

RESEARCH OUTPUTS / RÉSULTATS DE RECHERCHE

Adaptive expectations, confirmatory bias, and informational efficiency

Aldashev, Gani; Carletti, Timoteo; Righi, Simone

Publication date:
2010

Document Version
Early version, also known as pre-print

[Link to publication](#)

Citation for published version (HARVARD):

Aldashev, G, Carletti, T & Righi, S 2010, *Adaptive expectations, confirmatory bias, and informational efficiency*. naXys Technical Reports Series, no. 2, vol. 1, vol. 1, 2 edn, FUNDP. Namur center for complex systems, Namur.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.



Adaptive Expectations, Confirmatory Bias, and Informational Efficiency

by Gani Aldashev, Timoteo Carletti and Simone Righi

Report naXys-02-2010

23 September 2010



Namur Center for Complex Systems

University of Namur

8, rempart de la vierge, B5000 Namur (Belgium)

<http://www.naxys.be>

Adaptive Expectations, Confirmatory Bias, and Informational Efficiency*

Gani Aldashev[†] Timoteo Carletti[‡] Simone Righi[§]

September 23, 2010

Abstract

We study the informational efficiency of a market with a single traded asset. The price initially differs from the fundamental value, about which the agents have noisy private information (which is, on average, correct). A fraction of traders revise their price expectations in each period. The price at which the asset is traded is public information. The agents' expectations have an adaptive component and a social-interactions component with confirmatory bias. We show that, taken separately, each of the deviations from rationality worsen the information efficiency of the market. However, when the two biases are combined, the degree of informational inefficiency of the market (measured as the deviation of the long-run market price from the fundamental value of the asset) can be non-monotonic both in the weight of the adaptive component and in the degree of the confirmatory bias. For some ranges of parameters, two biases tend to mitigate each other's effect, thus increasing the informational efficiency.

Keywords: informational efficiency, confirmatory bias, agent-based models, asset pricing.

JEL codes: G14, D82, D84.

*The authors thank the National Bank of Belgium for financial support.

[†]Corresponding author. Department of Economics and CRED, University of Namur (FUNDP). Mailing address: Department of Economics, 8 Rempart de la Vierge, 5000 Namur, Belgium. Email: gani.aldashev@fundp.ac.be.

[‡]NaXys - Namur Center for Complex Systems, University of Namur (FUNDP). Mailing address: 8 Rempart de la Vierge, 5000 Namur, Belgium. Email: timoteo.carletti@fundp.ac.be.

[§]Department of Economics, University of Namur (FUNDP). Mailing address: Department of Economics, 8 Rempart de la Vierge, 5000 Namur, Belgium. Email: simone.righi@fundp.ac.be.

1 Introduction

In most economic interactions, individuals possess only partial information about the value of exchanged objects. For instance, when a firm "goes public", i.e. launches an initial public offering of its shares, no financial market participant has the complete information concerning the future value of the profit stream that the firm would generate. The fundamental question, going back to Hayek (1945), is then: To which extent the market can serve as the aggregator of this dispersed information? In other words, when is the market informationally efficient, i.e. that the market price converges to the value that would obtain if all market participants were to have full information about the fundamental value of the asset exchanged?

Most of the studies that address this question are based on the assumption that individual market participants are rational. Under full rationality, the seminal results on the informational efficiency of centralized markets have been established by Grossman (1976), Wilson (1977), Milgrom (1981), and, for decentralized markets, by Wolinsky (1990), Blouin and Serrano (2001), and Duffie and Manso (2007).

However, research in experimental economics and behavioral finance indicates that traders do not behave in the way consistent with the full-rationality assumption. For instance, Haruvy et al. (2007) find that traders have adaptive expectations, i.e. they give more importance to the past realized price of the asset than the fully-rational agent would. Along a different dimension, Rabin and Schrag (1999) discuss the evidence that individuals suffer from the so-called confirmatory (or confirmation) bias: they tend to discard the new information that substantially differs from their priors. Understanding whether (and under which conditions) the financial markets are informationally efficient when agents do not behave fully rationally remains an open question.

In this paper, we study the informational efficiency of a market with a single traded asset. The price initially differs from the fundamental value, about which the agents have noisy private information (which is, on average, correct). A fraction of traders revise their price expectations in each period. The price at which the asset is traded is public information. The agents' expectations have an adaptive component and a social-interactions component with confirmatory bias.

We show that, taken separately, each of the deviations from rationality worsen the information efficiency of the market. However, when the two biases are combined, the degree of informational inefficiency of the market (measured as the deviation of the long-run market

price from the fundamental value of the asset) can be non-monotonic both in the weight of the adaptive component and in the degree of the confirmatory bias. For some ranges of parameters, two biases tend to mitigate each other's effect, thus increasing the informational efficiency.

The paper is structured as follows. Section 2 presents the setup of the model. Section 3 derives analytical results for each bias taken separately. In Section 4, we present the simulation results when two biases are combined. Finally, Section 5 discusses the implication of our results and suggests some future avenues for research.

