Fernando De La Prieta ·
Alfonso González-Briones ·
Pawel Pawleski · Davide Calvaresi ·
Elena Del Val · Fernando Lopes ·
Vicente Julian · Eneko Osaba ·
Ramón Sánchez-Iborra (Eds.)

Communications in Computer and Information Science 1047

Highlights of Practical Applications of Survivable Agents and Multi-Agent Systems

The PAAMS Collection

International Workshops of PAAMS 2019
Ávila, Spain, June 26–28, 2019
Proceedings

Springer

Asymmetric Information and Learning by Imitation in Agent-Based Financial Markets

Luca Gerotto[1], Paolo Pellizzari[1(✉)], and Marco Tolotti[2]

[1] Department of Economics, Ca' Foscari University, Cannaregio 873, 30121 Venice, Italy
{luca.gerotto,paolop}@unive.it
[2] Department of Management, Ca' Foscari University, Cannaregio 873, 30121 Venice, Italy
tolotti@unive.it

Abstract. We describe an agent-based model of a market where traders exchange a risky asset whose returns can be partly predicted purchasing a costly signal. The decision to be informed (at a cost) or uninformed is taken by means of a simple learning by imitation mechanism that periodically occurs.

The equilibrium is characterized describing the stationary distribution of the price and the fraction of the informed traders. We find that the number of agents who acquire the signal decreases with its cost and with agents' risk aversion and, conversely, it increases with the signal-to-noise ratio and when learning is slow, as opposed to frequent. Moreover, price volatility appears to directly depend on the fraction of informed traders and, hence, some heteroskedasticity is observed when this fraction fluctuates.

Keywords: Agent-based modeling · Bounded rationality · Information in financial markets

1 Introduction

Information is of paramount importance in competitive financial markets and one of the most sought-after properties of a market is its ability to spread information in a timely manner. However, some of the paradoxical implications of information are well known and have been widely analyzed in the last decades (see, for example, Schredelseker [1], Kurlat and Veldkamp [2], Veldkamp [3], and the evergreen Grossman and Stiglitz [4], hereafter referred to as GS). One of the celebrated achievements in GS is to show that, in exchange economies, a trading equilibrium among informed traders cannot exist, as a perfectly informative price system would not compensate arbitrageurs for their (costly) activity of information gathering and processing.

F. De la Prieta et al. (Eds.): PAAMS 2019 Workshops, CCIS 1047, pp. 164–175, 2019.
https://doi.org/10.1007/978-3-030-24299-2_14

In the standard treatment, fully rational agents *ex-ante* solve for the equilibrium, trading (only) at the equilibrium price given the fraction λ of informed traders and computing the unique λ that makes the expected utility of any informed agent equal to that of the uninformed one. This is a grueling task involving sophisticated cognitive and technical abilities that may belong to an abstract *Homo Economicus* but are likely to be scant in more realistic depictions of human behavior.

In this paper, we present an agent-based model of a dynamic market where the risky return depends on a random component as well as on an informative signal, which can be purchased for a constant cost. Informed agents can exploit the reduction in uncertainty provided by the signal and increase their profit by making educated bids on the future payoff. Uninformed agents do not bear any information cost and face greater risks but, intuitively, there is an equilibrium share of informed traders in which the benefits of purchasing the signal exactly offsets the cost, thus making expected profits equal for informed and uninformed traders.

In our model, agents are boundedly rational and noisily attempt to maximize their cumulated wealth over some span of time, deciding whether to acquire the signal (paying the required cost) or not. We do not assume the existence of an utility function, nor the ability to solve sophisticated maximization schemes or understand the complex endogenous structure of the stochastic equilibrium that should materialize (prices depends on the fraction of informed traders which, in turn, depends on the profit that are driven by the individual decisions ultimately responsible for the price dynamics).

