (0.022) \quad (0.013) \end{matrix} \right\}; \mathbf{m}_{5,1} : \left\{ \begin{matrix} 0.447, 0.170 \\ (0.092) \quad (0.035) \end{matrix} \right\}; m_{6,1} : \left\{ \begin{matrix} 0.426, 0.911 \\ (0.019) \quad (0.014) \end{matrix} \right\}; m_{6,2} : \left\{ \begin{matrix} -0.026, 0.139 \\ (0.052) \quad (0.046) \end{matrix} \right\}; \\ \mathbf{m}_{6,3} : \left\{ \begin{matrix} 2.633, 0.932 \\ (0.265) \quad (0.200) \end{matrix} \right\}; \mathbf{m}_{6,4}^* : \left\{ \begin{matrix} -0.176, -0.003 \\ (0.049) \quad (0.044) \end{matrix} \right\}.$$

Observations: 421; Adj. $R^2 = 0.887$; F Statistic: 128.3 (df = 26; 395)

$$\mathbf{Tr. L} \quad \mathbf{m}_{1,1} : \left\{ \begin{matrix} 0.697, 0.487 \\ (0.035) \quad (0.024) \end{matrix} \right\}; m_{1,2} : \left\{ \begin{matrix} -0.079, 0.057 \\ (0.026) \quad (0.017) \end{matrix} \right\}; \mathbf{m}_{1,3} : \left\{ \begin{matrix} 2.621, 0.381 \\ (0.128) \quad (0.077) \end{matrix} \right\}; \mathbf{m}_{2,1}^* : \left\{ \begin{matrix} 0.663, 0.034 \\ (0.097) \quad (0.032) \end{matrix} \right\}; \\ \mathbf{m}_{3,1}^* : \left\{ \begin{matrix} 0.457, 0.155 \\ (0.090) \quad (0.066) \end{matrix} \right\}; \mathbf{m}_{3,2} : \left\{ \begin{matrix} -0.325, -0.074 \\ (0.020) \quad (0.020) \end{matrix} \right\}; \mathbf{m}_{3,3}^* : \left\{ \begin{matrix} 1.632, 0.179 \\ (0.138) \quad (0.097) \end{matrix} \right\}; \mathbf{m}_{4,1} : \left\{ \begin{matrix} 0.512, 0.223 \\ (0.040) \quad (0.026) \end{matrix} \right\}; \\ \mathbf{m}_{4,2} : \left\{ \begin{matrix} -0.205, -0.113 \\ (0.005) \quad (0.005) \end{matrix} \right\}; \mathbf{m}_{4,3} : \left\{ \begin{matrix} 1.531, 0.412 \\ (0.091) \quad (0.062) \end{matrix} \right\}; \mathbf{m}_{5,1}^* : \left\{ \begin{matrix} 0.676, -0.007 \\ (0.040) \quad (0.011) \end{matrix} \right\}; \mathbf{m}_{5,2} : \left\{ \begin{matrix} -0.362, -0.077 \\ (0.015) \quad (0.014) \end{matrix} \right\}; \\ \mathbf{m}_{5,3} : \left\{ \begin{matrix} 1.579, 0.249 \\ (0.086) \quad (0.036) \end{matrix} \right\}; \mathbf{m}_{5,3} : \left\{ \begin{matrix} 0.522, -0.068 \\ (0.041) \quad (0.025) \end{matrix} \right\}.$$

Observations: 441; Adj. $R^2 = 0.910$; F Statistic: 160.8 (df = 28; 413)

D Supplementary material

Instructions (*N.B.: for short-horizon forecasters*)

Welcome to our experiment! The experiment is **anonymous**, the data from your choices will only be linked to your station ID, never to your name. You will be paid privately at the end, after all participants have finished the experiment. During the experiment, you are not allowed to communicate with other participants. If you have a question at any time, please raise your hand and we will come to your desk.

Please read these instructions carefully and **answer the quiz (five questions)**. Once we have made sure that all participants have answered correctly, we will start the experiment. **At the end of the experiment** and before the payment, you will be asked to **fill out a short questionnaire**.

Thank you for your participation!

Your role: price forecasting on a chicken market

You are a farmer in a chicken market and have to *submit price forecasts*. There are 10 farmers in the market, every farmer is a participant like you. The group of farmers will not change during the experiment. Every chicken produces the same number of eggs at the beginning of each period. Eggs are **either** traded for chickens in the market **or** consumed by the farmers. Eggs cannot be carried over between periods. You do not need to make trading decisions, a computerized trader will do it for you based on your forecasts.