2 The model

Consider a market with N participants, each endowed with an initial level of wealth equal to $W_0 > 0$. The amount $L_0 \in (0, W_0]$ is in liquid form. Time is discrete (e.g. to mimic the daily opening and closure of a financial market), denoted with $t = 0, 1, \dots$. Market participants trade a single asset, whose price in period t we denote with P_t . This price is public information. Prices are normalized in such a way that they belong to the interval $[0, 1]$.

At the beginning of each period t , every agent i can place an order to buy or short sell 1 unit of the asset, on the basis of her expectation about the price for period t , denoted with $P_t^{e,i}$. Placing an order implies a fixed, small but positive *transaction cost* c , i.e. $0 < c \ll 1$. At the end of the period, each agent i learns the price P_t at which the trade is settled (as explained below).

The agent i then constructs her price expectation for the next period and decides to participate in the trading in period $t + 1$ according to the *expected next-period gain*, i.e. if

$$|P_{t+1}^{e,i} - P_t| - c > 0. \quad (1)$$

Moreover, she participates as a *buyer* if her price expectation for the next period exceeds the current price, i.e.

$$P_{t+1}^{e,i} > P_t, \quad (2)$$

or as a *seller* if, on the contrary,

$$P_{t+1}^{e,i} < P_t. \quad (3)$$

The way in which agents form their next-period price expectations differs from the standard rational-expectation benchmark in the following way. First deviation is the fact that

agents give weight to the past public prices, i.e. they have (partially) adaptive expectations. Secondly, they can influence each other's expectations via social interactions with confirmatory bias.

Formally, suppose that in every period a fraction, $\gamma \in [0, 1]$, of the agents makes a revision of their price expectations. An agent revises her price expectation by analyzing the past price of the asset and by randomly encountering some other agent (at zero cost), and possibly exchanging her own price expectation with this partner. In these encounters, the agents have a *confirmatory bias*, i.e. each agent tends to ignore the information coming from the other agent if it differs too much from her own. If, on the contrary, this difference is not too large, i.e. smaller than some fixed threshold, which we denote with σ , then the agent incorporates this information into her price expectation. The remaining $(1 - \gamma)N$ agents do not revise their expectations in the current period.

Summarizing, the expectation formation process of agent i meeting agent j is:

$$P_{t+1}^{e,i} = \alpha P_t + (1 - \alpha) \begin{cases} P_t^{e,i} & \text{if } |P_t^{e,i} - P_t^{e,j}| \geq \sigma \\ \frac{P_t^{e,i} + P_t^{e,j}}{2} & \text{otherwise} \end{cases}, \quad (4)$$

and it is analogous for $P_{t+1}^{e,j}$. Here, α measures the relative weight of the past price. If $\alpha = 1$, the agents have purely adaptive expectations (and social interactions play no role). If $\alpha = 0$ and $\sigma = 1$ the agents (that revise their expectations) completely disregards the past and fully integrate all the information coming from the social interactions.

Our objective is to analyze the price formation under the different values of the parameters α , σ , and γ .

Concerning the market microstructure, we assume that the market is centralized, with a simple price response to excess demand. In other words, the market mechanism is similar to the Walrasian auctioneer. More precisely, the price formation mechanism functions as follows:

1. There exists a hypothetical price at period $t + 1$ that would (approximately) equate the number of buy-orders and sell-orders. Let us denote it with P_{t+1}^* . From (2) and (3), P_{t+1}^* is the solution of the equation:

$$n_B(x) = n_S(x),$$

where $n_B(x)$ and $n_S(x)$ are the numbers of buyers and sellers at price x . Whenever

there are several solutions to this equation, P_{t+1}^* denotes the average of the values that solve the equation.

2. Out of equilibrium, the price adjustment depends on the size of the excess demand or excess supply relative to the size of the population; in other words, denoting $\beta(x) = |n_B(x) - n_S(x)|/N$, the price adjustment process is:

$$P_{t+1} = \beta(P_t)P_{t+1}^* + (1 - \beta(P_t))P_t. \quad (5)$$

Thus, the deviation from the equilibrium does not disappear instantly. However, the price moves in the direction that eliminates the excess demand or supply, and, moreover, the speed of adjustment depends on the size of disequilibrium (relative to the size of the population).¹

3. Given that each agent that participates in the market in period t places an order for one unit of the asset, the number of exchanges that occurs is $\min\{n_B(P_t), n_S(P_t)\}$. Then, each seller i updates her wealth by $W_{t+1}^i = W_t^i + P_t - P_{t+1} - c$, and her liquidity by $L_{t+1}^i = L_t^i + P_t - c$. Similarly, for a buyer j , we have $W_{t+1}^j = W_t^j - P_t + P_{t+1} - c$ and $L_{t+1}^j = L_t^j - P_t - c$.
4. If an agent's liquidity dries up to zero, then she leaves the market. At her place, at the beginning of the next period enters a new agent with wealth W_0 , liquidity L_0 , and the next-period price expectation randomly drawn from the $[0, 1]$ interval.

