Journal of Economic Dynamics & Control 180 (2025) 105181

Contents lists available at ScienceDirect

Journal of Economic Dynamics and Control

journal homepage: www.elsevier.com/locate/jedc

Learning to bet (rationally) with logs ☆

A. Carvajal [a], H. Zhou [b,*]

[a] University of California, Davis, United States of America
[b] School of Finance, Shanghai University of Finance and Economics, China

ARTICLE INFO

JEL classification:
D53
D84
G02
G14

Keywords:
Rational expectations
Level-k reasoning
Information revealed by prices
Bounded reasoning
Radner equilibrium

ABSTRACT

In an economy with uncertainty and asymmetric information, suppose that some agents learn the relation between fundamentals and prices by observing past market outcomes. They refine their understanding as they become more experienced, but their past "errors" contaminate the information they receive. Does this process converge to the "perfect" understanding of the market that underlies rational expectation equilibria? We address this question in a simplified setting that allows for explicit computation of the learning process: a two-state economy with logarithmic utilities and no background risk. Our first result is that as long as the wealth of the uninformed agents is less than half the aggregate wealth of the economy, the learning process indeed converges to rational expectations. This convergence, however, is non-monotonic, and the market oscillates between phases of excess price volatility and phases of excess volume of trade. The learning process, in addition, is costly for the uninformed agents. We interpret our results as underscoring the fragility of REE: markets operate orderly only when speculation is less significant than fundamental trade.

0. Introduction

Rational expectations equilibrium (REE, henceforth) has been an essential idea in economic theory, widely used in both microeconomics and macroeconomics. The seminal work by Muth (1961) introduces rational expectations to study how agents predict price movements, while Radner (1979) models REE as self-fulfilling beliefs, where agents maximize their utility based on their beliefs, and the market-clearing outcome confirms these beliefs. In particular, Radner models REE as a mapping from the set of states of the world to the set of prices and proves the generic existence and invertibility of a REE; prices are hence fully revealing, and all agents can figure out all the private information in the economy once they observe them.

However, two limitations of Radner (1979) are the implicit assumption that the rational mapping from states to prices is common knowledge among all agents and that the model is silent on how the agents would learn this mapping.[1] With this motivation, we incorporate adaptive learning into a general equilibrium setting and define an iterative process where the agents' mappings evolve as they learn and become "more sophisticated."

☆ We are grateful to the handling editor and two anonymous reviewers who pushed us to improve our paper and present it better. Hang Zhou acknowledges the funding support from the Fundamental Research Funds for the Central Universities, the National Natural Science Foundation of China [grant number: 72573104], and the Shanghai Rising-Star Program [grant number: 23QA1403200].

* Corresponding author.
E-mail addresses: acarvajal@ucdavis.edu (A. Carvajal), zhouhang@mail.shufe.edu.cn (H. Zhou).

[1] Dubey et al. (1987) further criticizes the REE approach in general equilibrium models with asymmetric information since it fails to explain how information gets encoded into asset prices.

https://doi.org/10.1016/j.jedc.2025.105181
Received 14 October 2024; Received in revised form 13 August 2025; Accepted 8 September 2025

The mathematics of our model resembles game-theoretical models of level-k reasoning. In those models, level-k reasoning is an alternative to Nash equilibrium that describes how strategic sophistication determines players' strategies.[2] Here, despite the resemblance, we do not interpret the evolution of beliefs as an issue of reasoning. Instead, we imagine the agents as learning through the observation of reality: after trading, the less informed agents discover the information they lack and use this observation to figure out the structure of the mapping that determines prices as a function of the available information. If these agents were fully rational, they would realize that just by learning about such a mapping, the relation between the two variables may change; indeed, REE is the formalization of the idea that the equilibrium mapping is one that is consistent with itself, in the sense that it is the mapping that results when all agents are in effect using it. On the other hand, our agents follow an adaptive learning process: once they learn a mapping, they start using it, and it takes them some time to observe the resulting relation between prices and information, namely the next mapping they will use.

Also, our results shed new light on the problem of information revelation with boundedly rational agents. We show that the level-k mapping, which links asset prices with market fundamentals, evolves in an oscillating manner. When the unsophisticated speculators' asset demand is insensitive to prices, the market-clearing prices become informative since the fundamentalists primarily drive market variations. Thus, the asset demand for sophisticated traders will be sensitive to prices, leading to excessive speculation and weakening asset prices' informativeness. Our mechanism generates oscillating behavior of price volatility as well.

1. Connections to the literature

This paper provides an answer to the question of convergence to REE by incorporating adaptive learning into a general equilibrium model. We show that convergence to REE requires the informed traders' aggregate asset demand to be more responsive to asset prices than that of the uninformed speculators.

Addressing the same difficulty of REE as here, McAllister (1990) incorporates rational expectations equilibrium in a decision-making framework and constructs the rational expectation mapping at the individual level. That author considers the space of uncertainty for each agent to be the product of the state space and the set of all possible asset positions of other traders. An REE consists of an admissible prior, a price vector, and optimal asset positions that clear all the markets. Later, Dutta and Morris (1997) generalized McAllister (1990) by relaxing the assumptions of common knowledge of consensus beliefs and degeneracy of expectations. They introduce the concepts of *belief equilibrium*, where agents may disagree on their prior beliefs, and of *common belief equilibrium*, where agents hold the same beliefs. There, an REE is a restriction to common belief equilibrium where the agents only consider the exogenous states of nature, and the mapping from states to prices is a consensus between all agents. In this line of research, Ben-Porath and Heifetz (2011) present, to the best of our knowledge, the most general result on the epistemic foundation of rational expectations equilibrium literature. They show that common knowledge of rationality and market clearing is sufficient to yield REE.

Unlike the literature above, this paper addresses how the agents learn the rational expectations mapping. There have been other results on the convergence to rational expectations in both the macro and the general equilibrium literature. Shiller (1978), for instance, studies convergence to the rational expectation forecast in Muth (1961), while DeCanio (1979) studies convergence to rational expectations in a linear forecast model with general autocorrelation structure. Closer to our paper, Bray (1982) studies a setting where agents learn the relation between asset returns and prices using OLS estimation. Bray's results suggest that the learning process could converge to rational expectations even if agents use misspecified models. Blume and Easley (1984) study a dynamic market process in which agents learn a payoff-relevant parameter by conditioning on past endogenously generated data. They define REE as the limit of the learning process once all the agents' beliefs converge to the true parameter almost surely.

Our paper is closely related to Grandmont (1998), which examines an adaptive learning process in which agents are uncertain about the stability of the equilibrium and extrapolate from past trends. The paper demonstrates that near-equilibrium dynamics play a critical role in the convergence to REE. Convergence is assured only if agents do not excessively extrapolate past trends. In comparison, our findings suggest that convergence (or divergence) occurs when speculators are less (or more) responsive than fundamentalists. Both papers highlight the potential link between divergence and the excessive amount of adaptive behavior.

This paper is also related to Feldman (1987) and Pesce et al. (2024), which study the stability of the REE.[3] Feldman (1987) shows that convergence to REE is possible even with heterogeneous initial belief. Unlike here, the agents in that paper are perfectly Bayesian but have yet to discover the values of some parameters. Pesce et al. (2024) study an economy where agents update their private information based on the observed REE prices and allocations. They, too, show convergence to the symmetric information REE, but their learning process is different: the agents always infer correct information and become more refined across rounds. In our case, no round of trade needs to result in an REE, and the agents here may infer wrong information from the prices they observe. The emphasis in Pesce et al. (2024) is on the convergence of their sequence of REE to a symmetric REE. For us, the focus is on whether a sequence of non-REE price functions converges to an REE.

Other researchers have applied the level-k framework to macroeconomics. Farhi and Werning (2019), for instance, explain the forward guidance puzzle with level-k thinking agents, whereas Angeletos and Lian (2017) argue that level-k thinking explains the

---

[2] Importantly, this approach is supported by ample experimental evidence—see Nagel (1995), Stahl and Wilson (1994) and Stahl and Wilson (1995). Crawford et al. (2013) provides a thorough review on level-k thinking and supporting experimental evidence.