Instead, we assume that traders evolve using a simple learning by imitation device: after a predetermined number of trading periods, some agents are randomly paired, compare their performance (namely, cumulated wealth) and the poorer trader ends up in copying the strategy of the richer; i.e., if the pair is formed by an informed and an uninformed agent, after learning, both will adopt the same (more favorable) strategy.[1]

We obtain three main results. First, for all the values of the parameters, the model converges to some equilibrium expressed in terms of the fraction of informed and the price of the risky asset. Second, the outcome is affected, as expected, by the informativeness of the signal, but also by a set of "learning" parameters of the model, and there are situations in which heteroskedasticity of prices is observed. Finally, the length of the period used to cumulate profits has a remarkable role and, say, short-term traders prefer not to use weakly informative signals whereas the same knowledge is purchased and exploited if more periods are allowed to average profits.

1.1 Related Literature

We depart from the seminal GS setup in that the "game" is no longer static but dynamically develops in a series of periods through an explicit learning

[1] To preserve diversity in the population, we add also a minimal degree of random "mutation" in every period, see the details below.

mechanism. Agents learn in a very simple and robust way, by checking whether another trader makes larger profits. By contrast, Routledge [5,6] uses full-blown genetic algorithms to equip agents with sophisticated learning skills. We enrich this extant literature showing that market outcomes are affected by learning in important ways and this holds even when learning takes place in extremely simple ways (or, if you like, also when the "genetic algorithm" has no proper selection or mutation operator and crossover is replaced by sheer imitation).

Perhaps more importantly, the agents of the model do not play a sequence of repeated but otherwise identical trading games, as done previously. We introduce a market maker in charge of adjusting the future price based on the excess demand produced by the traders in the current time. Hence, in a somewhat realistic fashion, price fluctuates because of random shocks, the changing fraction of informed/uninformed traders, as well as due to price adjustments stickily incorporating imbalances in demand for the risky asset.

Our paper can be related to the vast amount of work dealing with asymmetric information in financial markets. Several scholars have explored different setups and definition of *information*: Chen et al. [7], for instance, develop a model where returns are affected by the volume of Google searches of the asset ticker (the "driving force") and propose a simple three-bodies approximation of the resulting price dynamics. Billett et al. [8] examine how the reduction in analysts' coverage of one stock predict worse industry-adjusted performance, due to the reduction of the information available to investors.

Recently, Krichene and El-Aroui [9] presents a model in which, among other features, "information" is assumed to be essentially equivalent to traders' sentiment, which can spread and has the potential to trigger herds, bubbles and crashes. We feel that the previous and non-exhaustive list of models appears to depict *knowledge* (on a specific stock), more than *information* on the returns, that is expressed through economic analysis, extent of web coverage and perceptions/sentiment of investors. In our model, the signal directly refers to the future yield and may be thought as a simplified (or, if you wish, distilled) form of the just mentioned sources of knowledge.

The article is organized as follows. Next section describes the model and provides details both on the structure of the market and on the features of the agents. Section 3 presents the results of a NetLogo implementation of the model and relates the findings to the literature. We then have some concluding remarks and point out paths for future work.

2 The Model

Our model basically consists of two parts: a simple financial market where two assets are available and a staggered learning mechanism involving, at selected calendar dates, a random set of traders.

The Financial Market

We assume there are two assets in the market: a riskless asset with unit price and a risky asset, whose price p_t is determined by a market maker. While the safe asset yields as return the risk-free rate $R \geq 1$, the risky asset has an uncertain return u_t that can be decomposed in three components:

$$u_t = d + \theta_t + \epsilon_t, \tag{1}$$

where $d \geq 1$ is a constant; $\theta_t \sim N(0, \sigma_\theta^2)$ is a signal observable at the beginning of the period by paying an amount c; and $\epsilon_t \sim N(0, \sigma_\epsilon^2)$ is an ex-ante unforecastable zero-mean shock.

The market is populated by N traders who can be either informed or uninformed. At the beginning of each period t, a fraction λ_t of traders spends c to be *informed* and learn θ_t, whereas the remaining traders are *uninformed*.