Sequence of markets: You may play in several markets in a row

In each period after the first one, an outbreak of avian flu may occur with a constant and independent probability of 5%. If this happens, all chickens die and become worthless, and a new market starts. We will run **as many markets as possible** during the time for which you have been recruited. You will play in every market for at least **20 periods** because **you will only find out after 20 periods whether and in which period the chickens have died** from avian flu.

If the chickens have died within the first 20 periods, the market stops in period 20, you receive new chickens and **enter a new market**. If the chickens have not died within the first 20 periods, **you play**

for 20 more periods, till period 40, and then observe whether the chickens have died between period 21 and period 40. If this is the case, a new market starts. If not, you continue in the market for another 20 periods, etc., till the chickens die and the market ends. **All periods after the chickens died will not be counted towards your earnings.**

At the beginning of each market, the **number of eggs produced per chicken** and the **number of chickens that you have received** will be displayed on your computer interface, which is mainly self-explanatory. The number of eggs per chicken **remains fixed for the whole market**. You never observe the number of chickens of the other 9 farmers.

Your task: Forecasting the price in the next period

In each period, your task is to **forecast the price** of a chicken *in the next period*: in period 1, you have to forecast the price in period 2; in period 2, you have to forecast the price in period 3, etc. **Based on your forecasts, a computerized trader buys or sells chickens on your behalf.** Not all participants may have a computerized trader using forecasts for the next period, they may use a different time horizon.

The price of a chicken in terms of eggs is always adjusted so that the demand for chickens equals the supply (up to small random errors). The price depends on all participants' forecasts: if all participants forecast an **increase (resp. decrease) in the price**, the *current price* will tend to **increase (resp. decrease)**. **Once every participant has submitted a forecast**, the computers trade, **the current price of a chicken is determined** and you observe how many chickens you have bought or sold and your egg consumption. Your egg consumption is your amount of eggs at the beginning of the period plus the eggs you received from selling chickens (or minus the eggs you used to buy chickens). Your trader ensures that you always consume at least one egg and always have at least one chicken in any period.

Whether your trader buys or sells chickens depends both on your forecast and the forecasts of the other participants. The **higher your forecast compared to the forecasts of the other participants**, the **more chickens your trader buys**, and the **fewer eggs you consume now (but the more later on)**. The **lower your forecasts compared to the ones of the other participants**, the **more chickens your trader sells**, and the **more eggs you consume now (but the fewer later on)**.

Your payoff: forecasting accuracy and egg consumption

You may earn points in **two ways**. **First**, you may earn points based on your **price forecast accuracy**. The closer your forecast to the **realized price**, the higher your payoff. **There is a forecasting payoff table on your desk** that indicates how many points you make that way. If your prediction is perfect (zero error), you earn a maximum of 1100 points. If your forecast error is larger than 7, you earn zero point.

You receive your **first forecasting payoff** once the first price that you had to forecast becomes observable, that is **at the end of period 2**. Your corresponding forecast error is the difference between your forecast of the price in period 2 that you made in period 1 and the realized price in period 2.

If the chickens die, some of your last forecasts will not be rewarded because **only the price of living chickens counts towards the computation of the average price** that you had to forecast. For this reason, we pay **2 times your last rewarded forecast**, that is the one for which we can compute the price that you had to forecast. Your total forecasting payoff in each market is the sum of your realized forecasting payoff in that market.

Second, you may earn points with your **egg consumption**. **There is a consumption payoff table on your desk** that indicates how many points you earn that way. The more eggs you consume, the higher your consumption payoff. Notice that **consuming few eggs in one period** (e.g. 20) **and a lot in the next** (e.g. 980) **gives you less points** ($359 + 827 = 1186$, see your payoff table!) **than consuming an equal amount of eggs (500) in the two periods** ($746 \times 2 = 1492$). Your total consumption payoff in each market is the sum of your consumption payoff in every period for which the chickens are alive.

At the end of **each market**, you will be rewarded **either** with your total consumption **or** your total forecasting payoff **with equal probability**. This **does not depend** on how you have been rewarded in the previous markets. **If the chickens do not live for at least 2 periods, you will not have any forecasting score, so you will be rewarded with your consumption payoff**. All participants are paid in the same way.


Your total amount of points over all markets will be converted into euro and paid to you at the end of the experiment. One euro corresponds to 2000 points.