[3] Blume et al. (1982) survey the literature on convergence to the REE.

slow adjustments of general equilibrium effects. These studies, however, do not determine conditions that guarantee convergence, and Farhi and Werning (2019) in fact suggest that convergence to rational expectations occurs in a monotonic way.[4]

Our paper investigates the convergence of the "level-k" mapping in a general equilibrium setting, and our results are that convergence to REE (1) is guaranteed only under restrictive assumptions, and (2) if it occurs, it does so in a non-monotonic way. In the same vein, Zhou (2022) studies a static asset market with imperfect competition and speculators that perform level-k thinking, and shows that convergence to Nash equilibrium (also the REE) requires that the prices be less informative than the private information. This paper differs from Zhou (2022) in that we incorporate adaptive learning into a general equilibrium setting, focusing on how agents iteratively update their beliefs and mappings from prices to fundamentals. We explicitly model the learning process and its convergence properties, highlighting the oscillatory dynamics and the role of wealth distribution, whereas Zhou (2022) analyzes static strategic sophistication without a learning mechanism.

The key concern in our paper is how traders extract fundamental information from asset prices. REE is the very special case where all traders correctly understand the equilibrium mapping between prices and fundamental values, and trade based on the correct mapping. In the real financial markets, however, traders make mistakes. Previous research has studied various traders' behavioral biases. For instance, Eyster et al. (2019) consider traders who fail to fully appreciate what prices convey about others' private information. They show that in the case of symmetric traders with no random endowments and public signals, cursedness leads to substantial trade. Trading volume increases with the number of traders, and prices underreact to private signals, showing positive autocorrelation. Mondria et al. (2022) study a rational inattention model where traders pay a cost to interpret signals about the fundamental value. They find that investors face costs in interpreting asset price information, leading to noise in the price system, price momentum, excessive return volatility, and excessive trading volume. Besides, there is strategic complementarity in acquiring sophistication for taming sentiment, which may result in multiple equilibria. In a recent paper, Bastianello and Fontanier (2025) develop a theory of "partial equilibrium thinking" and study the mislearning from prices when traders are subject to partial equilibrium thinking. In that setting, each trader learns fundamental information from prices but does not realize that other traders are doing the same. This paper is related to ours as we assume that the speculators in our model fail to consider others' learning from prices while they update the mapping from prices to states.

2. A simple Radner economy

Consider a two-period exchange economy with uncertainty where the state space for the future period is $\{1,2\}$, there is only one commodity, and consumption takes place only in the second period.

There are two types of agents in the market. *Fundamentalists* know that the probability that future state $s=1$ will realize is $\pi \in \Delta$, while *Speculators* only know that π is a realization of random variable Π with support Δ. Importantly, we assume that the distribution of Π is common knowledge to all traders.

We will use the super-index $a \in \{F, S\}$ to denote each agent's type. Besides their information, the two types may differ in their preferences and future wealth. In future state s, agents of type a will be endowed with wealth ω^a_s. Ex-ante, they are expected utility maximizers with instantaneous utility function $u^a(x)$, which is assumed to be C^2, strictly increasing, and strictly concave. There is a continuum of agents of each type, with respective masses μ^F and μ^S.

In the present, the agents trade the elementary securities corresponding to the two future states. We normalize the price of the security for state $s=2$ to unity and denote by q the price of the security for $s=1$. When an agent of type a chooses her portfolio, the only constraint she faces is that $q \cdot y^a_1 + y^a_2 = 0$, where y^a_s denotes her holdings of the security that pays in state s. Market clearing requires that $\mu^F y^F_1 + \mu^S y^S_1 = 0$ and $\mu^F y^F_2 + \mu^S y^S_2 = 0$.

3. Learning to bet with logs

To obtain explicit computations, suppose that: (1) the agents face no background risk, in the sense that $\omega^a_1 = \omega^a_2 = \omega^a$ for both types; (2) both types have logarithmic preferences, $u^a(x) = \ln(x)$; and (3) Π follows the uniform distribution over the support $\Delta = [1/2 - \delta, 1/2 + \delta]$, where $0 \le \delta < 1/2$.

For reasons that will be clear later on, we maintain the assumptions that $\mu^F \omega^F > \mu^S \omega^S$. Also, to avoid corner solutions in prices, we maintain the assumption that

$$\delta < \frac{1}{2}\left[1 - \left(\frac{\mu^S \omega^S}{\mu^F \omega^F}\right)^2\right]. \tag{1}$$

---

[4] To be sure, that paper shows that the response of output and inflation to policy shocks under the level-k thinking process, for instance, converges to that of a standard REE model. Such convergence is monotonic in the sense that the economic responses unidirectionally approach the REE benchmark: there is no oscillation. In Farhi and Werning (2019), if prices are assumed to be a fixed constant, convergence is an unconditional property of the level-k thinking process. However, when the model incorporates inflation and sticky prices (a more realistic scenario), an important condition for convergence emerges. It is no longer guaranteed by the level-k process alone. The monetary authority must follow a "sufficiently responsive interest rate rule" (e.g., a Taylor rule that actively adjusts rates in response to inflation). Without such a rule, the level-k equilibria may not converge to a determinate rational-expectations equilibrium. No such conditions emerge in our paper.

3.1. Individual demands

At price q, when they know that the probability of state $s = 1$ is π, each fundamentalist demands $y^F(q,\pi)$ units of the asset that pays in that state, which solves

$$\max_y \left\{\pi \ln(\omega^F + y) + (1-\pi)\ln(\omega^F - qy)\right\}.$$

Direct computation yields

$$y^F(q,\pi) = \left(\pi \cdot \frac{q+1}{q} - 1\right)\omega^F, \tag{2}$$

which is decreasing in q and increasing in π. At the same prices, a speculator whose information is only that $\pi \in \mathcal{I} \subseteq \Delta$ demands

$$\begin{aligned} y^S(q,\mathcal{I}) &= \arg\max_y \left\{\mathrm{E}\left[\Pi \ln(\omega^S + y) + (1-\Pi)\ln(\omega^S - qy) \mid \mathcal{I}\right]\right\} \\ &= \left[\mathrm{E}(\Pi \mid \mathcal{I}) \cdot \frac{q+1}{q} - 1\right]\omega^S. \end{aligned} \tag{3}$$

This function, too, is decreasing in q and increasing in $\mathrm{E}(\Pi \mid \mathcal{I})$.

In addition to guaranteeing market clearing, a definition of equilibrium must stipulate how $\mathcal{I}$ is determined as a function of q.

3.2. Rational expectations equilibrium

In the current setting, an REE is a function $\bar{q}$, mapping probabilities into prices, such that, for every π, markets clear at price $\bar{q}(\pi)$ when the speculators condition on $\mathcal{I} = \bar{q}^{-1}(\bar{q}(\pi))$. The equilibrium is said to be fully revealing if it is injective, namely if $\bar{q}^{-1}(\bar{q}(\pi)) = \{\pi\}$.

If one conjectures that there is a fully revealing REE, $\bar{q}(\pi)$ must solve the market clearing condition $\mu^F y^F(\bar{q}(\pi),\pi) = -\mu^S y^S(\bar{q}(\pi),\{\pi\})$, which gives

$$\bar{q}(\pi) = \frac{\pi}{1-\pi}. \tag{4}$$

Since this function is injective, the REE is confirmed to be fully revealing. Since both demand functions are decreasing in q and increasing in π, this is the only REE. For future reference,

$$\bar{q}'(\pi) = \frac{1}{(1-\pi)^2}. \tag{5}$$

Once we take into account the randomness of the probability distribution over future states, the REE price is the random variable $\bar{Q} := \bar{q}(\Pi)$. The distribution of this random variable is

$$\bar{G}(q) := \Pr(\bar{Q} \leq q) = \Pr(\Pi \leq \bar{q}^{-1}(q)) = \Gamma\left(\frac{q}{1+q}\right)$$

where Γ is the CDF of Π.

Under the assumption that Π follows the uniform distribution over $\Delta = [1/2 - \delta, 1/2 + \delta]$, this distribution turns out to be a (shifted, re-scaled) Beta $(1,3)$ distribution.[5] This is a very tractable function, which we describe in Table 1. There, we consider two location parameters (the median and the mean) and two spread parameters (the range[6] and the variance).

Note that the traders' endowments are immaterial for the distribution of $\bar{Q}$. This is so because all the agents have the same homothetic preference and the REE is fully revealing. Also, since complete revelation takes place at equilibrium, there is no trade: for all π, $y^S(\bar{q}(\pi),\{\pi\}) = 0 = y^F(\bar{q}(\pi),\pi)$.