The information set Ω_t^j of agent j at time t can be either θ_t or the empty set, depending on whether he is informed. Accordingly, his demand for the risky asset is determined by the following heuristic:[2]

$$X_t(\Omega_t^j) = \frac{E(u_t|\Omega_t^j) - p_t R}{\alpha Var(u_t|\Omega_t^j)} = \begin{cases} \frac{d+\theta_t - p_t R}{\alpha\sigma_\epsilon^2} = X_t^I \text{ if } \Omega_t^j = \theta_t \\ \frac{d - p_t R}{\alpha(\sigma_\theta^2 + \sigma_\epsilon^2)} = X_t^U \text{ if } \Omega_t^j = \emptyset \end{cases} . \tag{2}$$

Intuitively, the higher the expected excess return $E(u_t|\Omega_t^j)$ with respect to the riskless asset, the higher the quantity demanded. Hence, the demand increases in d and, for informed traders, in θ_t, whereas it decreases in the price p_t and R, *ceteris paribus*. Under the assumption that agents are risk averse, their demand is negatively related to their degree of risk aversion α and to the perceived volatility of returns $Var(u_t|\Omega_t^j)$. The informed agents know the signal θ_t concerning u_t, which leads them to sell the risky asset when $d+\theta_t-p_tR$ is negative and buy the risky asset otherwise. They do so facing a residual risk that depends on the shock ϵ_t alone (being θ_t known). In contrast, the uninformed trader takes a short (long) position in the risky asset if $d - p_t R$ is negative (positive). The demand of the uninformed is typically lower than the one of the informed as the lack of knowledge of θ_t inflates the denominator to $\alpha(\sigma_\theta^2 + \sigma_\epsilon^2)$. As a consequence, the trading volume of the informed is most often much bigger than the one produced by the uninformed agents. Observe that all informed agents demand the identical amount X_t^I. The same holds for uninformed quantity X_t^U that is the same for any uninformed trader.

At the end of the period, the wealths of the informed and uninformed agents are, respectively:

$$w_t^I = (u_t - p_t R) X_t^I + (w_{t-1}^I - c) R \tag{3}$$

$$= \left(d + \theta_t + \epsilon_t\right)\left(\frac{d + \theta_t - p_t R}{\alpha\sigma_\epsilon^2}\right) + (w_{t-1}^I - c)R, \tag{4}$$

[2] This result corresponds to a simplified version of equations (8) and (8′) of Grossman and Stiglitz [4]. Note, in particular, that our agents are not able to exploit entirely the information revealed by the price p_t on the signal θ_t (hence, on u_t) as for the fully rational agents in GS.

and

$$w_t^U = (u_t - p_t R) X_t^U + w_{t-1}^U R \tag{5}$$

$$= \left(d + \theta_t + \epsilon_t\right)\left(\frac{d - p_t R}{\alpha(\sigma_\theta^2 + \sigma_\epsilon^2)}\right) + w_{t-1}^U R. \tag{6}$$

Once trading has occurred, the market maker reacts to any excess demand (supply) of the risky asset by proportionally increasing (decreasing) the price that will be available in the next trading period. As customarily done (see Cont and Bouchaud [10]), this is a simple device to introduce a reasonable price dynamics in the model, by means of adjustments of the current price based on the magnitude and sign of the current demand imbalance.

Given the average per-trader total excess demand $Q_t = \lambda_t X_t^I + (1 - \lambda_t) X_t^U$ at the end of the period, the market maker determines next price

$$p_{t+1} = p_t + k(Q_t - \eta_t), \tag{7}$$

where $\eta_t \sim N(0, \sigma_\eta^2)$ is an exogenous supply shock. The value of the parameter $k > 0$ determines the strength of the reaction of the market maker.