In this setting, consider an initial public offering (IPO) of the asset. At time $t = 0$, the asset gets introduced in the market at some price P_0 . Let's also suppose that, *on average*, the agents have the unbiased information about its fundamental value. In particular, let's suppose that the initial price expectations of the agents is a uniform distribution in the $[0, 1]$ interval, i.e. the fundamental value of the asset is $\frac{1}{2}$. However, the initial price P_0 differs from the fundamental value. The questions that we pose are:

- Does the market price P_t converge to the fundamental value of the asset?
- If not, how large is the deviation of the long-run price P_t as $t \rightarrow \infty$, from the fundamental value?

¹We avoid the shortcoming of assuming a constant $\beta(P_t)$. As discussed by LeBaron (2001), if $\beta(P_t)$ is assumed to be constant, the behavior of the simulated market is extremely sensitive to the value of β , which makes it difficult to interpret the results.

- How does this deviation depend on the weight of history α (i.e. the "adaptiveness" of agents' expectations), the prominence of the confirmatory bias of the traders σ , and the frequency with which agents adjust their expectations, γ ?

3 Analytical results

We can characterize analytically the answers to the above questions for some of the values of the parameters. This requires a further assumption that the number of market participants (N) and every agent's initial wealth and liquidity (W_0 and L_0) are sufficiently large.

3.1 Purely adaptive expectations

Consider first the case where agents discard the social interactions and consider only the past price. In other words, $\alpha = 1$ in (4). We analyze separately two sub-cases: (i) all agents revise their expectations in every period, i.e. $\gamma = 1$; and (ii) only a fraction of agents revise their expectations in every period, i.e. $\gamma < 1$.

(i) In the case $\gamma = 1$, Eq. (4) simply reduces to:

$$P_{t+1}^{e,i} = P_t \text{ for any } i \text{ and } t.$$

However, this is also the hypothetical price that equates buyers and sellers, i.e. $P_{t+1}^* = P_t$, and thus $\beta(P_t) = 0$. Finally, from (5) we get $P_{t+1} = P_t$. This means that the market price doesn't evolve: $P_t = P_0$ in every period. Intuitively, if all agents revise their expectations in every period and have purely adaptive expectations, once the initial price P_0 is announced, every agent immediately revises her next-period expectation, substituting it with P_0 . Given that every agents does so, no agents is interested in trading, and the price does not evolve.

(ii) Next, consider the case $\gamma < 1$, with γN being sufficiently large. Without loss of generality, suppose that $P_0 > 1/2$. We prove that the market reaches the long-run equilibrium, after a few periods, with the long-run market price deviating from the initial price by a value smaller than $c(1 - \gamma)$.

We need the following preliminary result.

Proposition 1. *Consider a population of agents divided into two groups: agents in the first group, whose size is N_1 , have expectations uniformly distributed in $[0, 1]$, and agents in the second group, whose size is $N_2 \gg N_1$, all have the price expectation equal to some fixed*

$\hat{P} \in (0, 1)$. Then, the price P^* defined at point 1 is given by $P^* = \hat{P} - c$ if $\hat{P} > 1/2$ and $P^* = \hat{P} + c$ if $\hat{P} < 1/2$. If $\hat{P} = 1/2$, then $P^* = \hat{P}$.

Proof. We consider only the case $\hat{P} > 1/2$ (the proof for the case $\hat{P} < 1/2$ is analogous, while the case $\hat{P} = 1/2$ is trivial). Define the functions $\theta_c(x, P)$

$$\theta_c(x, P) = \begin{cases} 1 & \text{if } x > P + c \\ 0 & \text{otherwise} \end{cases},$$

and $\eta_c(x, P) = 1 - \theta_c(x, P)$. Then for a sufficiently large N_2 , the numbers of sellers and buyers at price $x \in [0, 1]$ are respectively given by:

$$\begin{aligned} n_S(x) &= (x - c)N_1 + N_2\theta_c(x, \hat{P}) \quad \text{and} \\ n_B(x) &= (1 - x - c)N_1 + N_2\eta_c(x, \hat{P}). \end{aligned} \tag{6}$$

This follows from the trading protocol, given that for a price sufficiently close to \hat{P} (i.e. with a deviation less than c), only the first group of agents participates in the trading, and that the expectations are uniformly distributed in the first group. On the other hand, if $x < \hat{P} - c$ (or $x > \hat{P} + c$), the second group also participates in the trading as buyers (sellers).

Then, the difference in the number of buyers and sellers is

$$\Delta(x) = (1 - 2x)N_1 + N_2 \left(\eta_c(x, \hat{P}) - \theta_c(x, \hat{P}) \right), \tag{7}$$

and, therefore, P^* becomes the price at which the sign of $\Delta(x)$ changes (or the average of these values, if more than one exist). We can then easily prove that

$$\Delta_- = \lim_{x \rightarrow (\hat{P}-c)^-} \Delta(x) = (1 - 2\hat{P} + 2c)N_1 + N_2 > 0. \tag{8}$$

Finally, using the assumption $N_2 \gg N_1$, we get $\Delta_- > 0 > \Delta(x)$ for all $x > \hat{P} - c$. This implies that $P^* = \hat{P} - c$. ■

This proposition has the following

Corollary 2. *Suppose the assumptions of Proposition 1 hold. If a third group of agents (of arbitrary size) with price expectation \tilde{P} , such that $|\tilde{P} - \hat{P}| < c$, joins the market, then P^* does not change.*

We can now analyze the market dynamics under the assumptions $\alpha = 1$ and $\gamma < 1$ with γN large.