Example

The box below provides an example of a sequence of events in a market where the chickens die in period 15, so you play until period 20.

You enter **period 1**.


You submit a forecast of the price in period 2:



You then observe P_1 (the price in **period 1**), the number of chickens you traded, your corresponding egg consumption and consumption points in period 1.

You enter **period 2**.

You submit a forecast of the price in period 3:




You then observe P_2 . You then see your forecast error and forecasting payoff for your forecast made in period 1.

You also observe the number of chickens you traded, your corresponding egg consumption and consumption points in period 2.

⋮

You enter **period 11**.


You submit a forecast of the price in period 12:



You observe P_{11} , your forecast made in period 10, your forecast error and corresponding forecasting payoff. You also observe the number of chickens you traded, your corresponding egg consumption and consumption points in period 11.

⋮

You enter **period 20**, you submit a prediction for period 21 **and then** observe whether the chickens have died during the last 20 periods: the chickens died in **period 15**, this market ends.



Your last forecasting payoff is for your forecast in period 14 of the price P_{15} in period 15. **This payoff is multiplied by 2.**

Your other forecasting payoffs are from period 1 (for period 2) till period 13 (for period 14). **Your forecasts made from period 15 till 20 are not rewarded.** Your total consumption payoff is the sum of your consumption points from period 1 to 15. **Your egg consumption from period 16 till 20 is not rewarded.**

In this market, you earn **either** your consumption payoff **or** your forecasting payoff with equal probability. You then enter a new market.

Quiz

1. Assume you are at the end of period 13. Here is the sequence of realized prices:

$p_1 = 32, p_2 = 61, p_3 = 77, p_4 = 78, p_5 = 120, p_6 = 42, p_7 = 96, p_8 = 100, p_9 = 90, p_{10} = 70, p_{11} = 71, p_{12} = 46, p_{13} = 4.$

(a) In period 12, what did you have to predict?

- i. The price in period 13.
- ii. The price in period 12.
- iii. The difference between the price in period 12 and in period 13.
- iv. None of the above.

(b) Assume that, in period 12, you predicted a price of 2.5.

- i. What is your forecast error?
- ii. How many forecasting points do you earn?
(USE YOUR PAYOFF TABLE ON YOUR DESK!)

2. What is the probability for all the chickens to die if you are entering...

- (a) ... period 6?
- (b) ... period 45?

3. If the chickens die in period 23, when will you find out?

4. If you forecast a high price for the next period, ...

N.B.: Multiple answers may be possible.

- (a) Your forecast will not impact your demand for chickens.
- (b) Your trader is likely to buy chickens.
- (c) Your trader is likely to sell chickens.
- (d) You are likely to consume fewer eggs now but more later on.
- (e) You are likely to consume more eggs now but fewer later on.
- (f) This also depends on the other participants' forecasts.

5. If all participants tend to forecast an increase in the price in the future compared to past levels, what is the implication on the realized current market price?

N.B.: Multiple answers may be possible.

- (a) The realized market price is likely to increase.
- (b) The realized market price is likely to decrease.
- (c) The realized market price is likely to remain stable.
- (d) The realized market price will not be impacted by participants' forecasts.

Instructions (*N.B.: for long-horizon forecasters*)

Welcome to our experiment! The experiment is **anonymous**, the data from your choices will only be linked to your station ID, never to your name. You will be paid privately at the end, after all participants have finished the experiment. During the experiment, you are not allowed to communicate with other participants. If you have a question at any time, please raise your hand and we will come to your desk.

Please read these instructions carefully and **answer the quiz (five questions)**. Once we have made sure that all participants have answered correctly, we will start the experiment. **At the end of the experiment** and before the payment, you will be asked to **fill out a short questionnaire**.

Thank you for your participation!

Your role: price forecasting on a chicken market

You are a farmer in a chicken market and have to ***submit price forecasts***. There are 10 farmers in the market, every farmer is a participant like you. The group of farmers will not change during the experiment. Every chicken produces the same number of eggs at the beginning of each period. Eggs are **either** traded for chickens in the market **or** consumed by the farmers. Eggs cannot be carried over between periods. You do not need to make trading decisions, a computerized trader will do it for you based on your forecasts.

Sequence of markets: You may play in several markets in a row

In each period after the first one, an outbreak of avian flu may occur with a constant and independent probability of 5%. If this happens, all chickens die and become worthless, and a new market starts. We will run **as many markets as possible** during the time for which you have been recruited. You will play in every market for at least **20 periods** because **you will only find out after 20 periods whether and in which period the chickens have died** from avian flu.