The Learning Mechanism

Staggered learning in the model is introduced assuming that from time to time agents assess their performance comparing their wealth to the one of other peers. We have two parameters in the learning mechanism: T is the horizon over which profits are cumulated before a learning round by imitation begins, while h is the number of couples of agents who compare the respective performances and eventually copy the more profitable strategy. More formally, every T periods, $2h$ out of the N agents are randomly paired. Denote by $\mathcal{T} = \{zT,\ z \in \mathbb{N}\}$, the set of calendar dates at which random matchings happen. For large N, at dates $\tau \in \mathcal{T}$, approximately $h\lambda_\tau(1 - \lambda_\tau)$ out of the h pairs are composed by an informed and an uninformed trader.[3] They compare their performance (i.e., cumulated wealth) and the poorer trader ends up in copying the strategy of the richer. As a result, after learning, they both will be either in the set of the informed or uninformed agents for the next T periods.

When some traders modify their strategy at time $\tau \in \mathcal{T}$, a change in the proportion of informed at the aggregate level follows: $\lambda_{\tau+1}$ will be higher than λ_τ if the informed outperformed the uninformed over the previous T periods, and will be lower otherwise. To avoid trivial situations in which all agents are either informed or uninformed, thus making the learning procedure useless, we conclude the process assigning a random status (Informed or Uninformed) to one agent.

Finally, in order to study the effects of the parameter T (the duration of the accumulation period), after learning has taken place all agents start from the

[3] Clearly, if the pair is formed by two agents that were equally informed or uninformed in the last T periods, no change happens as both members in the couple have the same wealth.

same wealth level w_0, i.e. all the previous gains and losses are reset (or wealth is entirely consumed, if the reader prefers an equivalent interpretation).

It is worth pointing out that the previous learning mechanism is, basically, a crude learning by imitation procedure requiring very little sophistication on the part of the agents, who are only assumed to be able to know, every once in a while, the wealth of a peer and whether he has been paying to obtain information in the (recent) past. We believe this is a cognitive plausible representation of agents and does not require the precise understanding of the endogenous equilibrium possibly arising or, say, the skills needed to maximize a utility function. Even if there are technical and conceptual similarities with genetic algorithms (whose variety is, incidentally, sweeping), our setup is greatly simplified as selection is totally random and independent of fitness and the crossover operator is replaced by sheer imitation.

Summarizing, we can argue that the equilibrium (λ, p) emerging in our model, expressed in terms of proportion of informed traders and price of the risky asset, is surely reachable by boundedly rational agents supported by a credible and simple learning heuristic. In the remainder of this article, we will analyze by means of simulation the behavior of the model and the properties of the equilibrium for different configurations of the parameters.

3 Results

We present here the results of many simulations of the model that was coded in NetLogo, see Wilensky [11]. The parameters of the model belong to two main sets. The first one is related to structural features including d, R, the information cost c and the variances σ_ϵ^2 and σ_θ^2. The second group of parameters are associated to the learning procedure: the horizon T and the number of couples h. In what follows, we assume the individual parameter α, the risk aversion of the agents, to be constant across the population.

We run 27000 simulations for 10000 periods,[4] setting $N = 1000, d = 1.1, R = 1.01$ and systematically allowing $\sigma_\epsilon^2, \sigma_\theta^2$ to take all the values in $\{0.03, 0.06, ..., 0.30\}$, $c \in \{0.1, 0.2, ..., 0.5\}$, $T \in \{1, 4, 16\}$, $h \in \{15, 30\}$ and $\alpha \in \{1, 2, 3\}$. We check for the absence of transient effects in the simulations, running the experiments beginning with a fraction λ_0 of informed agents in $\{0.3, 0.5, 0.7\}$.

To ease exposition, we define a benchmark model in which $\alpha = 2, \sigma_\epsilon^2 = 0.06, \sigma_\theta^2 = 0.09, c = 0.3, T = 1, h = 15, \lambda_0 = 0.5, k = 0.05, \sigma_\eta^2 = 0.005$. Figure 1 depicts the time series of prices and the evolution of λ_t in a representative run of the benchmark case. The average price in this specific simulation is 1.091 and λ fluctuates around an average value of 0.57.

Table 1 shows the estimate of the mean and standard deviation of the stationary distribution of the random variables p^* (equilibrium price) and λ^*. We denote them, respectively, by μ_p and σ_p (for price) and μ_λ and σ_λ (for the

[4] For simulations involving $T = 16$, 12000 periods.

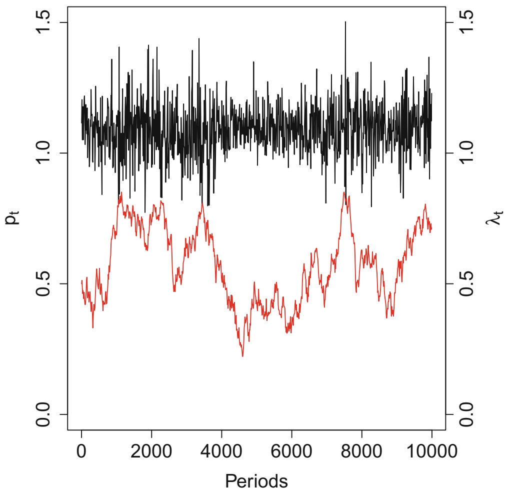

Fig. 1. Price p_t (black, left axis) and fraction of informed traders λ_t (red, right axis) in the benchmark configuration ($\alpha = 2, \sigma_\epsilon^2 = 0.06, \sigma_\theta^2 = 0.09, c = 0.3, T = 1, h = 15, \lambda_0 = 0.5$) (Color figure online)

fraction of informed). The table displays results obtained in the benchmark configuration, as well as in other situations in which one parameter alone is tilted with respect to the benchmark.

Table 1. Main outcomes of the model for the benchmark configuration and selected variations. Every entry in the table is the average of 6 independent simulations (3 values for $\lambda_0 \times 2$ values for h). Each row shows the sample mean and the standard deviation of the stationary distribution of price and of the share of informed traders

	μ_p	σ_p	μ_λ	σ_λ
Benchmark	1.089	0.091	0.557	0.168
$\alpha = 1$	1.089	0.239	0.958	0.015
$\sigma_\theta^2 = 0.12$	1.091	0.160	0.925	0.024
$c = 0.4$	1.089	0.038	0.111	0.039

The time-average of the price is remarkably close to $d/R = 1.089$ for all our simulations and all cases depicted in Table 1. Indeed, the mean aggregate demand in Eq. (2) is proportional to $d - Rp$ for all traders and any prolonged deviation of p from d/R is corrected by the market maker, who would detect and act on any sustained imbalance. In the benchmark case (first row), approximately 56% of traders are informed at each period, with the presence of oscillations that are visible in Fig. 1, as well as in the standard deviation σ_λ taking the value 0.17.

A further point is related to the heteroskedasticity in prices: it can be seen in Fig. 1 that the standard deviation of prices is high in every time interval featuring an high share of informed individuals, λ_t. The intuition is that if a large portion of traders acquires the information θ_t and acts accordingly, high pressure (either upward or downward) is put on prices. Conversely, if only a few individuals have superior information, they exploit the signal without heavily affecting the market.

Table 1 also exemplifies other general and sensible outcomes of the model. Other things being fixed, less risk-averse traders (with $\alpha = 1$) acquire the information in over 95% of cases. Raising σ_θ^2 to 0.12 (third row) makes the signal more informative and appealing, thus raising μ_λ from 0.56 to 0.93. As expected, if the cost is increased from 0.3 to 0.4 (fourth row), the fraction of informed traders drops to about 11%.

One of the most interesting features of the model is the relationship between μ_λ, the "informativeness" of the signal, σ_θ^2, and the baseline level of the idiosyncratic noise, σ_ϵ^2. In Fig. 2 we plot the average equilibrium fraction of informed agents as a function of σ_θ^2, for different values of σ_ϵ^2.