During the first period $N_1 = (1 - \gamma)N$ agents do not revise their expectations. These expectations are uniformly distributed in $[0, 1]$. Contrarily, $N_2 = \gamma N$ agents revise their expectations, which now becomes the IPO price P_0 (i.e. $P_1^{e,i} = P_0$). Proposition 1 ensures that $P_1^* = P_0 - c$. Moreover, the size of market disequilibrium is small: using the definition, we get $\beta(P_0) = (2P_0 - 1)(1 - \gamma)$. Finally, the end-of-period 1 price P_1 will be

$$P_1 = \beta(P_0)(P_0 - c) + (1 - \beta(P_0))P_0 = P_0 - \beta(P_0)c. \quad (9)$$

Note that this price is c -close to P_0 , given that $\beta(P_0)$ is small.

During the next period, γN agents revise their expectations, while $(1 - \gamma)N$ agents do not revise. Then, on average, we have that the second group contains $N_2 = \gamma(1 - \gamma)N$ agents, for which $P_2^{e,i} = P_1^{e,i} = P_0$, the first group contains $N_1 = (1 - \gamma)^2$ agents (who do not revise their initial expectations), and, moreover, there exists a third group of agents, of size $N_3 = \gamma(2 - \gamma)$, for whom $P_2^{e,i} = P_1 = P_0 - \beta(P_0)c$. We can then apply Corollary 2 and conclude that $P_2^* = P_0 - c$. Computing the next-period market disequilibrium $\beta(P_1)$, we can easily observe that $\beta(P_1) \sim (1 - \gamma)^2$. Therefore, the next-period price P_2 will be

$$P_2 = \beta(P_1)(P_0 - c) + (1 - \beta(P_1))P_1 \sim P_0 - \beta(P_0)c + \mathcal{O}(1 - \gamma)^2 \sim P_0 - c(1 - \gamma) + \mathcal{O}(1 - \gamma)^2. \quad (10)$$

Thus, the market price varies as long as there exist agents that have not yet revised their initial expectations. However, the market price does not move too far from P_0 . Assuming the extreme-case scenario where in every period the same $(1 - \gamma)$ agents happen to be the ones that do not revise their expectations, the number of periods that pass before the market price converges to its steady-state value equals $-\log N / \log(1 - \gamma)$.

Numerical simulations presented in Figure 1 confirm our theoretical findings.

[Insert Figure 1 about here]

3.2 Social interactions

Next, consider the setting where the agents' expectations have no adaptive component (i.e. $\alpha = 0$), and the agents that revise their expectations rely on the social interactions with other agents. Then, the relevant parameter is the extent of the confirmatory bias (σ) that the agents have. We derive analytical results for the the cases of the extreme form of confirmatory bias ($\sigma \ll 1$) and for that of virtually no bias ($\sigma \sim 1$).

3.2.1 Social interactions: large confirmatory bias

Consider the extreme form of confirmatory bias, i.e. that whenever two agents meet, neither of them adjusts her price expectation, no matter how close their past-period expectations are. We will prove that the market is fully informationally efficient in the long-run, but convergence to this efficient outcome takes arbitrary long time.

In every period, γN agents engage in social interactions (without influencing each other's price expectations). This implies that no agent revises her price expectation. Therefore, the mean price expectation (which we denote with \tilde{P}^e) does not change either; namely, from (4),

$$\tilde{P}_{t+1}^e = \tilde{P}_t^e.$$

Under the assumption that the initial price expectations are uniformly distributed, we obtain $P_{t+1}^* = \tilde{P}_{t+1}^e = 1/2$. The market disequilibrium is thus given by

$$\beta(P_t) = |1 - 2P_t|. \quad (11)$$

Therefore, the market price evolves according to the equation

$$P_{t+1} = \frac{|1 - 2P_t|}{2} + (1 - |1 - 2P_t|) P_t = P_t + \frac{|1 - 2P_t|(1 - 2P_t)}{2}. \quad (12)$$

Let us define the mapping

$$f(P) = \begin{cases} P - \frac{(2P-1)^2}{2} & \text{if } P \geq 1/2 \\ P + \frac{(2P-1)^2}{2} & \text{if } P < 1/2 \end{cases}. \quad (13)$$

The evolution of the market price is determined by the dynamic system

$$P_{t+1} = f(P_t). \quad (14)$$

This mapping has a unique fixed point at $P = 1/2$. Moreover, this is an attractor, whose strength decreases the closer we are to the fixed point: $P_t - 1/2 \sim a/t$.