3.3. Completely inexperienced speculators

A versed speculator would understand that the fundamentalists' asset demands, and hence the market-clearing prices, depend on the realized probability π. We model the least experienced version of a speculator as one who completely fails to recognize such a dependence. We call this level of naïveté *level-*0 *understanding*, and define the corresponding demand of this speculator as $y_0^S(q) = y^S(q;\Delta)$. By linearity, that is

$$y_0^S(q) = \arg\max_y \left\{\mathrm{E}(\Pi)\ln(\omega^S + y) + [1 - \mathrm{E}(\Pi)]\ln(\omega^S - qy)\right\} = \left(\frac{q+1}{2q} - 1\right)\omega^S,$$

since, under the simplifying assumptions, the (unconditional) expectation of probability Π is $1/2$.

---

[5] Which is also the Kumaraswamy $(1,3)$ double-bounded distribution.

[6] Namely, the difference between the extrema of the support.

Table 1
The probability distribution of the REE price, $\bar{Q}$.

Support	$\bar{S} := \left[\frac{1-2\delta}{1+2\delta}, \frac{1+2\delta}{1-2\delta}\right]$
CDF	$\bar{G}(q) := \frac{1}{2\delta}\left(\frac{q}{1+q} + \delta - \frac{1}{2}\right)$
PDF	$\bar{g}(q) := \frac{1}{2\delta(1+q)^2}$
Median	1
Mean	$\bar{m} := \frac{1}{2\delta}\ln\left(\frac{1+2\delta}{1-2\delta}\right) - 1$
Range	$\bar{r} := \frac{8\delta}{1-4\delta^2}$
Variance	$\bar{v} := \frac{4}{1-4\delta^2} - \frac{1}{4\delta^2}\ln^2\left(\frac{1+2\delta}{1-2\delta}\right)$

With all speculators at this level of understanding, equilibrium requires that the corresponding price function $q_0 : \Delta \to \mathbb{R}$ be the solution to $\mu^F y^F(q_0(\pi), \pi) = -\mu^S y_0^S(q_0(\pi))$ at each π. Denoting the ratio $(\mu^F \omega^F)/(\mu^S \omega^S) =: \rho$, by direct computation

$$q_0(\pi) = \frac{1 + 2\pi\rho}{1 + 2(1-\pi)\rho}, \tag{6}$$

which is positive and injective. As before, with the monotonicity of both demand functions, this is the only market-clearing mapping. Unlike the REE, though, this function depends on the endowments of the two agent types.

As before, once we take into account the randomness of the probability distribution over future states, the market-clearing price is the random variable $Q_0 := q_0(\Pi)$. The distribution of this random variable is

$$G_0(q) := \Pr(Q_0 \leq q) = \Pr(\Pi \leq q_0^{-1}(q)) = \Gamma\left(\frac{-1 + (1+2\rho)q}{2\rho(1+q)}\right),$$

recalling that Γ is the CDF of Π.

3.4. Adaptive learning by the speculators

To model the evolution of the speculators' understanding of the relation between prices and probabilities, it is convenient to see these agents as a sequence of non-overlapping generations of individuals, and assume that each generation participates in T rounds of trade, starting from the level-0 generation we just described.

Suppose that the sequence of independent realizations of Π, across the T rounds of trade in which this generation participates, is $\langle \pi_t \rangle_{t=1}^T$. At round t, the resulting market-clearing price is $q_0(\pi_t)$, so the history of resulting prices is the data set $D_0 := \langle q_0(\pi_t) \rangle_{t=1}^T$. For any price q, the speculators can compute the empirical distribution

$$\hat{G}_{0,T}(q) = \tfrac{1}{T} \textstyle\sum_{t=1}^T \mathbb{1}_{q_0(\pi_t) \leq q}.$$

By the Glivenko-Cantelli theorem, $\hat{G}_{0,T} \to G_0$ uniformly, almost surely, as $T \to \infty$. We assume that T is large enough for the data set D_0 to identify the distribution G_0. Since both G_0 and the distribution of Π are monotonic and function Γ is commonly known, the next generation of speculators can identify the market-clearing price function by $q_0(\pi) = G_0^{-1}(\Gamma(\pi))$. The following is the commutative diagram for this construction[7]:

$$\begin{tikzcd}
\Delta \arrow[r, "q_0"] \arrow[d, "\Gamma"] & q_0[\Delta] \arrow[dl, "G_0 = \Gamma \circ q_0^{-1}"] \\
(0,1) &
\end{tikzcd}$$

Having discovered the function $q_0(\pi)$, this second generation of speculators infers the information $\mathcal{I} = q_0^{-1}(q)$, upon observation of q, when choosing their optimal demand according to (3). This gives a new demand function

$$y_1^S(q) = \arg\max_y \left\{ \mathrm{E}(\Pi \mid \mathcal{I}) \ln(\omega^S + y) + [1 - \mathrm{E}(\Pi \mid \mathcal{I})] \ln(\omega^S - qy) \right\} = \left(q_0^{-1}(q) \cdot \frac{q+1}{q} - 1 \right) \omega^S,$$

so the market-clearing condition $\mu^F y^F(q_1(\pi), \pi) = -\mu^S y_1^S(q_1(\pi))$ requires a new price function, $q_1(\pi)$.

Given a new sequence of i.i.d. realizations of Π for the rounds of trade in which this generation participates, say $\langle \hat{\pi}_t \rangle_{t=1}^T$, the resulting history of prices is the data set $D_1 = \langle q_1(\hat{\pi}_t) \rangle_{t=1}^T$. As $T \to \infty$, the empirical distribution $\hat{G}_{1,T}$ of this experiment converges

---

[7] Algorithmically, an observer must construct the one-to-one correspondence between each of the p-quantiles, $G_0^{-1}(p)$ and $\Gamma^{-1}(p)$, of the two distributions.

almost surely to the actual distribution of market-clearing prices, which we will denote G_1. Once this object has been identified, the next generation of speculators learns $q_1 = G_1^{-1} \circ \Gamma$.

Recursively, the k^{th} generation of speculators has discovered the equilibrium mapping $q_{k-1}(\pi)$ and uses the information $\mathcal{I} = q_{k-1}^{-1}(q)$, which results in a new market-clearing price function $q_k(\pi)$. The $T \to \infty$ rounds of trade of this generation yield a data set $\boldsymbol{D}_k$ of market-clearing prices, with empirical distribution $\hat{G}_{k,T}$. We assume that $T \to \infty$, which by the law of large numbers converges to the actual market-clearing distribution G_k. Once again, the next generation thus identifies $q_k = G_k^{-1} \circ \Gamma$.[8]

This adaptive learning process resembles level-k thinking in the game theory literature, although speculators do not explicitly model others' strategic choices. Two key features emerge. First, when a level-k speculator updates to the $k+1^{\text{st}}$-generation model, this process effectively approximates a best-response to a conjectured strategy profile where all other speculators remain at the level-k. This mimics the recursive reasoning structure of level-k models without requiring explicit strategic conjectures. Second, the adaptation process inherently incorporates *partial equilibrium thinking*, as speculators abstract from the possibility that others might also be adapting. This adjustment process is similar to the level-k thinking's iterated best response structure.

3.5. *Learning and the evolution of prices*

For any natural number k, we say that a speculator has *level-k understanding* if, at prices q, she uses information $\mathcal{I} = q_{k-1}^{-1}(q)$ in her choice of an optimal portfolio. This is an investor who understands the dependence of prices on the information available to the fundamentalists through the price function $q_{k-1}(\pi)$, so her optimal demand for the first security is $y_k^S(q) = y^S(q, q_{k-1}^{-1}(q))$. This demand gives raise to a new pricing function, $q_k : \Delta \to \mathbb{R}$, defined implicitly by the market clearing condition $\mu^F y^F(q_k(\pi); \pi) = -\mu^S y_k^S(q_k(\pi))$, assuming that such a function exists.[9]

Starting from Eq. (6) and using mathematical induction over k, we obtain

$$q_k(\pi) = \frac{(-1)^k + 2\pi\rho^{k+1}}{(-1)^k + 2(1-\pi)\rho^{k+1}}. \tag{7}$$

This function is positive and injective,[10] and is the only market-clearing mapping since both demand functions are decreasing in q. Also,

$$q_k'(\pi) = \frac{4\rho^{k+1}[(-1)^k + \rho^{k+1}]}{[(-1)^k + 2(1-\pi)\rho^{k+1}]^2}. \tag{8}$$

As before, once the probability distribution's randomness is considered, the level-k equilibrium prices induce the random variable $Q_k := q_k(\Pi)$. As in the case of REE, this random variable follows a (shifted, re-scaled) Beta $(1,3)$ distribution,

$$\begin{aligned} G_k(q) &= \Pr(Q_k \le q) \\ &= \Pr\left(\Pi \le \frac{(-1)^{k+1} + [(-1)^k + 2\rho^{k+1}]q}{2\rho^{k+1}(1+q)}\right) \\ &= \Gamma\left(\frac{(-1)^{k+1} + [(-1)^k + 2\rho^{k+1}]q}{2\rho^{k+1}(1+q)}\right), \end{aligned}$$

which we describe in Table 2.[11]

Note again that, unlike in the REE, the distribution of level-k prices depends on the relative aggregate wealth of the fundamentalists to the speculators, ρ.