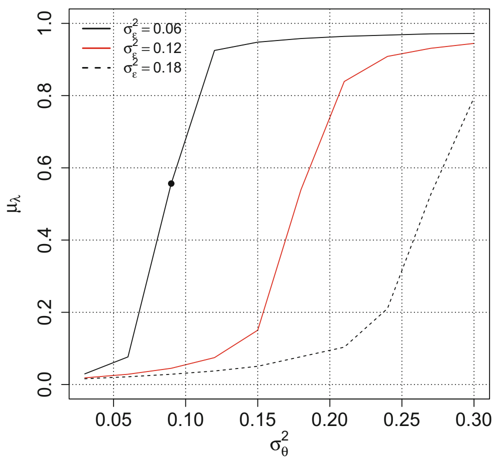

Fig. 2. Equilibrium fraction of the informed agents as a function of σ_θ^2, for $\sigma_\epsilon^2 = 0.06, 0.12, 0.18$. The thick point represents the benchmark configuration ($\alpha = 2, \sigma_\epsilon^2 = 0.06, \sigma_\theta^2 = 0.09, c = 0.3, T = 1$). (Color figure online)

Without doubt, scarcely informative signals lead to small values of μ_λ and, conversely, increasing σ_θ^2 to substantial levels (with respect to σ_ϵ^2) ultimately pushes the fraction of the informed ones to very high values approaching, but never reaching, 100%. Comparing the left solid line with the right dashed one, for instance, it can be seen that tripling σ_ϵ^2 for a given informativeness, greatly reduces the number of informed traders. This follows from the reduced utility of the signal embedded in a setup where the background noise is prevalent. In

Fig. 3, we further investigate how μ_λ depends on σ_ϵ^2 and σ_θ^2, plotting the set of couples of the parameters for which μ_λ takes the values 0.1, 0.5 and 0.9, respectively (i.e., we plot three contour levels of the function $\mu_\lambda(\sigma_\epsilon^2, \sigma_\theta^2)$, keeping fixed all the other parameters at the level taken in the benchmark case).

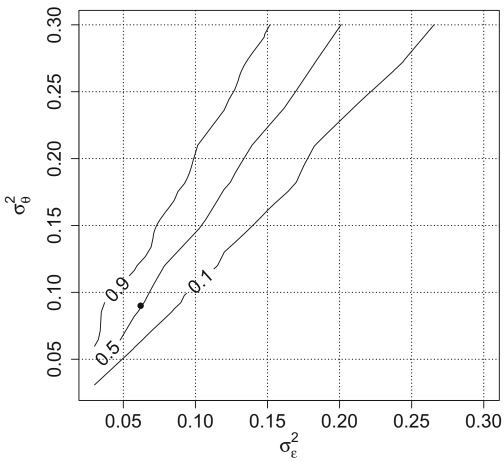

Fig. 3. The set of $(\sigma_\epsilon^2, \sigma_\theta^2)$ where μ_λ takes the values 0.1, 0.5 and 0.9 (from right to left). The thick point represents the benchmark case and, for example, when σ_ϵ^2 and σ_θ^2 are 0.10, we have $\mu_\lambda \sim 10\%$

The rightmost line, relative to combinations for which $\mu_\lambda = 10\%$, shows that very few traders acquire the information when the informativeness is about the same size of the background noise (for the given values of the other parameters). The steeper leftmost line depicts configurations in which 90% of the agents are informed: this roughly occurs when σ_θ^2 is about the double of σ_ϵ^2. More importantly, the almost linear shape of the three contour lines strongly suggests that the signal-to-noise ratio $\rho = \sigma_\theta^2/\sigma_\epsilon^2$ is crucial in shaping the emerging equilibrium. For instance, whenever $\rho = 3/2$, we obtain $\mu_\lambda \sim 50\%$ (or, put differently, one in two agents buys the signal). Essentially, and rather sensibly, it looks as if the important thing in a market where information is costly is the signal-to-noise ratio faced by the traders (and not the peculiar variances of the sources of noise determining the dividends paid by the risky asset).[5]

We turn then to another significant insight provided by the extended set of simulations we have run. An increase in the time T between two learning rounds has a deep effect on the outcomes. Recall that h traders are randomly paired with other agents every T periods and learn by imitation, copying the behavior

[5] It should be stressed that the result obviously depends also on the other parameters, say α and c, but the conclusion that the signal-to-noise ratio plays an important role robustly holds for all the combinations we have simulated (specific details are not discussed for brevity).