Finally, note that if $P_0 = 1/2$, then (12) implies that $P_t = 1/2$ for all t .

[Insert Figure 2 about here]

3.2.2 Social interactions: small confirmatory bias

Let us now consider the opposite extreme, i.e. the agent that uses for updating her expectation that of her partners in social interactions even when such expectations diverge radically

from hers. Assume for the moment that all agents revise their expectations in every period ($\gamma = 1$). We will prove that in this case the market is fully informationally efficient in the long-run, and that convergence to this efficient outcome occurs within in a finite number of periods (that essentially depends on the transaction cost c).

Given that the expectation-revision rule (4) with $\alpha = 0$ and $\gamma = 1$ preserves the mean price expectation, we trivially get

$$\tilde{P}_{t+1}^e = \tilde{P}_t^e = 1/2.$$

This follows from the fact that the initial price expectations are uniformly distributed in $[0, 1]$, hence with average value $1/2$, which also equals the hypothetical Walrasian-auctioneer price $P_t^* = 1/2$. Moreover, the equation (4) implies that price expectations follow the Deffuant dynamic (Deffuant et al. 2000, Weisbuch et al. 2002). In other words, the dispersion of price expectations, denoted with Δ_{P_e} , shrinks to zero according to (see the left panel of Figure 3):

$$\Delta_{P_e}(t) \sim \frac{1}{2^{t/2}}.$$

[Insert Figure 3 about here]

Because of the transaction cost is positive, the market activity stops once all the expectations fall inside the interval whose width is smaller than $2c$. This happens after a time $\hat{T} \sim -2(1 + \log_2 c)$.

Let us assume now that the price expectations have a dispersion large enough, so that the market activity still does not stop. Then, we can easily compute the market disequilibrium $\beta(P_t)$:

$$\beta(P_t) = 2^{t/2}|1 - 2P_t|,$$

which implies that the next-period price is given by:

$$P_{t+1} = P_t + 2^{t/2-1}|1 - 2P_t|(1 - 2P_t). \quad (15)$$

Let us introduce the auxiliary variable x , defined as $P_t - 1/2 = x_t/2^{t/2}$. This allows us to fully describe the market price evolution with the dynamic system given by the function $g(x)$:

$$g(x) = \begin{cases} \sqrt{2}x - 2\sqrt{2}x^2 & \text{if } x \geq 0 \\ \sqrt{2}x + 2\sqrt{2}x^2 & \text{if } x < 0 \end{cases}. \quad (16)$$

This mapping has three fixed points: $x = 0$ (unstable), and $x = \pm(2 - \sqrt{2})/4$, that are stable.

We can thus conclude that, as long as the market activity continues, the market price converges to $1/2$, given that in all cases $P_t = \frac{x_t}{2^{t/2}} + \frac{1}{2} \rightarrow \frac{1}{2}$. These findings are supported by numerical simulations, whose results we report in Figure 4.

[Insert Figure 4 about here]

In a similar fashion, we can study the case $\gamma \neq 1$. In this case, in every period γN agents revise their expectations and, because they have a very small confirmatory bias (i.e. $\sigma \sim 1$), they influence each other's expectations. Then, overall there is a tendency for the expectations to converges (because of the process driven by Eq. (4) with $\alpha = 0$). We should keep in mind, however, that in every period $(1 - \gamma)N$ agents do not revise their expectations.

Zero weight given to the past prices in forming the next-period expectation ($\alpha = 0$) implies that the mean price expectation does not change, $\tilde{P}_{t+1}^e = \tilde{P}_t^e$. Hence $\tilde{P}_t^e = 1/2$ and $P_t^* = 1/2$ for all t . Moreover the expectations are distributed in an interval whose width (denoted with $\Delta_{P_e}(t)$) shrinks to zero, but slower than in the case $\gamma = 1$. The simulations presented in the right panel of Figure 3, allow us to see that this narrowing of expectation dispersion follows approximately the law

$$\Delta_{P_e}(t) \sim \frac{1}{2^{q_\gamma t}},$$

where $q_\gamma = a\gamma + b$, $a = 0.61 \pm 0.02$ and $b = -0.13 \pm 0.01$, independent of σ .

Thus, the price disequilibrium can be estimated as:

$$\beta(P_t) \sim 2^{q_\gamma t} |1 - 2P_t|,$$

which implies the following price dynamics:

$$P_{t+1} = P_t + 2^{q_\gamma t - 1} |1 - 2P_t| (1 - 2P_t).$$

Introducing a new variable y_t such that $P_t = 1/2 + y_t/2^{q_\gamma t}$, we obtain the following difference equation for the evolution of y_t :

$$y_{t+1} = 2^{q_\gamma} y_t - 2^{q_\gamma + 1} |y_t| y_t.$$

This mapping has three fixed points, $y = 0$ (unstable), and $y = \pm(1 - 2^{q_\gamma})/2^{q_\gamma + 1}$ (stable). We can finally conclude that, similar to the results above, the market price converges to $1/2$ as long as the market run. The market activity stops once all the expectations fall inside the interval whose width is smaller than $2c$. This happens after time $\hat{T} \sim -(1 + \log_2 c)/q_\gamma$.