If the chickens have died within the first 20 periods, the market stops in period 20, you receive new chickens and **enter a new market**. If the chickens have not died within the first 20 periods, **you play for 20 more periods**, till period 40, and then observe whether the chickens have died between period 21 and period 40. If this is the case, a new market starts. If not, you continue in the market for another 20 periods, etc., till the chickens die and the market ends. **All periods after the chickens died will not be counted towards your earnings.**

At the beginning of each market, the **number of eggs produced per chicken** and the **number of chickens that you have received** will be displayed on your computer interface, which is mainly self-explanatory. The number of eggs per chicken **remains fixed for the whole market**. You never observe the number of chickens of the other 9 farmers.

Your task: Forecasting the average price over the next 10 periods

In each period, your task is to **forecast the average price** of a chicken *over the next 10 periods*: in period 1, you have to forecast the average price over period 2 to period 11 (i.e. the **average price over the next 10 periods**); in period 2, you have to forecast the average price over period 3 to period 12, etc. **Based on your forecasts, a computerized trader buys or sells chickens on your behalf**. Not all participants may have a computerized trader using forecasts for the next 10 periods, they may use a different time horizon.

The price of a chicken in terms of eggs is always adjusted so that the demand for chickens equals the supply (up to small random errors). The price depends on all participants' forecasts: if all participants forecast an **increase (resp. decrease) in the average price over the next 10 periods**, the **current price** will tend to **increase (resp. decrease)**. **Once every participant has submitted a forecast**, the computers trade, **the current price of a chicken is determined** and you observe how many chickens you have bought or sold and your egg consumption. Your egg consumption is your amount of eggs at the beginning of the period plus the eggs you received from selling chickens (or minus the eggs you used to buy chickens). Your trader ensures that you always consume at least one egg and always have at least one chicken in any period.

Whether your trader buys or sells chickens depends both on your forecast and the forecasts of the other participants. The **higher your forecast compared to the forecasts of the other participants, the more chickens your trader buys, and the fewer eggs you consume now (but the more later on)**. The **lower your forecasts compared to the ones of the other participants, the more chickens your trader sells, and the more eggs you consume now (but the fewer later on)**.

Your payoff: forecasting accuracy and egg consumption

You may earn points in **two ways**. **First**, you may earn points based on your **price forecast accuracy**. The closer your forecast to the realized **average** price, the higher your payoff. **There is a forecasting payoff**

table on your desk that indicates how many points you make that way. If your prediction is perfect (zero error), you earn a maximum of 1100 points. If your forecast error is larger than 7, you earn zero point.

You receive your **first forecasting payoff** once the first average price that you had to forecast becomes observable, that is **in period 11**. Your corresponding forecast error is the difference between your forecast of the average price over the periods 2 – 11 that you made in period 1 and the realized average price over the periods 2 – 11.

If the chickens die, some of your last forecasts will not be rewarded because **only the price of living chickens counts towards the computation of the average price** that you had to forecast. For this reason, we pay **11 times your last rewarded forecast**, that is the one for which we can compute the average price that you had to forecast. Your total forecasting payoff in each market is the sum of your realized forecasting payoff in that market.

Second, you may earn points with your **egg consumption**. **There is a consumption payoff table on your desk** that indicates how many points you earn that way. The more eggs you consume, the higher your consumption payoff. Notice that **consuming few eggs in one period** (e.g. 20) **and a lot in the next** (e.g. 980) **gives you less points** ($359 + 827 = 1186$, see your payoff table!) **than consuming an equal amount of eggs (500) in the two periods** ($746 \times 2 = 1492$). Your total consumption payoff in each market is the sum of your consumption payoff in every period for which the chickens are alive.

At the end of **each market**, you will be rewarded **either** with your total consumption **or** your total forecasting payoff **with equal probability**. This **does not depend** on how you have been rewarded in the previous markets. **If the chickens do not live for at least 11 periods, you will not have any forecasting score, so you will be rewarded with your consumption payoff**. All participants are paid in the same way.


Your total amount of points over all markets will be converted into euro and paid to you at the end of the experiment. One euro corresponds to 2000 points.

Example

The box below provides an example of a sequence of events in a market where the chickens die in period 15, so you play until period 20.

You enter **period 1**.