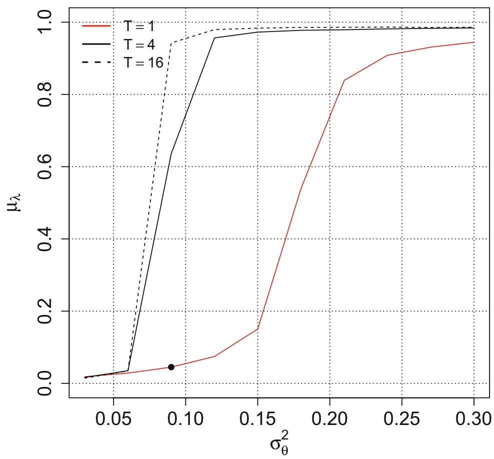

Fig. 4. Equilibrium fraction of the informed agents as a function of σ_θ^2, for $T = 1, 4, 16$. The thick point represents the benchmark configuration ($\alpha = 2, \sigma_\epsilon^2 = 0.12, \sigma_\theta^2 = 0.09, c = 0.3, T = 1$)

of the peer if the accumulated wealth over T periods proves to be higher. So far we have discussed the case $T = 1$, corresponding to a market populated by short-term traders, who revise their decision to buy (or not to buy) the signal based on the profits gained in a single period. Figure 4, resembling what was done in Fig. 2, depicts the average equilibrium fraction of informed agents as a function of σ_θ^2, for different values of T. For this analysis, we alter the benchmark configuration setting $\sigma_\epsilon^2 = 0.12$ (see the red line in Fig. 2) to help the reader spot the increase in μ_λ implied by an increase in T. Observe, for instance, the thick point in the figure where less than 10% of traders acquire the information when $T = 1$: keeping the other parameters fixed, it suffices to raise T from 1 to 4 to boost the number of the informed ones to over 60%. The effect is much stronger if $T = 16$, a situation in which λ abruptly increases to almost 100% as soon as the informativeness σ_θ^2 reaches 0.09.

Even though a greater T may be somewhat interpreted in terms of "stubbornness" on the part of the agents, who insist in using their strategy and less frequently try to learn from their peers, an alternative explanation is in order: if learning takes place less frequently, traders have more time to learn whether the signal is, on average, worth the cost. In the present setup, the second effect is clearly predominant over the first and the benefits of slower (and more accurate) learning outweigh what can be gained with frequent (but necessarily noisier) assessments of trading performances.

Our findings highlight that weakly informative signals, which will not be acquired if short-term gains are sought for, are nevertheless valuable in the long run (i.e., provided that multiple periods are considered and used to average the profits and adjust behaviour). A similar idea was presented in the seminal paper on the Santa Fe Artificial Stock Market, Arthur et al. [12], in which it is pointed out that "where investors explore alternative expectational models at a low rate, the market settles into the rational-expectations equilibrium", whereas high-frequency exploration leads to more hectic behavior, with rich psychology and "technical trading emerges, temporary bubbles and crashes occur". In a related fashion, fast revisions of the decision to acquire information in our model lead to episodic adoption of the signal, volatility bursts and reduced use of aggregate information; on the contrary, slow learning with larger T increases the fraction of traders who get and use the information and, in this sense, produces what can be deemed as more rational outcomes.

4 Conclusions

The model presented in this paper features boundedly rational agents who have the option to acquire a costly informative signal on the return of the risky asset. Paying for the information would make predictions more accurate and reduce the residual risk, increasing the traded volume. Uninformed agents, on the contrary, face greater uncertainty and typically buy or sell less but do not bear any additional cost. Some agents in the population learn by imitation whether their choice to get (or not to get) the information is convenient by comparing their profits every T periods with the ones of another peer.