4 Simulation results

When the price expectations of agents have both the adaptive component and confirmatory bias, obtaining analytical results is beyond reach. We thus proceed by running numerical simulations. In what follows, we vary the values of α and σ , from 0 to 1, in steps of 0.01. For each pair of values (α, σ) the market is simulated 10 times. The cost of a trading transaction is fixed at $c = 0.005$. Each simulation can run for 100 steps, being this a time interval large enough for the market price to converge to the steady-state. Note that in the simulations we define the steady-state as the situation in which the difference between the market prices in periods t and $t + 1$ differ by a value smaller than 0.0001. We then look at the degree of market informational inefficiency in the long-run, i.e. how far the market price diverges from the fundamental value of the asset (averaging across the 10 simulations). We also look at the average volatility of the market price, as measured by the standard deviation in the market price in the last 90% of the steps of the simulation.

The agents have a relatively low level of wealth. Remember that if the outcomes of the trading strategy of a trader lead to losses that, accumulated over several periods, exhaust her wealth, she quits the market and is replaced by another trader with a randomly drawn initial price expectation. Given that traders have a relatively low level of wealth, a certain number of them will quit the market and this implies that the turnover rate of traders is relatively high. This means that some amount of noise gets continuously injected into the market.

[Insert Figure 5 about here]

Figure 5 (Panels A, B, and C) report the informational inefficiency of the market (as measured by the divergence of the steady-state market price from the fundamental value of the asset) for the cases in which the fraction of agents that revise their expectations in every period is $\gamma = 0.2, 0.5,$ and $0.8,$ respectively. The market inefficiency is a function of the weight of the adaptive component in the price expectations of traders α and of the degree of confirmatory bias σ . Colors closer to dark blue indicate lower level of market inefficiency, while those closer to dark red indicate higher inefficiency. Figure 6 describes the volatility of the market price, while Figure 7 shows the average number of traders that exit the market as their wealth hits the zero bound.

[Insert Figures 6 and 7 about here]

Analyzing these figures, we obtain the following findings.

Fixing the value of σ , as we move from the extreme-left point ($\alpha = 0$) to the right, the average deviation of the long-run market price from the fundamental value (0.5) first decreases and then increases - at least for some values of σ . In other words:

Proposition 3. *Market inefficiency can be non-monotonic in the weight of the adaptive component (α) in the price expectations of agents.*

In all three figures, we find that for very high values of α , the degree of the informational inefficiency of the market is very high. Clearly, when traders put a large weight to the past price in forming expectations, the initial price becomes very important. When receiving information which indicates that the value of the asset is low (even in the absence of confirmatory bias), traders tend to give little weight to it - basically, all that matters is the past price. In this case, the initial price strongly influences the aggregate expectation formation process (the expectations of all agents quickly converge upwards to some point between the initial price and the fundamental value) and, given that in our case the initial price strongly differs from the fundamental value, the long-run market price stays largely above the fundamental value.

Consider now the situation with the most extreme form of confirmatory bias, i.e. all traders completely ignore the information that comes from others. As α declines, the traders give less weight to the past prices and more weight to their own expectations of the previous period. Therefore, the agents whose initial expectations are very low do not move their next-period expectations upwards too much. At the same time, the market price keeps falling, driven by the Walrasian auctioneer (which also implies the downward move in the expectations of the agents whose initial expectations are high). These two inter-related processes - the upward drift of price expectations of initially low-expectation agents and the downward pressure on the market price converge to some value relatively close to the fundamental one.

As α declines further, we observe that the market inefficiency rises again. This is due to the fact that for the lower values of α , the first process (upward move in expectations of the initially low-expectation agents) becomes slower than the second one (i.e. downward move in the market price). Thus, the low-expectation agents keep making negative profits, eventually hit the zero-wealth bound, and exit (we can note this by looking at Figure 7: the number of agents that exit the market increases at the bottom-left part of the figure). There

is a sufficiently high exit rate of these agents from the market so as to soften the downward move in the market price, which means that the long-run price at which the system settles down is higher than in the situation in which the exit of traders is negligible.

As the degree of confirmatory bias of agents becomes smaller (i.e. the value of σ increases), the channel that leads to the exit of low-expectation traders softens down, as there is now an additional mechanism that creates an upward pressure on the expectations of those traders: the integration of information that comes from their peers. Notice (in Figure 7) that the exit rate is lower at the higher values of σ .

Furthermore, comparing across the three panels of Figure 5, one notes that as the fraction of agents that revise their expectations in each period (γ) increases, the areas of σ in which the non-monotonicity in α occurs becomes smaller:

Proposition 4. *The tendency of the market inefficiency to be non-monotonic in α is stronger, the lower is the fraction of agents that revise their price expectations in each period.*

This happens because the higher frequency of expectation revision (larger γ) and the likelihood to integrate the information coming from other traders (larger σ) act in a complementary fashion: if the rate of revision of price expectations is relatively low, the "openness of mind" (i.e. low degree of confirmatory bias) has a relatively small effect on the mitigation of the exit channel. It is only when the agents revise their expectations relatively frequently, that the "openness of mind" starts to have a real bite, and the upward-sloping part on the left side of the relation between market inefficiency and α starts to disappear.