4. Beliefs bias while learning

During the learning process, the speculators fail to realize that the mapping they use to discern information for market prices becomes outdated once they start using it. When they have level-k understanding and the probability is π, markets clear at the price $q_k(\pi)$ but the speculators infer the probability $q_{k-1}^{-1}(q_k(\pi))$. We call the difference the speculators' bias:

$$b_k(\pi) = \begin{cases} \frac{1}{2} - \pi, & \text{if } k = 0; \\ q_{k-1}^{-1}(q_k(\pi)) - \pi, & \text{if } k \ge 1. \end{cases}$$

---

[8] Using the language of machine learning, while the k^{th} model is being used, the next generation model is being *trained*, using the data set D_k.

[9] Importantly, the speculators do *not* realize that their use of function q_{k-1} changes the equilibrium prices at each value of π—namely, that it induces the new mapping q_k.

[10] When k is odd, Eq. (1) and $\rho > 1$ guarantee these properties. When it is even, the properties hold regardless of these assumptions.

[11] The very devoted reader may want to verify the identification argument given in Section 3.4. To make their life easy, note that the p-quantile of Q_k is

$$G_k^{-1}(p) = \frac{(-1)^k + [1 + 2\delta(2p-1)]\rho^{k+1}}{(-1)^k + [1 - 2\delta(2p-1)]\rho^{k+1}}.$$

Evaluating this function at $p = \Gamma(\pi) = (\pi + \delta - 1/2)/2\delta$ yields $q_k(\pi)$.

Table 2
The probability distribution of Q_k.

Support	$\mathcal{S}_k := \left[\frac{(-1)^k + (1-2\delta)\rho^{k+1}}{(-1)^k + (1+2\delta)\rho^{k+1}}, \frac{(-1)^k + (1+2\delta)\rho^{k+1}}{(-1)^k + (1-2\delta)\rho^{k+1}}\right]$
CDF	$G_k(q) := \frac{1}{2\delta}\left\{\frac{(-1)^{k+1} + [(-1)^k + 2\rho^{k+1}]q}{2\rho^{k+1}(1+q)} + \delta - \frac{1}{2}\right\}$
PDF	$g_k(q) := \frac{2(-1)^k\rho^{k+1}}{2\delta\rho^{k+1}(1+q)^2}$
Median	1
Mean	$m_k := \frac{(-1)^k + \rho^{k+1}}{2\delta\rho^{k+1}} \ln\left[\frac{(-1)^k + (1+2\delta)\rho^{k+1}}{(-1)^k + (1-2\delta)\rho^{k+1}}\right] - 1$
Range	$r_k := \frac{8\delta\rho^{k+1}[(-1)^k + \rho^{k+1}]}{1 + (-1)^k 2\rho^{k+1} + (1-4\delta^2)\rho^{2k+2}}$
Variance	$v_k := \frac{4[(-1)^k+\rho^{k+1}]^2}{[(-1)^k+(1+2\delta)\rho^{k+1}][(-1)^k+(1-2\delta)\rho^{k+1}]} - \frac{[(-1)^k+\rho^{k+1}]^2}{4\delta^2\rho^{2(k+1)}} \ln^2\left[\frac{(-1)^k+(1+2\delta)\rho^{k+1}}{(-1)^k+(1-2\delta)\rho^{k+1}}\right].$

Our first observation is that, across levels of understanding, the sign of the bias oscillates for most values of the probability:

Proposition 1. *If k is even and $\pi > 1/2$, or k is odd and $\pi < 1/2$, then $b_k(\pi) > 0$. If k is even and $\pi < 1/2$, or k is odd and $\pi > 1/2$, then $b_k(\pi) < 0$. For all k, $b_k(1/2) = 0$.*

Proof. By direct computation,

$$b_k(\pi) = \frac{(-1)^k(1-2\pi)(1+\rho^k)}{2[(-1)^k + \rho^{k+1}]}. \quad \square$$

This observation underlies much of the intuition for the results that follow.

5. Convergence to rational expectations

Given $\pi \in \Delta$, the deviation of level-k prices from REE is

$$q_k(\pi) - \bar{q}(\pi) = \frac{(-1)^k(1-2\pi)}{(1-\pi)[(-1)^k + 2(1-\pi)\rho^{k+1}]}. \tag{9}$$

It follows immediately that:

Proposition 2. *If k is even and $\pi < 1/2$, or k is odd and $\pi > 1/2$, then $q_k(\pi) - \bar{q}(\pi) > 0$. If k is even and $\pi > 1/2$, or k is odd and $\pi < 1/2$, then $q_k(\pi) - \bar{q}(\pi) < 0$. For all k, $q_k(1/2) - \bar{q}(1/2) = 0$.*

Proof. The denominator on the right-hand side of Eq. (9) is unambiguously positive, given Eq. (1). The premises of the proposition determine the sign of the numerator. $\square$

Asymptotically, moreover,

Proposition 3. *For all π, $q_k(\pi) - \bar{q}(\pi) \to 0$.*

Proof. Since, by assumption, $\rho > 1$, we have that $\rho^{k+1} \to \infty$ and then, by Eq. (9), $|q_k(\pi) - \bar{q}(\pi)| \to 0$ as $k \to \infty$. $\square$

Importantly, since Δ is compact, the second of these results implies that the sequence of level-k price functions converges uniformly to the REE. The first result, on the other hand, implies that convergence occurs in an oscillatory manner. In fact,

Proposition 4. $\pi \gtreqless \frac{1}{2} \Leftrightarrow q_0(\pi) \lesseqgtr q_2(\pi) \lesseqgtr q_4(\pi) \lesseqgtr \ldots \lesseqgtr \bar{q}(\pi) \lesseqgtr \ldots \lesseqgtr q_5(\pi) \lesseqgtr q_3(\pi) \lesseqgtr q_1(\pi)$.

Proof. This follows by direct computation, using Eq. (9) and

$$q_k(\pi) - q_{k-2}(\pi) = \frac{(-1)^k(1-2\pi)(1-\rho)^2\rho^{k-1}}{1/2 + (-1)^k(1-\pi)(1+\rho^2)\rho^{k-1} + 2(1-\pi)^2\rho^{2k}}. \quad \square$$

Intuitively, the mechanism is as follows. Suppose that the fundamentalists learn the realization $\pi > \mathrm{E}(\Pi)$. When the speculators have level-0 understanding, they fail to understand that high prices occur because the fundamentalists have observed high values

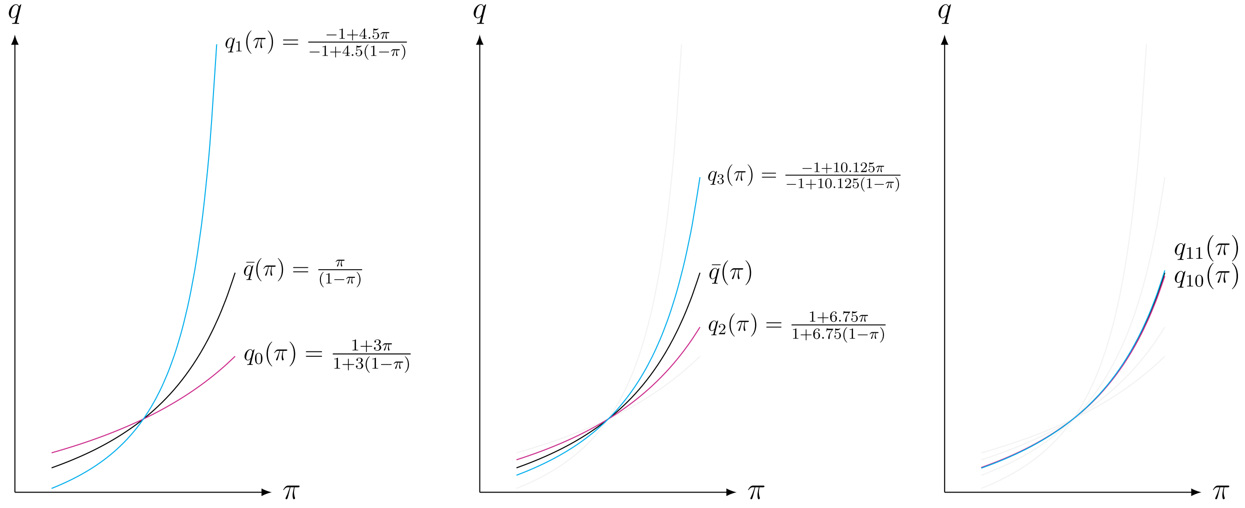

Fig. 1. Evolution of level-k mappings for $\rho = 3/2$ and $\delta = 1/4$. The first two graphs illustrate Propositions 2, 4, and 5. The third graph illustrates Proposition 3.

of π. Absent the feedback effect that would arise from belief updating, the asset's price need not increase as much as it would under REE. The q_0 mapping is less responsive to variations in π than $\bar{q}$: by direct comparison of Eqs. (5) and (8), $q_0'(\pi) \leq \bar{q}'(\pi)$ at all π. (See also the first graph in Fig. 1.)

Consider now what happens when the speculators use q_0 to interpret market signals. For them, a slight price increase is associated with a significantly higher probability. To deter them from increasing their position in the asset, a large substitution effect – an even higher price – is hence needed.[12] As a consequence, q_1 is *more* responsive to the probability than $\bar{q}$: $q_1'(\pi) \geq \bar{q}'(\pi)$. (See, again, the first graph in Fig. 1.)

Level-2 speculators then use a very responsive mapping q_1. In their beliefs, a significant increase in the price is consistent with a slight increase in the probability observed by the speculators, so their demands become unresponsive to price changes. As in the case of the level-0 mapping, the result is that the q_2 mapping is less responsive to variations in π than $\bar{q}$, although more than q_0. (Compare now the first and second graphs in Fig. 1.)

Following this intuition, even levels of understanding display unresponsive price mappings, which are followed by more responsive mappings at odd levels of understanding:

Proposition 5. *For all* $\pi \in \Delta$, $q_0'(\pi) \leq q_2'(\pi) \leq q_4'(\pi) \leq \ldots \leq \bar{q}'(\pi) \leq \ldots \leq q_5'(\pi) \leq q_3'(\pi) \leq q_1'(\pi)$.

Proof. Equations (5) and (8) imply that

$$q_k'(\pi) - \bar{q}'(\pi) = \frac{4(-1)^{k+1}\pi(1-\pi)\rho^{k+1} - 1}{(1-\pi)^2[(-1)^k + 2(1-\pi)\rho^{k+1}]^2}$$

and

$$q_k'(\pi) - q_{k-2}'(\pi) = \frac{4(-1)^k[(1+4\pi(1-\pi)\rho^{2k}]\rho^{k-1}(1-\rho^2)}{[(-1)^k + 2(1-\pi)\rho^{k+1}]^2[(-1)^k + 2(1-\pi)\rho^{k-1}]^2}.$$

The result follows by direct computation. □

The bounds we imposed on the probability distribution guarantee that all the mappings are positive and increasing. Combined with the fact that the fundamentalists own more wealth in the aggregate than the speculators, these properties yield an oscillating but converging learning process.

Unfortunately, all the market-clearing price functions depend on the probability π nonlinearly. For the REE case, the mean of $\bar{Q}$ equals $\bar{q}(\mathrm{E}(\Pi))$ only in the limit as $\delta \downarrow 0$. This feature is problematic since $\bar{Q}$ is the relative price of two assets being traded in a symmetric setting, so one could expect their relative price to be equal to one. This problem can be solved through a different normalization for a given model.[13] In our case, this solution would require different normalizations for each level of understanding and for the REE solution, which seems ad-hoc.

---

[12] There exists a tension between the information effect and the substitution effect, but the assumption that the aggregate wealth of the fundamentalists is higher, namely that $\rho > 1$, means that the substitution effect dominates.

[13] For example, in the REE case we could choose the constant $\kappa = \mathrm{E}\,(\Pi/1-\Pi)$ to be the price of the second asset.

Absent these normalizations, the expected equilibrium price has ambiguous comparative statics. Note that

$$m_k \geq \bar{m} \Leftrightarrow \frac{(-1)^k + \rho^{k+1}}{\rho^{k+1}} \cdot \ln\left[\frac{(-1)^k + (1+2\delta)\rho^{k+1}}{(-1)^k + (1-2\delta)\rho^{k+1}}\right] \geq 1 \cdot \ln\left(\frac{1+2\delta}{1-2\delta}\right).$$

However,

$$\frac{(-1)^k + \rho^{k+1}}{\rho^{k+1}} \geq 1 \Leftrightarrow k \text{ is even} \Leftrightarrow \frac{(-1)^k + (1+2\delta)\rho^{k+1}}{(-1)^k + (1-2\delta)\rho^{k+1}} \leq \frac{1+2\delta}{1-2\delta},$$

so the determination of the sign of the difference $m_k - \bar{m}$ in closed form results elusive.

The other measure of the location of the distributions of market-clearing prices, their median, does not exhibit the same behavior. At REE and all levels of understanding, the median of equilibrium prices equals unity. Statistically, this occurs simply because the price functions are increasing.[14]

6. Learning and price dispersion

Tables 1 and 2 report the supports and two measures of the dispersion of equilibrium prices: their range and their variance. Considering first the extrema of the supports,

Proposition 6. $\max \mathcal{S}_k \geq \max \mathcal{S}_{k+2} \geq \max \bar{\mathcal{S}}$ *and* $\min \mathcal{S}_k \leq \min \mathcal{S}_{k+2} \leq \min \bar{\mathcal{S}}$ *if, and only if,* k *is odd.*

Proof. Using Tables 1 and 2, note that

$$\max \mathcal{S}_k \geq \max \bar{\mathcal{S}} \Leftrightarrow \frac{(-1)^k + (1+2\delta)\rho^{k+1}}{(-1)^k + (1-2\delta)\rho^{k+1}} \geq \frac{1+2\delta}{1-2\delta} \Leftrightarrow (-1)^k(1-2\delta) \geq (-1)^k(1+2\delta),$$

while

$$\max \mathcal{S}_k \geq \max \mathcal{S}_{k+2} \Leftrightarrow \frac{(-1)^k + (1+2\delta)\rho^{k+1}}{(-1)^k + (1-2\delta)\rho^{k+1}} \geq \Leftrightarrow \frac{(-1)^k + (1+2\delta)\rho^{k+3}}{(-1)^k + (1-2\delta)\rho^{k+3}} \Leftrightarrow (-1)^k \geq (-1)^k \rho^2.$$

The computations for the minima are similar. □

It follows immediately that:

Corollary 1. $\mathcal{S}_0 \subseteq \mathcal{S}_2 \subseteq \mathcal{S}_4 \subseteq \ldots \subseteq \bar{\mathcal{S}} \subseteq \ldots \subseteq \mathcal{S}_5 \subseteq \mathcal{S}_3 \subseteq \mathcal{S}_1$ *and* $r_0 \leq r_2 \leq r_4 \leq \ldots \leq \bar{r} \leq \ldots \leq r_5 \leq r_3 \leq r_1$.

The intuition for this comparative statics comes from the mechanism described above. By Proposition 5, minor variations in the price suffice to clear the market when k is even, while larger ones are needed when k is odd.

Unfortunately, the difficulties encountered for the comparative statics of the means of the equilibrium distributions affect the comparative statics of their variances. For example, in trying to compare $\bar{v}$ and v_0, we need to determine whether

$$\frac{4}{1-4\delta^2} - \frac{1}{4\delta^2}\ln^2\left(\frac{1+2\delta}{1-2\delta}\right) \gtreqless \frac{4(1+\rho)^2}{[1+(1+2\delta)\rho][1+(1-2\delta)\rho]} - \frac{(1+\rho)^2}{4\delta^2\rho^2}\ln^2\left[\frac{1+(1+2\delta)\rho}{1+(1-2\delta)\rho}\right].$$

Since $\rho > 1$,

$$\frac{4}{1-4\delta^2} > \frac{4(1+\rho)^2}{[1+(1+2\delta)\rho][1+(1-2\delta)\rho]}$$

and

$$\frac{1}{4\delta^2} < \frac{(1+\rho)^2}{4\delta^2\rho^2},$$

but the comparison becomes analytically intractable since

$$\frac{1+2\delta}{1-2\delta} > \frac{1+(1+2\delta)\rho}{1+(1-2\delta)\rho}.$$

Computing the equilibria explicitly, however, suggests that the comparative statics of the variance resemble the ones of the range, as per Table 3. We consider three values of $\rho > 1$ and, for each of them, two values of $\delta < (1-\rho^{-2})/2$. We do not display the median of the distribution, as it is always 1, and confirm the analytical results regarding the support and the range. Interestingly, the results yield

---

[14] Carvajal and Zhou (2025) offer a more detailed discussion on the convergence result. In particular, that paper shows that the convergence result holds for any initial belief E(Π).