We have shown that, similarly to Grossman and Stiglitz [4], the model converges to an equilibrium where stationary distributions for the price and the fraction of informed traders can be described. Numerical simulations demonstrate how the adoption of information increases with the informativeness of the signal or, more precisely, with the signal-to-noise ratio. Moreover, less information is used by more risk-averse agents and when the signal is more costly.

The volatility of the price, which is adjusted by a market maker based on the excess demand, is not constant and notably depends on the fluctuation of the fractions of informed, as a larger (smaller) magnitude of returns is observed on average when λ is big (small). The channel through which high values of λ cause large shocks is the larger volume of trades prompted by informed agents.

Finally, we have shown how the interval T between two learning rounds affects the equilibrium dynamics. An increase in T effectively results in slower learning, a situation where traders stick to the same strategy in a stubborn way for several periods, apparently renouncing the frequent opportunity to revise their strategy. However, slow learning also gives the chance to assess the value of the signal in a much more accurate (and less noisy) way. In our model, the latter effect prevails on the former and a larger T gives rise to markets where (many) more agents acquire the signal, suggesting that even weakly informative signals can be profitable if the trading gain they produce is assessed over a long horizon.

Further research would be needed to clarify how the market structure affects the results that, in the present framework, are mainly driven by the signal-to-noise ratio and by the frequency with which learning is activated. Indeed, the market maker used in the model adjusts the price only after demands are revealed and transactions are cleared. As such, he may delay or hamper the chance to infer the signal from the price and in principle a formal auction may disclose information in a more efficient and timely manner (even though it is difficult to ascertain under which dimensions this would be good or bad).

References

1. Schredelseker, K.: Is the usefulness approach useful? Some reflections on the utility of public information. In: McLeay, S., Riccaboni, A. (eds.) Contemporary Issues in Accounting Regulation, pp. 135–153. Springer, Boston (2001). https://doi.org/10.1007/978-1-4615-4589-7_8
2. Kurlat, P., Veldkamp, L.: Should we regulate financial information? J. Econ. Theory **158**, 697–720 (2015)
3. Veldkamp, L.L.: Media frenzies in markets for financial information. Am. Econ. Rev. **96**(3), 577–601 (2006)
4. Grossman, S.J., Stiglitz, J.E.: On the impossibility of informationally efficient markets. Am. Econ. Rev. **70**(3), 393–408 (1980)
5. Routledge, B.R.: Adaptive learning in financial markets. Rev. Financ. Stud. **12**(5), 1165–1202 (1999)
6. Routledge, B.R.: Genetic algorithm learning to choose and use information. Macroecon. Dyn. **5**(02), 303–325 (2001)
7. Chen, T.T., Zheng, B., Li, Y., Jiang, X.F.: Information driving force and its application in agent-based modeling. Phys. A: Stat. Mech. Its Appl. **496**, 593–601 (2018)
8. Billett, M.T., Garfinkel, J.A., Yu, M.: The effect of asymmetric information on product market outcomes. J. Financ. Econ. **123**(2), 357–376 (2017)
9. Krichene, H., El-Aroui, M.A.: Artificial stock markets with different maturity levels: simulation of information asymmetry and herd behavior using agent-based and network models. J. Econ. Interact. Coord. **13**(3), 511–535 (2018)
10. Cont, R., Bouchaud, J.P.: Herd behavior and aggregate fluctuations in financial markets. Macroecon. Dyn. **4**(2), 170–196 (2000)
11. Wilensky, U.: NetLogo. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL (1999). http://ccl.northwestern.edu/netlogo/
12. Arthur, W.B., Holland, J.H., LeBaron, B., Palmer, R., Tayler, P.: Asset pricing under endogenous expectations in an artificial stock market. In: Arthur, W., Lane, D., Durlauf, S. (eds.) The Economy as an Evolving, Complex System II, pp. 15–44. Addison Wesley, Redwood City (1997)