Let's now fix the value of $\alpha = 0.25$ on Figure 5A, $\alpha = 0.1$ on Figure 5B, and $\alpha = 0.05$ on Figure 5C. As we move from the point at the bottom ($\sigma = 0$) upwards, the average deviation of the steady-state market price from the fundamental value (0.5) first decreases and then increases. In other words:

Proposition 5. *Market inefficiency can be non-monotonic in the degree of confirmatory bias of agents.*

The first part of the non-monotonic relationship is easy to explain: as an agent suffers less from confirmatory bias, she starts to integrate at least some of the information about the fundamentals contained in the price expectations of another trader (incidentally, this phenomenon occurs only when the adaptive component in the price expectation is relatively small).

But why the market inefficiency would *rise* as the agents become even more "open-minded"? To understand this, we need to note that this phenomenon occurs only when the adaptive component is not too small. Then, the fact that the initial price differs substantially from the fundamental value plays a role. The agents have an early-stage upward drift in expectations. At the same time, the market price starts to fall. If the agents are very "open-minded", this implies that they 'excessively' integrate the early upward price drift into their expectations, which, in turn, implies that the price at which the market settles in the long run is relatively high. If, instead, the agents' confirmatory bias is stronger, the decline in the market price is faster than the 'propagation' of the upward-drifting expectations: this is why the market settles at the price relatively close to the fundamental value.

This analysis suggests a very interesting and potentially more general insight: when market participants suffer from more than one deviation from fully rational behavior (in our case, adaptive expectations and confirmatory bias), at least in some range of bias parameters, the two biases mitigate each other. Given our analysis, it should not be difficult to construct examples of asset markets with traders that have multiple sources of biases, that exhibit the same price behavior as under full rationality, and in which the price behavior would deviate from the full-rationality benchmark as soon as one of the bias sources is eliminated.

Next, looking across the three panels of Figure 5, we also note that the values of α at which we have found the non-monotonicity of the market inefficiency decrease at higher values of γ . In other words:

Proposition 6. *The weight of the adaptive component in the expectations (α) at which the non-monotonicity of the market inefficiency in the degree of confirmatory bias occurs decreases with the fraction of agents that revise their expectations in every period.*

To capture the intuition behind this result, we need to conduct the following thought experiment. Let's fix a point with sufficiently high values of α and σ , for example (0.4, 0.4). Next, let's increase the frequency of revision of expectations by the agents, from $\gamma = 0.2$ to $\gamma = 0.8$. We then observe that the market inefficiency increases. This indicates the importance of the frequency with which agents revise their price expectations for the propagation of the 'excessive' integration of the upward drift into the expectations, as noted above. In other words, at higher frequency of expectation revision, this 'excessive' integration of the upward drift channel swamps the opposite (i.e. the quantity-of-information) channel more easily. The quantity-of-information channel starts to play a role only when the early-stage upward

drift is sufficiently small (i.e., history weighs relatively little in the expectation formation).

If we measure the degree of market inefficiency, while varying σ along a fixed α , in ranges different from those where the non-monotonicity occurs (for example, $\alpha = 0.1$ and $\alpha = 0.3$ on Figure 5A), we see that at the higher values of α the relationship between the degree of market inefficiency and σ is negative, while at the lower values of α , this relationship is positive. In other words:

Proposition 7. *The slope of the relationship of market inefficiency in the degree of confirmatory bias (σ) can be of opposite sign at different values of the weight of the adaptive component (α).*

The above discussion has already hinted at the potential explanation why this reversal of the relationship occurs. At sufficiently high values of α , the early-stage upward drift is very important and the smaller confirmatory bias of agents only helps to propagate this drift into the price expectations. At sufficiently low values of α , the early-stage upward drift matters much less and the smaller confirmatory bias becomes beneficial for the informational efficiency of the market, because it helps to integrate more of the relatively unbiased information into the expectations. In other words, in both cases the smaller confirmatory bias (i.e. higher σ) plays the role of the catalyzer; what differs in the two cases is the initial unbiasedness of expectations.

5 Conclusion

This paper has studied the informational efficiency of an agent-based financial market with a single traded asset. The price initially differs from the fundamental value, about which the agents have noisy private information (which is, on average, correct). A fraction of traders revise their price expectations in each period. The price at which the asset is traded is public information. The agents' expectations have an adaptive component (i.e. the past price influences their future price expectation to some extent) and a social-interactions component with confirmatory bias (i.e. agents exchange information with their peers and tend to discard the information that differs too much from their priors).