Every player submits a forecast of the average price over periods 2 to 11:



You then observe P_1 (the price **in period 1**), the number of chickens you traded, your corresponding egg consumption and consumption points in period 1.

You enter **period 2**.

Every player submits a forecast of the average price from period 3 to 12:




You then observe P_2 , the number of chickens you traded, your corresponding egg consumption and consumption points in period 2.

⋮

You enter **period 11**.

Every player submits a forecast of the average price from period 12 to 21:




You observe P_{11} and the average price $\tilde{P}_{2,11}$ over period 2 to 11. You then see your forecast error and forecasting payoff for your forecast made in period 1.

You also observe the number of chickens you traded, your corresponding egg consumption and consumption points in period 11.

⋮

You enter **period 20**, submit a forecast for $\tilde{P}_{21,30}$ **and then** observe whether the chickens have died over the last 20 periods: the chickens died **in period 15**, this market ends.



Your last forecasting payoff is for your forecast in period 5 of the average price $\tilde{P}_{6,15}$ between periods 6 and 15 because the chickens have died in period 15, so we do not have prices afterwards. **This payoff is multiplied by 11.**

Your other forecasting payoffs are in period 1 (for periods 2 to 11), period 2 (for periods 3 to 12), period 3 (for periods 4 to 13) and period 4 (for periods 5 to 14). **Your forecasts made from period 6 till 20 are not rewarded.**

Your total consumption payoff is the sum of your consumption points from period 1 to 15. **Your egg consumption from period 16 till 20 is not rewarded.**

In this market, you earn **either** your consumption payoff **or** your forecasting payoff with equal probability. You then enter a new market.

Quiz

1. Assume you are entering period 14. Here is the sequence of realized prices:

$p_1 = 32, p_2 = 61, p_3 = 77, p_4 = 78, p_5 = 120, p_6 = 42, p_7 = 96, p_8 = 100, p_9 = 90, p_{10} = 70, p_{11} = 71, p_{12} = 46, p_{13} = 4.$

- (a) In period 2, what did you have to predict?
 - i. The price in period 3.
 - ii. The average price over periods 3 to 12.
 - iii. The average price over periods 2 to 11.
 - iv. The difference between the price in period 3 and in period 11.
 - v. The difference between the price in period 2 and in period 11.
 - vi. None of the above.
- (b) Compute the average price over period 3 to 12:
(USE THE CALCULATOR ON YOUR DESK!)
- (c) Assume that, in period 2, you predicted an average price over periods 3 to 12 of 82.5.
 - i. What is your forecast error?
 - ii. How many forecasting points do you earn?
(USE YOUR PAYOFF TABLE ON YOUR DESK!)

2. What is the probability for all the chickens to die if you are entering...

- (a) ... period 6?
- (b) ... period 45?

3. If the chickens die in period 23, when will you find out?

4. If you forecast a high average price over the next 10 periods, ...

N.B.: Multiple answers may be possible.

- (a) Your forecast will not impact your demand for chickens.
- (b) Your trader is likely to buy chickens.
- (c) Your trader is likely to sell chickens.
- (d) You are likely to consume fewer eggs now but more later on.
- (e) You are likely to consume more eggs now but fewer later on.
- (f) This also depends on the other participants' forecasts.

5. If all participants tend to forecast an increase in the price over the next 10 periods compared to past levels, what is the implication on the realized current market price?

N.B.: Multiple answers may be possible.

- (a) The realized market price is likely to increase.
- (b) The realized market price is likely to decrease.
- (c) The realized market price is likely to remain stable.
- (d) The realized market price will not be impacted by participants' forecasts.