Table 3
Moments of the distribution of equilibrium prices.

$\rho = 1.5$									
$\delta = 0.025$					$\delta = 0.25$				
k	Support	Mean	Range	Variance	k	Support	Mean	Range	Variance
0	[0.942, 1.062]	1.0006	0.1201	0.0012	0	[0.538, 1.857]	1.0635	1.3187	0.1377
1	[0.835, 1.198]	1.0054	0.3629	0.0109	1	[0.053, 19.000]	2.2716	18.9474	10.3493
2	[0.926, 1.080]	1.0010	0.1545	0.0020	2	[0.443, 2.256]	1.1091	1.8125	0.2508
3	[0.883, 1.133]	1.0026	0.2502	0.0052	3	[0.232, 4.306]	1.3433	4.0739	1.0474
10	[0.906, 1.104]	1.0016	0.1982	0.0033	10	[0.338, 2.955]	1.1919	2.6164	0.4887
100	[0.905, 1.105]	1.0017	0.2005	0.0033	100	[0.333, 3.000]	1.1972	2.6667	0.5055
REE	[0.905, 1.105]	1.0017	0.2005	0.0033	REE	[0.333, 3.000]	1.1972	2.6667	0.5055
$\rho = 5$									
$\delta = 0.045$					$\delta = 0.45$				
k	Support	Mean	Range	Variance	k	Support	Mean	Range	Variance
0	[0.860, 1.162]	1.0038	0.3017	0.0076	0	[0.143, 7.000]	1.5945	6.8571	2.4112
1	[0.829, 1.207]	1.0059	0.3783	0.0119	1	[0.032, 31.000]	2.6629	30.9677	19.6153
2	[0.836, 1.196]	1.0053	0.3600	0.0108	2	[0.057, 17.667]	2.2163	17.6101	9.3788
3	[0.835, 1.198]	1.0054	0.3635	0.0110	3	[0.052, 19.293]	2.2833	19.2408	10.5643
10	[0.835, 1.198]	1.0054	0.3629	0.0109	10	[0.053, 19.000]	2.2716	18.9474	10.3493
100	[0.835, 1.198]	1.0054	0.3629	0.0109	100	[0.053, 19.000]	2.2716	18.9474	10.3493
REE	[0.835, 1.198]	1.0054	0.3629	0.0109	REE	[0.053, 19.000]	2.2716	18.9474	10.3493
$\rho = 15$									
$\delta = 0.0495$					$\delta = 0.495$				
k	Support	Mean	Range	Variance	k	Support	Mean	Range	Variance
0	[0.830, 1.205]	1.0058	0.3745	0.0116	0	[0.037, 26.826]	2.5441	26.7888	16.3027
1	[0.819, 1.221]	1.0066	0.4017	0.0134	1	[0.003, 357.400]	4.9118	357.3972	324.4529
2	[0.820, 1.220]	1.0066	0.3998	0.0133	2	[0.005, 193.302]	4.3190	193.2970	167.0155
3	[0.820, 1.220]	1.0066	0.3999	0.0133	3	[0.005, 199.392]	4.3487	199.3869	172.7888
10	[0.820, 1.220]	1.0066	0.3999	0.0133	10	[0.005, 199.000]	4.3468	198.9950	172.4170
100	[0.820, 1.220]	1.0066	0.3999	0.0133	100	[0.005, 199.000]	4.3468	198.9950	172.4170
REE	[0.820, 1.220]	1.0066	0.3999	0.0133	REE	[0.005, 199.000]	4.3468	198.9950	172.4170

$$m_0 \le m_2 \le m_4 \le \ldots \le \bar{m} \le \ldots \le m_5 \le m_3 \le m_1$$

and

$$v_0 \le v_2 \le v_4 \le \ldots \le \bar{v} \le \ldots \le v_5 \le v_3 \le v_1 .$$

We conjecture that these comparative statics hold in general.

7. Learning and trade volume

Since the speculators have an imperfect understanding of the market, there is trade at equilibrium. Given π, the volume of trade is

$$\mu^F |y^F(q_k(\pi), \pi)| + \mu^s |y^S_k(q_k(\pi))|.$$

By market clearing, this number equals

$$2\mu^F |y^F(q_k(\pi), \pi)| = 2\mu^F \omega^Y \left| \pi \frac{q_k(\pi)+1}{q_k(\pi)} - 1 \right| .$$

To study the behavior of this variable, we dismiss the constant and focus on the random variable

$$T_k := \left| \Pi \frac{Q_k+1}{Q_k} - 1 \right| .$$

Its mean is

$$\mathrm{E}(T_k) = \frac{(-1)^k + \rho^{k+1}}{4\delta\rho^{2k+2}} \cdot \ln \left\{ \frac{[(-1)^k + \rho^{k+1}]^2}{[(-1)^k + \rho^{k+1}]^2 - 4\delta^2\rho^{2k+2}} \right\} .$$

Table 4 shows these moments for the same parametric configurations as Table 3. With the same reasoning as before, wide variations in the price are necessary in rounds where the reasoning of the speculators dampens trade and vice versa. This yields, in general, that

Table 4
Volume of trade: $\mathrm{E}(T_k)$.

$\rho = 1.5$		
k	$\delta = 0.025$	$\delta = 0.25$
0	0.40018	0.41916
1	0.80326	1.64023
2	0.22874	0.24747
3	0.24663	0.31157
10	0.01144	0.01310
20	0.00020	0.00023
100	0.00000	0.00000
REE	0	0
$\rho = 5$		
k	$\delta = 0.045$	$\delta = 0.45$
0	0.16714	0.24494
1	0.04185	0.10009
2	0.00797	0.01588
3	0.00161	0.00330
10	0.00000	0.00000
REE	0	0
$\rho = 15$		
k	$\delta = 0.0495$	$\delta = 0.495$
0	0.06277	0.14339
1	0.00449	0.02031
2	0.00030	0.00118
3	0.00002	0.00008
4	0.00000	0.00001
5	0.00000	0.00000
REE	0	0

$$\mathrm{E}(T_0) > \mathrm{E}(T_2) > \mathrm{E}(T_4) > \ldots \downarrow 0$$

and

$$\mathrm{E}(T_1) > \mathrm{E}(T_3) > \mathrm{E}(T_5) > \ldots \downarrow 0.$$

The same effects apply to the median of the distribution, broadly speaking.[15]

8. Noise trade

In our setting so far, the only random variable that affects prices is a fundamental of the economy. Models of financial markets often incorporate further randomness to obtain prices that fail to reveal the fundamentals fully. We now incorporate this sort of "noise" trade to study its effects on the convergence, or lack thereof, of our learning process. For this exercise, we ignore the wealth dynamics of Section 9.

For simplicity of notation only, assume that $\mu^S w^S = 1$, so that $\rho = \mu^F w^F$. Suppose that, regardless of π and q, there is a demand for Z units of the asset that pays in state 1 (and of $-qZ$ for the one that pays in state 2). This extra demand is noise trade: a random variable with support $[-\zeta, \zeta]$ and expectation $\mathrm{E}(Z) = 0$. For computational reasons, $\zeta < (1/2 - \delta)\rho$.

The market-clearing condition thus becomes $\mu^F y_1^F + \mu^S y_1^S + z = 0$, for all realizations z of Z.

8.1. Rational expectations equilibrium

The optimal asset demands of the two types of traders continue to be given by Eqs. (2) and (3).[16] Adding noise trade to the results in Section 3.2, the REE is a function $\bar{q}(\pi, z)$, defined over all realizations π and z of Π and Z, such that markets clear when the fundamentalists use beliefs $\mathrm{E}(\Pi \mid \bar{Q} = q)$ in Eq. (3), where $\bar{Q}$ is the random variable $\bar{q}(\Pi, Z)$.

---

[15] The formula for the median is

$$\mathrm{med}(T_k) = \frac{2\sqrt{1 + 2(-1)^k + (1 + 4\delta^2)\rho^{2k+2}} - 2 + (-\rho)^{k+1}}{4\delta\rho^{2k+2}}.$$

[16] We maintain our assumption that the speculators are expected utility maximizers. We'll further comment on this assumption below.

By direct computation, the REE is

$$\bar{q}(\pi, z) = \frac{\pi\rho}{(1-\pi)\rho - z}$$

and

$$\mathrm{E}\,(\Pi \mid \bar{Q} = q) = \frac{q}{1+q}.$$

Obviously, the REE is no longer fully revealing. Upon observing q, the speculators infer the probability $\mathrm{E}\,[\Pi \mid \bar{Q} = q]$, and, interestingly, trade null positions. This means that all the trade taking place at the REE prices is between the fundamentalists, who distinguish the fundamental information from the noise, and the noise traders: by direct substitution, $y^F(\pi, \bar{q}(\pi, z)) = -z/\mu^F$ for all π and z.

It is important to note that an implicit assumption of our setting underlies this result. With prices that are not fully revealing, upon observing q, the speculators cannot rule out any realization of Π in the set

$$\{\pi \in \Delta \mid \exists z \in [-\zeta, \zeta] : \bar{Q}(\pi, z) = q\}.$$

Since $(\pi, 1-\pi)$ is a probability distribution, their optimization problem is thus subject to ambiguity. In maintaining the assumption that the speculators are expected utility maximizers and continuing to use Eq. (3), we have assumed that these individuals are ambiguity-neutral. While the extension of these results to ambiguity-averse speculators is interesting, it is beyond the scope of this paper.

8.2. Learning

The learning process and equilibrium prices of Section 3.4 are affected similarly to the REE. We maintain the assumption that completely inexperienced speculators fail to realize that asset prices incorporate information, and use the unconditional expectation $\mathrm{E}\,(\Pi) = 1/2$ regardless of q. The market clearing condition then defines the price function $q_0(\pi, z)$ and the random variable $Q_0(\Pi, Z)$, which level-1 speculators use to infer $\mathrm{E}\,(\Pi \mid Q_0 = q)$ when choosing their optimal portfolio upon observing q, which gives the price function $q_1(\pi, z)$, and so on.

As before, the sequence of level-k prices can be obtained by induction:

$$q_k(\pi, z) = \frac{(-1)^k + 2\pi\rho^{k+1}}{(-1)^k + 2(1-\pi)\rho^{k+1} - z}.$$

Their beliefs, thus, evolve according to

$$\mathrm{E}\,(\Pi \mid Q_k = q) = \frac{q}{1+q} + \frac{(-1)^k}{2\rho^{k+1}} \cdot \frac{q-1}{q+1},$$

where Q_k is the random variable $q_k(\Pi, Z)$.

8.3. (Non-)convergence to REE

It is immediate from above that, since $\rho > 1$,

$$Q_k \xrightarrow[\text{s.}]{} \frac{\Pi}{1-\Pi} = \bar{q}(\Pi, 0). \tag{10}$$

while

$$\mathrm{E}\,(\Pi \mid Q_k = q) \to \frac{q}{1+q} = \mathrm{E}\,(\Pi \mid \bar{Q} = q). \tag{11}$$

The second of these results, Eq. (11) says that the sequence of the fundamentalists' beliefs converges pointwise to the REE beliefs. Eq. (10), however, says that the sequence of level-k equilibrium prices converges *surely* to a random variable that differs from the REE prices. The limit of the sequence $\langle Q_k \rangle_{k=1}^{\infty}$ is the random variable that would model the REE *in the absence of noise trade*.

This is important because it implies a different trade pattern even at the limit. We observed above that in the REE, the only trade that takes place is between the fundamentalists and the noise traders. If we compute the asset holdings of the fundamentalists at level-k prices, on the other hand,

$$y^F(\pi, q_k(\pi, z)) = -\frac{(-1)^k(1-2\pi) + z}{(-1)^k + 2\pi\rho^{k+1}} \cdot \frac{\rho}{\mu^F}.$$

Thus, $y^F(\pi, q_k(\pi, z)) \to 0$, for all π and z, as $k \to \infty$. While the speculators' expectations converge pointwise to the same function as in the REE, at the limit it is them, and not the fundamentalists, who carry out all the trade with noise traders.

9. Wealth evolution

Convergence of the sequence of level-k prices to REE, as per Proposition 3, relies critically on our assumption that $\rho > 1$: if this is not the case, Eq. (9) fails to converge to zero.

This observation motivates an exploration of how the wealth of the two types of agents in our economy evolves as they trade. While our setting is not genuinely dynamic, we approach this question using the average payoffs of the equilibrium portfolios to update the agents' endowments as follows. For the sake of simplicity, we assume that $\mu^F = \mu^S$ and there is no noise trade.

9.1. Portfolio payoffs

When their level of understanding is k and the probability is π, each speculator's portfolio earns, in expectation, a payoff of

$$p_k(\pi) := \pi \cdot y_k^S(q_k(\pi)) + (1-\pi) \cdot [-q_k(\pi) y_k^S(q_k(\pi))] = [\pi - (1-\pi) q_k(\pi)] \cdot y_k^S(q_k(\pi)),$$

where we are using their budget constraint to write the holdings of the asset for state $s = 2$ in terms of her holdings of the asset for state $s = 1$. Once we consider the probability distribution's randomness over the two states, the expected payoff is the random variable $P_k = p_k(\Pi)$.

Denoting by ω_k^a the agents' wealth when they have level-k understanding, define the random variables $W_{k+1}^S = \omega_k^S + P_k$ and $W_{k+1}^F = \omega_k^F - P_k$. These variables update the two types of endowments for the next level of understanding. The ratio of the two aggregate endowments becomes the random variable

$$R_{k+1} := \frac{W_{k+1}^F}{W_{k+1}^S} = \frac{\rho_k - \frac{P_k}{\omega_k^S}}{1 + \frac{P_k}{\omega_k^S}}, \tag{12}$$

where $\rho_k = \omega_k^F / \omega_k^S$.

To study the dynamics of wealth, we use this construction to define $\rho_{k+1} := \mathrm{E}(R_{k+1})$, recursively, starting from $\rho_0 = \omega_0^F / \omega_0^S$. Using the same learning process as before, the sequence of pricing functions becomes

$$q_k(\pi) = \frac{(-1)^k + 2\pi \prod_{\ell=0}^k \rho_\ell}{(-1)^k + 2(1-\pi) \prod_{\ell=0}^k \rho_\ell}. \tag{13}$$

For the convergence analysis to remain valid under these dynamics, we require that $\rho_k > 1$ at all k and $\lim_{k\to\infty} \prod_{\ell=0}^k \rho_\ell = \infty$. The following result shows that this is indeed the case:

Proposition 7. *Given $\rho_0 > 1$, the sequence $\langle \rho_k \rangle_{k=0}^\infty$ is increasing.*

Proof. Using Eqs. (3) and (13), by direct computation

$$\frac{P_k}{\omega_k^S} = -\frac{(1-2\Pi)^2}{1 + 2(-1)^k \prod_{\ell=0}^k \rho_\ell + 4\Pi(1-\Pi) \prod_{\ell=0}^k \rho_\ell} \rho_k.$$

Substitution into Eq. (7) then yields

$$\frac{R_{k+1}}{\rho_k} = \frac{1 + 2(-1)^k \prod_{\ell=0}^k \rho_\ell + 4\Pi(1-\Pi) \prod_{\ell=0}^k \rho_\ell^2 + (1-2\Pi)^2}{1 + 2(-1)^k \prod_{\ell=0}^k \rho_\ell + 4\Pi(1-\Pi) \prod_{\ell=0}^k \rho_\ell^2 - (1-2\Pi)^2 \rho_k}. \tag{14}$$

When k is even, it is immediate that $R_{k+1}/\rho_k > 1$ with probability 1. To see that the same is true when k is odd, note that the sign of the term

$$1 - 2\prod_{\ell=0}^k \rho_\ell + 4\pi(1-\pi) \prod_{\ell=0}^k \rho_\ell^2$$

depends on the value of π. By direct computation, though, this value is positive as long as $\prod_{\ell=0}^k \rho_\ell > 1$.

The argument that $\langle \rho_k \rangle_0^\infty$ is increasing now proceeds by induction: (1) Given $\rho_0 > 1$,

$$\frac{R_1}{\rho_0} = \frac{1 - 2\rho_0 + 4\Pi(1-\Pi)\rho_0 + (1-2\Pi)^2}{1 - 2\rho_0 + 4\Pi(1-\Pi)\rho_0 - (1-2\Pi)^2 \rho_0} > 1$$

and, hence, $\rho_1 = \mathrm{E}(R_1) > \rho_0$. (2) For any level k, given $\prod_{\ell=0}^k \rho_\ell > 1$, the expression on the right-hand side of Eq. (14) is larger than one and, as before, $\rho_{k+1} = \mathrm{E}(R_{k+1}) > \rho_k$. ☐

The learning process, thus, is costly to the speculators even though they are acting optimally *with the information they have*.

9.2. Exogenous wealth growth

Suppose that instead of accruing profits and losses, the relative wealth of the two types changes exogenously as k evolves. Suppose that the growth rates are α for the fundamentalists and β for the speculators: $\omega^F_{k+1} = (1+\alpha)\omega^F_k$, while $\omega^S_{k+1} = (1+\beta)\omega^S_k$. Then, we can replace Eq. (12) by

$$\rho_{k+1} = \frac{\omega^F_{k+1}}{\omega^S_{k+1}} = \rho_k \frac{1+\alpha}{1+\beta} = \rho_0 \left(\frac{1+\alpha}{1+\beta}\right)^k$$

and continue to use Eq. (13) to assess the evolution of prices, as long as other assumptions made before continue to hold.

It is immediate our analysis requires that $\beta \leq \alpha$.[17] Under this condition, it is immediate that $\rho_k > 1$ for all k and $\lim_{k\to\infty} \prod_{\ell=0}^k \rho_\ell = \infty$, so the sequence of equilibrium prices converges to REE.

9.3. Profits and exogenous growth

Finally, let us assume that the same exogenous growth we just considered takes place while portfolio profits are accrued. The corresponding random variables are $W^S_{k+1} = \omega^S_k + P_k + \beta\omega^S_0$ and $W^F_{k+1} = \omega^F_k - P_k + \alpha\omega^F_0$. The equation that replaces Eq. (12) is, therefore,

$$R_{k+1} = \frac{\rho_k - \frac{P_k}{\omega^S_k} + \alpha\frac{\omega^F_0}{\omega^S_k}}{1 + \frac{P_k}{\omega^S_k} + \beta\frac{\omega^S_0}{\omega^S_k}}.$$

If $\alpha \geq \beta$, so that the fundamentalists' wealth exogenously grows faster, the two mechanisms of the previous subsections operate in the same direction: $\langle\rho_k\rangle_{k=1}^\infty$ is bounded below strictly above 1, $\lim_{k\to\infty} \prod_{\ell=0}^k \rho_\ell = \infty$, and the sequence of equilibrium prices converges to REE.

If, on the other hand, $\alpha < \beta$, convergence again fails. This is because P_k is bounded below while $(\beta/\alpha)^k \to \infty$ so $R_{k+1} < 1$, with probability 1, for k large enough.

10. Concluding remarks

Our paper explores how adaptive learning by bounded rational agents shapes convergence to rational expectations equilibrium (REE) in a Radner economy. By modeling a two-state exchange economy with asymmetric information and logarithmic preferences, we demonstrate that convergence to REE occurs only if uninformed speculators' wealth remains below half of aggregate wealth. This process, however, is non-monotonic: markets oscillate between phases of excess price volatility and heightened trade volume. These oscillations arise from the speculators' iterative belief updates, where past errors distort inferences, creating feedback between price responsiveness and informational efficiency.[18]

The findings underscore the fragility of REE. Stability hinges on wealth distribution: markets self-correct only when fundamentalists dominate. Exogenous growth in speculators' wealth or noise trade sometimes disrupts convergence, revealing how structural parameters govern equilibrium dynamics. Critically, learning imposes costs on uninformed agents: their suboptimal trades erode wealth and deepen the informational asymmetries.

By integrating level-k reasoning with adaptive learning in a general equilibrium model, our paper bridges strategic and macroeconomic frameworks, offering micro-foundations for how agents interpret price signals. The oscillatory nature of convergence challenges assumptions of monotonic adjustment in macroeconomic models, highlighting the volatility rooted in learning frictions. Ultimately, our work emphasizes that convergence to rational expectations emerges not inherently, but through disciplined learning under precise economic constraints.

References

Angeletos, G.-M., Lian, C., 2017. Dampening general equilibrium: from micro to macro. Working paper.

Bastianello, F., Fontanier, P., 2025. Expectations and learning from prices. Rev. Econ. Stud. 92 (3), 1341–1374.

Ben-Porath, E., Heifetz, A., 2011. Common knowledge of rationality and market clearing in economies with asymmetric information. J. Econ. Theory 146 (6), 2608–2626.

Blume, L., Bray, M., Easley, D., 1982. Introduction to the stability of rational expectations equilibrium. J. Econ. Theory 26 (2), 313–317.

Blume, L., Easley, D., 1984. Rational expectations equilibrium: an alternative approach. J. Econ. Theory 34 (1), 116–129.

Bray, M., 1982. Learning, estimation, and the stability of rational expectations. J. Econ. Theory 26 (2), 318–339.

---

[17] If this is not the case, $\rho_k \to 0$ and Condition (1), *mutatis mutandis*, becomes untenable.

[18] To be sure, note that our results do not depend on the specific learning process modeled in Section 3.4. Suppose that the speculators can observe π after trade has occurred. (For example, if they can access whatever private information the fundamentalists observed *before* trade.) In this case, just one observation suffices to identify the q_k function almost surely, if one assumes that the speculators know the functional form of that function. All the convergence results in this paper remain valid, although the interpretation of the wealth evolution exercises becomes rather unclear.

Carvajal, A., Zhou, H., 2025. On the sophistication of financial investors and the information revealed by prices. J. Dyn. Games 12 (4), 406–414.

Crawford, V., Costa-Gomes, M., Iriberri, N., 2013. Structural models of nonequilibrium strategic thinking: theory, evidence, and applications. J. Econ. Lit. 51 (1), 5–62.

DeCanio, S., 1979. Rational expectations and learning from experience. Q. J. Econ. 93 (1), 47–57.

Dubey, P., Geanakoplos, J., Shubik, M., 1987. The revelation of information in strategic market games: a critique of rational expectations equilibrium. J. Math. Econ. 16 (2), 105–137.

Dutta, J., Morris, S., 1997. The revelation of information and self-fulfilling beliefs. J. Econ. Theory 73 (1), 231–244.

Eyster, E., Rabin, M., Vayanos, D., 2019. Financial markets where traders neglect the informational content of prices. J. Finance 74 (1), 371–399.

Farhi, E., Werning, I., 2019. Monetary policy, bounded rationality, and incomplete markets. Am. Econ. Rev. 109 (11), 3887–3928.

Feldman, M., 1987. An example of convergence to rational expectations with heterogeneous beliefs. Int. Econ. Rev. 28 (3), 635–650.

Grandmont, J.-M., 1998. Expectations formation and stability of large socioeconomic systems. Econometrica 66 (4), 741–781.

McAllister, P.H., 1990. Rational behavior and rational expectations. J. Econ. Theory 52 (2), 332–363.

Mondria, J., Vives, X., Yang, L., 2022. Costly interpretation of asset prices. Manag. Sci. 68 (1), 52–74.

Muth, J., 1961. Rational expectations and the theory of price movements. Econometrica 29 (3), 315–335.

Nagel, R., 1995. Unraveling in guessing games: an experimental study. Am. Econ. Rev. 85 (5), 1313–1326.

Pesce, M., Urbinati, N., Yannelis, N., 2024. On the limit points of an infinitely repeated rational expectations equilibrium. Econ. Theory, 1–28.

Radner, R., 1979. Rational expectations equilibrium: generic existence and the information revealed by prices. Econometrica 47 (3), 655–678.

Shiller, R., 1978. Rational expectations and the dynamic structure of macroeconomic models: a critical review. J. Monet. Econ. 4 (1), 1–44.

Stahl, D., Wilson, P., 1994. Experimental evidence on players' models of other players. J. Econ. Behav. Organ. 25 (3), 309–327.

Stahl, D., Wilson, P., 1995. On the players' model of other players: theory and experimental evidence. Games Econ. Behav. 10 (1), 218–254.

Zhou, H., 2022. Informed speculation with k-level reasoning. J. Econ. Theory 200, 105384.