We find that the degree of informational inefficiency of the market (measured as the deviation of the long-run market price from the fundamental value of the asset) can be non-monotonic both in the weight of the adaptive component and in the degree of the




Figure 1: Purely adaptive expectations ($\alpha = 1$). Left panel: semilog plot of the market inefficiency as a function of γ . Right panel: time to convergence to the steady state as a function of γ . There are $N = 1000$ agents, the transaction cost $c = 0.005$. Each point represents the average over 50 simulations. \circ : numerical simulations, solid line : analytical results.

confirmatory bias. For some ranges of parameters, two biases tend to mitigate each other's effect, thus increasing the informational efficiency.

Our findings complement the well-known results in the theory of markets showing that the allocative efficiency can be obtained even under substantial deviation from individual rationality of agents (Gode and Sunder 1993, 1997). We show that deviations from individual rationality, under certain conditions, can also facilitate the informational efficiency of markets. The key condition for this property is that the various behavioral biases that agents possess should mutually dampen their effects on the price dynamics.

Given the potential importance of this insight for financial economics, the natural extension of this work is to test its' predictions experimentally. This would require to construct experimental financial markets with human traders, similar to the setting of Haruvy et al. (2007), with the additional feature of allowing agents to share their information (in some restricted form). The outcomes of interest in such an experiment would be both the evolution of market price of the asset and the elicited price expectations of traders.

References

- [1] Blouin, M., and Serrano, R. 2001. A decentralized market with common values uncertainty: Non-steady states. *Review of Economic Studies* **68**: 323-346.
- [2] Deffuant, G., Neau, D., Amblard, F., and Weisbuch, G. 2000. Mixing beliefs among interacting agents. *Advances in Complex Systems* **3**: 87-98.




Figure 2: Social interaction without adaptive component ($\alpha = 0$ and $\sigma \sim 0$). Left panel: semilog plot of the difference between the simulated and analytical market inefficiency with $\sigma = 5.0 \cdot 10^{-4}$ and $P_0 = 0.6$ or 0.8 . Right panel: difference between the simulated and analytical market inefficiency as a function of σ . There are $N = 1000$ agents, the transaction cost $c = 0.005$.




Figure 3: Evolution of $[\log_2 \Delta_{P_e}(t)]/t$ as a function of σ (left panel) and as a function of γ (right panel). There are $N = 1000$ agents, the transaction cost $c = 0.005$.




Figure 4: Social interaction without adaptive component ($\alpha = 0$ and $\sigma \sim 0$). Left panel: difference between the simulated and analytical market inefficiency as a function of σ . Right panel: the ration of the simulated time to convergence over the analytical one, σ . There are $N = 1000$ agents, the transaction cost $c = 0.005$.




Figure 5: The degree of market efficiency as a function of the adaptive component (α) and the confirmatory bias (σ). Panel A: $\gamma = 0.2$, Panel B $\gamma = 0.5$ and Panel C : $\gamma = 0.8$. The initial price is 0.9.




Figure 6: Volatility of the market price (calculated over the last 90% of the time span) as a function of the adaptive component (α) and the confirmatory bias (σ). Panel A: $\gamma = 0.2$, Panel B $\gamma = 0.5$ and Panel C : $\gamma = 0.8$. The initial price is 0.9.




Figure 7: Number of agents exiting the market as a function of the adaptive component (α) and the confirmatory bias (σ). Panel A: $\gamma = 0.2$, Panel B $\gamma = 0.5$ and Panel C : $\gamma = 0.8$. The initial price is 0.9.

- [3] Duffie, D., and Manso, G. 2007. Information percolation in large markets. *American Economic Review* **97**: 203-209.
- [4] Gode, D., and Sunder, S. 1993. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy* **101**: 119-137.
- [5] Gode, D., and Sunder, S. 1997. What makes markets allocationally efficient? *Quarterly Journal of Economics* **112**: 603-630.
- [6] Grossman, S. 1976. On the efficiency of competitive stock markets when traders have diverse information. *Journal of Finance* **31**: 573-585.
- [7] Haruvy, E., Lahav, Y., and Noussair, C. 2007. Traders' expectations in asset markets: Experimental evidence. *American Economic Review* **97**: 1901-1920.
- [8] Hayek, F. 1945. The uses of knowledge in society. *American Economic Review* **35**: 519-530.
- [9] LeBaron, B. 2001. A builder's guide to agent-based financial markets. *Quantitative Finance* **1**: 254-261.
- [10] Milgrom, P. 1981. Rational expectations, information acquisition, and competitive bidding. *Econometrica* **49**: 921-943.
- [11] Rabin, M., and Schrag, J. 1999. First impressions matter: A model of confirmatory bias. *Quarterly Journal of Economics* **114**: 37-82.
- [12] Weisbuch, G., Deffuant, G., Amblard, F., and Nadal, J. 2002. Meet, discuss and segregate! *Complexity* **7**: 55-63.
- [13] Wilson, R. 1977. Incentive efficiency of double auctions. *Review of Economic Studies* **44**: 511-518.
- [14] Wolinsky, A. 1990. Information revelation in a market with pairwise meetings. *Econometrica* **58**: 1-23.