Forecasting payoff table

$$\text{Your payoff} : 1100 - \frac{1100}{49} (\text{Your forecast error})^2$$

2000 points = 1 euro

error	points	error	points	error	points	error	points
0	1100	1.85	1023	3.7	793	5.55	409
0.05	1100	1.9	1019	3.75	784	5.6	396
0.1	1100	1.95	1015	3.8	776	5.65	383
0.15	1099	2	1010	3.85	767	5.7	371
0.2	1099	2.05	1006	3.9	759	5.75	358
0.25	1099	2.1	1001	3.95	750	5.8	345
0.3	1098	2.15	996	4	741	5.85	332
0.35	1097	2.2	991	4.05	732	5.9	319
0.4	1096	2.25	986	4.1	723	5.95	305
0.45	1095	2.3	981	4.15	713	6	292
0.5	1094	2.35	976	4.2	704	6.05	278
0.55	1093	2.4	971	4.25	695	6.1	265
0.6	1092	2.45	965	4.3	685	6.15	251
0.65	1091	2.5	960	4.35	675	6.2	237
0.7	1089	2.55	954	4.4	665	6.25	223
0.75	1087	2.6	948	4.45	655	6.3	209
0.8	1086	2.65	942	4.5	645	6.35	195
0.85	1084	2.7	936	4.55	635	6.4	180
0.9	1082	2.75	930	4.6	625	6.45	166
0.95	1080	2.8	924	4.65	615	6.5	152
1	1078	2.85	918	4.7	604	6.55	137
1.05	1075	2.9	911	4.75	593	6.6	122
1.1	1073	2.95	905	4.8	583	6.65	107
1.15	1070	3	898	4.85	572	6.7	92
1.2	1068	3.05	891	4.9	561	6.75	77
1.25	1065	3.1	884	4.95	550	6.8	62
1.3	1062	3.15	877	5	539	6.85	47
1.35	1059	3.2	870	5.05	527	6.9	31
1.4	1056	3.25	863	5.1	516	6.95	16
1.45	1053	3.3	856	5.15	505	error ≥ 7	0
1.5	1049	3.35	848	5.2	493		
1.55	1046	3.4	840	5.25	481		
1.6	1043	3.45	833	5.3	469		
1.65	1039	3.5	825	5.35	457		
1.7	1035	3.55	817	5.4	445		
1.75	1031	3.6	809	5.45	433		
1.8	1027	3.65	801	5.5	421		

Consumption payoff table

Your payoff : $\log(\text{your egg consumption}) \times 120$

2000 points = 1 euro

eggs	points	eggs	points	eggs	points	eggs	points
1	0	420	725	1000	829	12000	1127
2	83	440	730	1250	856	14000	1146
3	132	460	736	1500	878	16000	1162
4	166	480	741	1750	896	18000	1176
5	193	500	746	2000	912	20000	1188
6	215	520	750	2250	926	22000	1200
7	234	540	755	2500	939	24000	1210
8	250	560	759	2750	950	26000	1220
9	264	580	764	3000	961	28000	1229
10	276	600	768	3250	970	30000	1237
20	359	620	772	3500	979	32000	1245
30	408	640	775	3750	988	34000	1252
40	443	660	779	4000	995	36000	1259
50	469	680	783	4250	1003	38000	1265
60	491	700	786	4500	1009	40000	1272
70	510	720	790	4750	1016	42000	1277
80	526	740	793	5000	1022	44000	1283
90	540	760	796	5250	1028	46000	1288
100	553	780	799	5500	1034	48000	1293
120	575	800	802	5750	1039	50000	1298
140	593	820	805	6000	1044	52000	1303
160	609	840	808	6250	1049	54000	1308
180	623	860	811	6500	1054	56000	1312
200	636	880	814	6750	1058	58000	1316
220	647	900	816	7000	1062	60000	1320
240	658	920	819	7250	1067	62000	1324
260	667	940	822	7500	1071	64000	1328
280	676	960	824	7750	1075	66000	1332
300	684	980	827	8000	1078	68000	1335
320	692	1000	829	8250	1082	70000	1339
340	699	1020	831	8500	1086	72000	1342
360	706	1040	834	8750	1089	74000	1345
380	713	1060	836	9000	1093	76000	1349
400	719	1080	838	10000	1096	80000	1352

Once you have finished the experiment, please fill out this questionnaire!

Table number (letter and number on the yellow card): _____

Gender

- ☐ male
☐ female
☐ other

Age: _____

Nationality: _____

Which of the following *comes closest* to your field of study?

- ☐ Economics, Business
☐ Psychology, Social Sciences, Law, Humanities
☐ Mathematics, Physics, IT
☐ Medicine, Biology, Chemistry,
☐ Other: _____
☐ no studies

How would you describe your command of English?

- ☐ Excellent
☐ Very good
☐ Good
☐ Satisfactory
☐ Poor

How clear were the instructions of the experiment?

- ☐ Very clear
☐ Clear
☐ Understandable
☐ Slightly confusing
☐ Confusing

Have you participated in a *similar* economic experiment before?

- ☐ yes
☐ no

Did you perceive the length of the markets to be:

- ☐ As announced in the instructions
☐ Longer than announced in the instructions
☐ Shorter than announced in the instructions

Could you, in few words, summarize your strategy(ies) in this experiment?

If you would like to leave any comments for us, please do so here:
