

Antico, Andrea; Bottazzi, Giulio; Giachini, Daniele

Working Paper

On the evolutionary stability of the sentiment investor

LEM Working Paper Series, No. 2022/09

Provided in Cooperation with:

Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies

Suggested Citation: Antico, Andrea; Bottazzi, Giulio; Giachini, Daniele (2022) : On the evolutionary stability of the sentiment investor, LEM Working Paper Series, No. 2022/09, Scuola Superiore Sant'Anna, Laboratory of Economics and Management (LEM), Pisa

This Version is available at:

<https://hdl.handle.net/10419/259552>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

INSTITUTE
OF ECONOMICS



Scuola Superiore
Sant'Anna

LEM | Laboratory of Economics and Management

Institute of Economics
Scuola Superiore Sant'Anna

Piazza Martiri della Libertà, 33 - 56127 Pisa, Italy
ph. +39 050 88.33.43
institute.economics@sssup.it

LEM

WORKING PAPER SERIES

On the evolutionary stability of the sentiment investor

Andrea Antico ^a
Giulio Bottazzi ^a
Daniele Giachini ^a

^a Institute of Economics and EMbeDS Department, Scuola Superiore Sant'Anna, Pisa, Italy.

2022/09

March 2022

ISSN(ONLINE) 2284-0400

On the evolutionary stability of the sentiment investor

Andrea Antico^{*1}, Giulio Bottazzi^{†1}, and Daniele Giachini^{‡1}

¹Institute of Economics & Department EMbeDS, Scuola Superiore Sant'Anna, Piazza Martiri della Libertà 33, 56127 Pisa - Italy

May 30, 2022

Abstract

The behavioral finance literature attributes the persistent market misvaluation observed in real data to the presence of deviations from rational thinking of the actors involved. Cognitive biases and the use of simple heuristics can be described using expected utility maximizing agents that adopt incorrect beliefs. Along these lines, Barberis et al. (1998) introduce a model which is able to replicate the behavior of both under-reaction and over-reaction to news. The representative agent they consider is characterized by an imperfect learning model. An interesting question that emerges is if, and to what degree, the heuristic mechanism they propose is evolutionary stable, that is how resilient is their representative agent to other agents possibly trading in the market. In fact, if the biased agent asymptotically disappears from the market, then misvaluation patterns generated by its behavior do not survive in the long term. The present paper investigates this question comparing the performance of the agent described in Barberis et al. (1998) with the one of a pure Bayesian competitor.

JEL Classification: C60, D53, D81, D83, D91, G11, G12

Keywords: Learning, Market Selection, Investor Sentiment, Model Misspecification, Financial Markets

*andrea.antico@santannapisa.it

†giulio.bottazzi@santannapisa.it

‡daniele.giachini@santannapisa.it

1 Introduction

In the last decades, Behavioral Finance has provided important contributions to our understanding of financial markets' dynamics. Its fundamental tenet, built upon the ideas of Tversky and Kahneman (1974), is that pricing anomalies are due to widespread cognitive biases, as they generate limited information perception and processing and imperfections in the decision making process (see Hirshleifer, 2015, for a recent review of the topic). Two of the most pervasive and widely investigated pricing anomalies are short-run momentum and long-run reversal: extreme excess returns tend to be followed by excess returns with the same sign over short time periods (3 to 12 months) and revert over longer horizons (see, e.g., De Bondt and Thaler, 1985, 1987; Jegadeesh and Titman, 1993, 2001; Rouwenhorst, 1998; Moskowitz and Grinblatt, 1999; Balvers et al., 2000; Gropp, 2004; Griffin et al., 2005; Chui et al., 2010; Mukherji, 2011; Asness et al., 2013). These pieces of evidence, in striking contrast with the Efficient Market Hypothesis of Fama (1970), sprinkled the interest of behavioral economists and generated a breadth of behavioral theories explaining how the two effects originate (see, among the others, Barberis et al., 1998; Daniel et al., 1998; Hong and Stein, 1999; Daniel and Hirshleifer, 2015; Bottazzi et al., 2019; Luo et al., 2021).

Driven by psychological evidence, Barberis et al. (1998) - henceforth BSV - present a model where the two anomalies are explained in terms of conservatism (Edwards, 1982) and representativeness heuristic (Tversky and Kahneman, 1974). Specifically, they assume a sentiment investor whose learning behavior is characterized by both under- and overreaction and show that these behavioral characteristics are able to replicate momentum and reversal in price dynamics. Their sentiment investor is a representative agent and, as such, its cognitive patterns perfectly translate into price dynamics. But what happens if another agent, with a different behavior, enters the market? The answer is non trivial and the issue is relevant, since the BSV sentiment investor may be driven out by the newcomer and, hence, have a negligible influence on prices in the long run. Survival of biased agents is not a new topic. It is one of the crucial research questions of the literature stemming from Blume and Easley (1992) and investigating the Market Selection Hypothesis of Alchian (1950) and Friedman (1953). The basic framework is a pure-exchange economy with infinite horizon and infinite-lived agents where prices are set according to Walrasian market clearing conditions under either general equilibrium (see, e.g., Sandroni, 2000, 2005; Blume and Easley, 2006, 2009; Jouini and Napp, 2011; Kogan et al., 2006, 2017; Dindo and Massari, 2020; Beddock and Jouini, 2021; Bottazzi and Giachini, 2022) or temporary equilibrium (see, e.g., Evstigneev et al., 2009, 2016; Bottazzi et al., 2018; Bottazzi and Gi-

achini, 2017, 2019a,b).¹ One of the most important results is that, in a general equilibrium model with complete markets, only the agents with the most accurate beliefs survive in the long-run (Sandroni, 2000; Blume and Easley, 2006). Thus, agents affected by behavioral biases in perceiving and processing information – as in BSV – seems to be fated to vanish when an agent processing information correctly – i.e. according to Bayes’s rule – enters such a market. This is explicitly investigated in Sandroni (2005) and the author shows that a Bayesian agent eventually dominates over biased ones both in *learnable* and *unlearnable* settings. Such a finding is challenged by Massari (2020), who shows that an under-reacting agent never vanishes against a Bayesian one under model misspecification, that is, when the true data generating process does not belong to the set of models the agents believe possible. Thus, in such a case, a learning rule presenting some behavioral bias can produce accurate beliefs and generate an advantage in the selection struggle.

In this work we test whether the biased probability updating process described by BSV and adapted to a general equilibrium setting by Antico et al. (2022) can survive when competing against a Bayesian. More specifically, we consider an Arrow-Debreu economy with complete markets where two agents seek to maximize their expected utility of consumption over an infinite horizon under the beliefs generated by their learning processes. Both agents rely upon two models of the true data generating process. The first is a Markov chain where the probability of changing state is larger than the one of persisting. The second is also a Markov chain, but the probability of persisting is larger than the one of switching. The first agent computes the weights assigned to the two models and the conditional probabilities according to Bayes’ rule, while the second one uses the adapted BSV rule. We assume that the states of nature are drawn according to a Bernoulli trial. Thus, as in BSV, the true data generating process is different and simpler than the two Markovian models the agents believe possible. Hence, our framework is much closer to the setting of Massari (2020) than the one of Sandroni (2005). Our analysis shows that the sentiment investor is able to dominate the Bayesian agent for some combinations of parameter values while it is driven out of the market for others. The main difference is that BSV’s rule prescribes to persistently mix the two Markovian models, while Bayesian learning generically selects the best (mispecified) model among those available. As a consequence, the BSV updating

¹The topic has also been investigated in different frameworks. For instance, Benos (1998) and Kyle and Wang (1997) base their model on Kyle (1985) and analyze the trading of a risky security in a partial equilibrium setting with risk neutral traders. The price of the security is decided by a risk neutral market maker, hence no Walrasian mechanism exists. In a one-stage game, the authors show that an overconfident trader can attain a higher profit than a rational one. Benos (1998) shows that the same result holds when the market maker is risk averse. The choice of a non-Walrasian price fixing condition seems to be essential for these results.

rule appears more advantageous when the transition probabilities characterizing the Markov models are heavily misspecified but taking a convex combination of them one gets close to the truth. Thus, our results are in-between those of Sandroni (2005) and Massari (2020). On the one hand, the nature of model misspecification we assume decreases the evolutionary fitness Bayesian learning has in well specified settings. On the other, the peculiar learning behavior of the sentiment investor makes its beliefs very accurate in some cases and particularly incorrect in others.

2 Intertemporal consumption model

Consider a pure-exchange economy in discrete time (indexed by $t = 0, 1, \dots$), where N agents trade in complete markets over an infinite horizon and consume an homogeneous good in every period. At each time $t > 0$ the economy can be in one between two states, denoted by $s_t \in \{0, 1\}$. The infinite sequence of states $\sigma = (s_1, s_2, \dots, s_t, \dots)$ is a path or history of the economy and $\sigma_t = (s_1, s_2, \dots, s_t)$ is the partial history until time t included. The set of all the possible paths is denoted by Σ while Σ^t is the set of the 2^t partial histories until time t . The cylinder with base σ_t is $\mathcal{C}(\sigma_t) = \{\sigma \in \Sigma | \sigma = (\sigma_t, \dots)\}$ and \mathcal{F}_t is the σ -algebra generated by the cylinders $\mathcal{C}(\sigma_t)$. Thus, $(\mathcal{F}_t)_{t=0}^\infty$ is a filtration and \mathcal{F} is the σ -algebra generated by the union of filtrations. In accordance with the assumptions of BSV on (Σ, \mathcal{F}) , we assume an i.i.d. probability measure p such that the realization of a state does not depend on the partial history, $p(s_t = 1 | \sigma_t) = \pi \forall t, \sigma$. The operator E denotes the expectation and, when there is no subscript or superscript, it is computed with respect to p . Each agent i has a subjective probability measure p_i on (Σ, \mathcal{F}) , discussed in Section 3. We impose that subjective probability measures are absolutely continuous with respect to the true measure, that is $p_i(\sigma_t) > 0$ for any partial history σ_t .

Agents have logarithmic Bernoulli utility and a common and constant intertemporal utility discount factor $\beta > 0$. Let $e_i(\sigma_t) > 0$ be the endowment of agent i at partial history σ_t . The optimal consumption plan $c_i(\sigma_t) > 0$ of agent i solves

$$\begin{aligned} \max_{\{c_i(\sigma_t), \forall t, \sigma\}} U_i &= \sum_{t=0}^{\infty} \sum_{\sigma_t \in \Sigma^t} \beta^t p_i(\sigma_t) \log(c_i(\sigma_t)) \\ \text{subject to} \quad & \sum_{t=0}^{\infty} \sum_{\sigma_t \in \Sigma^t} q(\sigma_t) (e_i(\sigma_t) - c_i(\sigma_t)) \geq 0, \end{aligned} \tag{1}$$

where $q(\sigma_t)$ is the price of the Arrow-Debreu security paying 1 if partial history σ_t is realized and zero otherwise. We shall perform our analysis under general equilibrium: consumption plans solve the problem in (1) for any $i \in \{1, 2, \dots, N\}$

and markets clear in every period,

$$\sum_{i=1}^N c_i(\sigma_t) = \sum_{i=1}^N e_i(\sigma_t). \quad (2)$$

Assuming that agents have constant and homogeneous endowment, $e_i(\sigma_t) = e \forall i, \sigma_t$, the optimal consumption plans can be explicitly derived (Bottazzi and Giachini, 2022)

$$c_i(\sigma_t) = \frac{p_i(\sigma_t)}{\sum_{j=1}^N p_j(\sigma_t)} N e, \quad (3)$$

while the Arrow-Debreu security price reads

$$q(\sigma_t) = \frac{\beta^t}{N} \sum_{i=1}^N p_i(\sigma_t). \quad (4)$$

3 Beliefs structure

Agents base their individual probability models on two Markovian models described by the transition matrices

$$M_h : \begin{matrix} & s_{t+1} = 0 & s_{t+1} = 1 \\ \begin{matrix} s_t = 0 \\ s_t = 1 \end{matrix} & \begin{pmatrix} \pi_h & 1 - \pi_h \\ 1 - \pi_h & \pi_h \end{pmatrix} & \end{matrix}, \quad h = 1, 2. \quad (5)$$

Denoting with $M_h(s_{t+1} | s_t)$ the probability of observing the state s_{t+1} realized after the state s_t according to model $h = 1, 2$, it is $M_h(s_{t+1} = s_t | s_t) = \pi_h$ and $M_h(s_{t+1} \neq s_t | s_t) = 1 - \pi_h$. Following BSV, we assume $\pi_1 < 0.5 < \pi_2$, such that model 1 assigns a larger probability to remaining in the same state, while model 2 assigns a larger probability to switching to the other state with respect to the one lastly realized. At time t , if the partial history σ_t is realized, agent i builds its prediction about the next state as a convex weighting of the two models

$$p_i(s_{t+1} | \sigma_t) = \sum_{h=1}^2 w_{i,h}(\sigma_t) M_h(s_{t+1} | s_t) \quad (6)$$

with $w_{i,h}(\sigma_t) \in [0, 1]$, $\sum_{h=1}^2 w_{i,h}(\sigma_t) = 1$, and $p_i(\sigma_t) = \prod_{\tau=1}^t p_i(s_\tau | \sigma_{\tau-1})$. While agents share the same baseline Markov models, they differ in the way they derive their weights. In particular, we shall consider two alternative weight updating rules. The first rule is the standard Bayesian learning, while the second rule is derived from the under-reaction/over-reaction behavioral bias model introduced by BSV and adapted to a general equilibrium setting by Antico et al. (2022).

3.1 Bayesian agent

A Bayesian agent i assigns the weight proportionally to the likelihood of the models. Let $L_h(\sigma_t) = \prod_{\tau=1}^t M_h(s_{\tau+1} | s_\tau)$ be the likelihood of model M_h , then it is

$$w_{i,h}(\sigma_t) = \frac{w_{i,h}(\sigma_0) L_h(\sigma_t)}{\sum_{k=1}^2 w_{i,k}(\sigma_0) L_k(\sigma_t)}. \quad (7)$$

Given the Markov property of the two models one can easily derive a recursive definition

$$w_{i,1}(\sigma_{t+1}) = \frac{w_{i,1}(\sigma_t) M_1(s_{t+1}|s_t)}{\sum_{k=1}^2 w_{i,k}(\sigma_t) M_k(s_{t+1}|s_t)} \quad (8)$$

and $w_{i,2}(\sigma_{t+1}) = 1 - w_{i,1}(\sigma_{t+1})$. Notice that if on a given history σ it is, on average, $M_1(s_t|s_{t-1}) > M_2(s_t|s_{t-1})$ (or equivalently $\lim_{t \rightarrow \infty} L_2(\sigma_t)/L_1(\sigma_t) = 0$), then $\lim_{t \rightarrow \infty} w_{i,1}(\sigma_t) = 1$. That is, the Bayesian agent assigns all the weight to the best model in terms of likelihood. This is always the case if only one among the considered models is correct.

3.2 BSV agent

Following the prescription of BSV and its adaptation in Antico et al. (2022), we consider an alternative weights updating process which depends on two parameters $\lambda_1, \lambda_2 \in [0, 1]$:²

$$w_{i,1}(\sigma_{t+1}) = \frac{C_{1-\lambda_1, \lambda_2}(w_{i,1}(\sigma_t)) M_1(s_{t+1}|s_t)}{C_{1-\lambda_1, \lambda_2}(w_{i,1}(\sigma_t)) M_1(s_{t+1}|s_t) + C_{\lambda_1, 1-\lambda_2}(w_{i,1}(\sigma_t)) M_2(s_{t+1}|s_t)}, \quad (9)$$

with $C_{\alpha, \beta}(x) = \alpha x + \beta(1-x)$. This rule generalizes (8), which can be recovered setting $\lambda_1 = \lambda_2 = 0$.³ The values of λ_1 and λ_2 determine how the rule behaves (see the proof of proposition 2 of BSV and Section 4 of Antico et al., 2022, for further discussions). Here we focus on two cases: $0 < \lambda_1 + \lambda_2 < 1$ and $\lambda_1 + \lambda_2 = 1$.⁴ If $0 < \lambda_1 + \lambda_2 < 1$ holds, then there exists two numbers $0 < \underline{w} < \bar{w} < 1$ such that

$$\underline{w} = C_{1-\lambda_1, \lambda_2}(\underline{w}) \pi_1 / (C_{1-\lambda_1, \lambda_2}(\underline{w}) \pi_1 + C_{\lambda_1, 1-\lambda_2}(\underline{w}) \pi_2),$$

²In the notation of BSV the weight $w_{i,1}(\sigma_t)$ is the quantity q_t reported at the end of page 322.

³BSV justify the prescription in (9) as a simil-Bayesian updating with respect to a fictional two-state model the agent has in mind. This justification, admittedly unconvincing, is immaterial for the present analysis. In other terms, we discuss the merits of the behavioral model described by (9) without discussing its degree of rationality with respect to some mental model possibly adopted by the agent.

⁴The case $\lambda_1 + \lambda_2 > 1$ is not considered because it generates a paradoxical behavior: the higher the evidence in favor of one model the lower the weight attached to it, see also Antico et al. (2022).

$$\bar{w} = C_{1-\lambda_1, \lambda_2}(\bar{w})(1 - \pi_1) / (C_{1-\lambda_1, \lambda_2}(\bar{w})(1 - \pi_1) + C_{\lambda_1, 1-\lambda_2}(\bar{w})(1 - \pi_2)).$$

If $w_{i,1}(\sigma_0) \in (\underline{w}, \bar{w})$ then $w_{i,1}(\sigma_t) \in (\underline{w}, \bar{w}) \forall \sigma_t$ and the updating rule in (9) defines a reinforcement learning process (Barberis et al., 1998). We shall assume $w_{i,1}(\sigma_0) \in (\underline{w}, \bar{w})$ when $0 < \lambda_1 + \lambda_2 < 1$. Hence, irrespective of the realization of the process, the BSV learning rule never converges to a single model, even when the realized sequence heavily favors one model over the other. This particular feature of the BSV updating is in striking contrast with Bayesian learning and can provide some evolutionary advantage in cases of model misspecification as the one we are considering here. If, instead, $\lambda_1 + \lambda_2 = 1$ holds, the updating rule becomes

$$w_{i,1}(\sigma_{t+1}) = \frac{\lambda M_1(s_{t+1}|s_t)}{\lambda M_1(s_{t+1}|s_t) + (1 - \lambda)M_2(s_{t+1}|s_t)} \quad (10)$$

with $\lambda = \lambda_2 = 1 - \lambda_1$. Thus, in this case, the weights attached to the two models do not change over time. Hence, no matter the sequence of states observed, the agent will maintain the same mixture of models.

4 Market selection and evolutionary stability

Our analysis aims at characterizing the asymptotic distribution of consumption shares of a BSV agent when it is not alone in the market (as, for instance, in the setting of Barberis et al., 1998). In particular, we want to evaluate its evolutionary stability – that is, its ability to asymptotically maintain a positive consumption share – when it trades with a Bayesian. We start our investigation with a numerical exercise. We assume $N = 2$ and impose that agent 1 is Bayesian and updates its weights according to (8). Agent 2, instead, follows the BSV behavioral bias model updating its weights as in (9). In Figure 1 we report some example of the time behavior of $c_2(\sigma_t)/(2e)$, the consumption share of the BSV agent. Each plot contains ten independent simulations. In all the examples we keep $\pi = 0.5$. In panel 1a and 1b we assume $\lambda_1 + \lambda_2 = 1$, while in panel 1c it is $\lambda_1 + \lambda_2 < 1$. Notice that, sooner or later, the consumption share converges to 1 for some parameter values and to 0 for other parameter values. That is, fluctuations in consumption share disappear over time letting either agent 1 or agent 2 – depending on the parameter settings – consume the whole aggregate endowment. Comparing the moment in which fluctuations fade out, one also notices that parameter values influence the speed of convergence toward the final consumption share distribution.

Before analyzing asymptotic consumption dynamics in detail, we need to clarify the meaning of some terms we shall use to indicate the different selection outcomes.

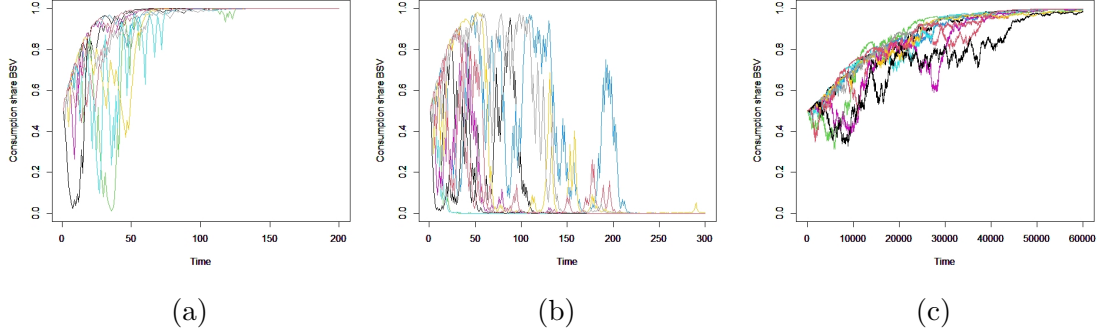


Figure 1: Consumption share dynamics of agent 2. **Parameter values:** (a): $\pi = 0.5$, $\pi_1 = 0.3$, $\pi_2 = 0.8$, $\lambda_1 = 0.25$, $\lambda_2 = 0.75$; (b): $\pi = 0.5$, $\pi_1 = 0.2$, $\pi_2 = 0.6$, $\lambda_1 = 0.25$, $\lambda_2 = 0.75$; (c): $\pi = 0.5$, $\pi_1 = 0.49$, $\pi_2 = 0.51$, $\lambda_1 = 0.2$, $\lambda_2 = 0.2$.

Definition 4.1. Agent i *vanishes* on a path σ if and only if

$$\lim_{t \rightarrow \infty} c_i(\sigma_t) = 0.$$

Agent *survives* on a path σ if and only if

$$\limsup_{t \rightarrow \infty} c_i(\sigma_t) > 0.$$

Agent i *dominates* on a path σ if and only if

$$\lim_{t \rightarrow \infty} c_i(\sigma_t) = 2e.$$

If the previous limits hold p -almost surely, we say, respectively, that agent i *vanishes*, agent i *survives*, agent i *dominates*.

From (3) we have that $c_i(\sigma_t)/c_j(\sigma_t) = p_i(\sigma_t)/p_j(\sigma_t)$. The probability $p_i(\sigma_t)$ is the likelihood of path σ_t in the model used by agent i . More accurate models have an higher likelihood.

Definition 4.2. If on path σ it is $\lim_{t \rightarrow \infty} p_j(\sigma_t)/p_i(\sigma_t) = 0$ then agent i is more accurate than agent j on σ . Agent i is more accurate than agent j if the previous limit applies p -almost surely, that is over a set of paths σ of measure one.

Since the consumption of agents is bounded, we have the following.

Proposition 4.1. *If agent i is more accurate than agent j on path σ , than agent j vanishes on path σ . If agent i is more accurate than agent j , than agent j vanishes.*

If there are only two agents $i = 1, 2$, and agent 1 vanishes, then agent 2 dominates and vice-versa. Thus, the analysis of the relative consumption patterns in a pairwise comparison of agents is reduced to the analysis of the relative accuracy of their respective learning models. In the next section we will investigate agents' relative accuracy using a mix of analytical and numerical results. The analytical investigation is facilitated by the following consideration (Sandroni, 2000; Blume and Easley, 2006). Define the *average relative entropy* of agent i on history σ as

$$\overline{D}_p(p_i, \sigma) = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \log \frac{p(s_\tau | \sigma_{\tau-1})}{p_i(s_\tau | \sigma_{\tau-1})}. \quad (11)$$

The expression in (11) is always non-negative and it is zero only if the agent's beliefs asymptotically converge to the true probabilities. If the limit in (11) exists for both agent i and j , then it is

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log \frac{p_i(\sigma_t)}{p_j(\sigma_t)} = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t (\log p_i(s_\tau | \sigma_{\tau-1}) - \log p_j(s_\tau | \sigma_{\tau-1})) = \overline{D}_p(p_j, \sigma) - \overline{D}_p(p_i, \sigma), \quad (12)$$

where we have added and subtracted the logarithm of true probability $p(s_\tau | \sigma_{\tau-1})$. The average relative entropy in (11) can be considered a measure of accuracy of agent i 's beliefs along path σ and can be used instead of the likelihood to compute models' relative accuracy. If agent i has p -almost surely a lower relative entropy than agent j , then agent j vanishes. The advantage of using the relative entropy as a measure of accuracy is that in the case of an i.i.d. true process like ours, applying the Strong Law of Large Numbers for Martingale Differences, one has

$$\overline{D}_p(p_i, \sigma) = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \pi \log \frac{\pi}{p_i(1 | \sigma_{\tau-1})} + (1 - \pi) \log \frac{1 - \pi}{p_i(0 | \sigma_{\tau-1})}, \quad (13)$$

which is generally an easier quantity to compute.

4.1 Bayes accuracy

As already mentioned, a Bayesian learner selects the model with larger likelihood on a path. An immediate consequence of this is that the Bayesian agent asymptotically selects the most accurate model. If for $k, j \in \{1, 2\}$ it is $\lim_{t \rightarrow \infty} L_k(\sigma_t)/L_j(\sigma_t) = 0$, then from (7) it is $\lim_{t \rightarrow \infty} w_{i,h}(\sigma_t) \rightarrow 1$. Thus, the Bayesian agent is as accurate as the most accurate model. Computing the average relative entropy of the underlying Markov models is relatively easy. From (13) it is

$$\overline{D}_p(M_h, \sigma) = \pi^2 \log \frac{\pi}{\pi_h} + \pi(1 - \pi) \log \frac{\pi(1 - \pi)}{(1 - \pi_h)^2} + (1 - \pi)^2 \log \frac{1 - \pi}{\pi_h}. \quad (14)$$

It follows that, to compute the relative entropy of a Bayesian agent it is enough to compute (14) for $h = 1, 2$, rank the two, and assign to the agent the lowest one.

4.2 BSV accuracy

Computing the average relative entropy of the BSV agent is more difficult. Indeed, as argued, an agent that follows the BSV weight updating never converges to a single model but keeps mixing the two. As already mentioned, its conditional probability can eventually converge to a fixed mixture of the two model, but that only happens along paths with zero measure, such as sequences of all equal states or perfectly alternated states. In the case $\lambda_1 + \lambda_2 = 1$, we are able to analytically compute the average relative entropy along sequences with full probability. That is made possible by the particular structure of weights which does not depend upon the previous weighting but only on the last realized states. For the more general case $\lambda_1 + \lambda_2 < 1$, instead, we shall rely on a numerical exercise.

4.2.1 Case $\lambda_1 + \lambda_2 = 1$

Setting $\lambda_2 = 1 - \lambda_1 = \lambda$ and substituting (10) in (6), one gets

$$p_2(s_{t+1}|\sigma_t) = p_2(s_{t+1}|(s_t, s_{t-1})) = \frac{\lambda M_1(s_t|s_{t-1})M_1(s_{t+1}|s_t) + (1 - \lambda)M_2(s_t|s_{t-1})M_2(s_{t+1}|s_t)}{\lambda M_1(s_t|s_{t-1}) + (1 - \lambda)M_2(s_t|s_{t-1})}.$$

From (13) it is

$$\bar{D}_p(p_2, \sigma) = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \left(\pi \log \frac{\pi}{p_2(1|(s_t, s_{t-1}))} + (1 - \pi) \log \frac{1 - \pi}{p_2(0|(s_t, s_{t-1}))} \right),$$

that, since the true process is i.i.d., reduces p -almost surely to

$$\begin{aligned} \bar{D}_p(p_2, \sigma) &= \pi^2 D_p(p_2, (1, 1)) \\ &\quad + \pi(1 - \pi) (D_p(p_2, (1, 0)) + D_p(p_2, (0, 1))) + (1 - \pi)^2 D_p(p_2, (0, 0)), \end{aligned} \quad (15)$$

with

$$D_p(p_2, (s_t, s_{t-1})) = \pi \log \frac{\pi}{p_2(1|(s_t, s_{t-1}))} + (1 - \pi) \log \frac{1 - \pi}{p_2(0|(s_t, s_{t-1}))}.$$

4.2.2 Case $0 < \lambda_1 + \lambda_2 < 1$

When $0 < \lambda_1 + \lambda_2 < 1$, we cannot compute the average relative entropy of the BSV agent in closed form. Hence, we estimate $\bar{D}_p(p_2, \sigma)$ numerically. The use

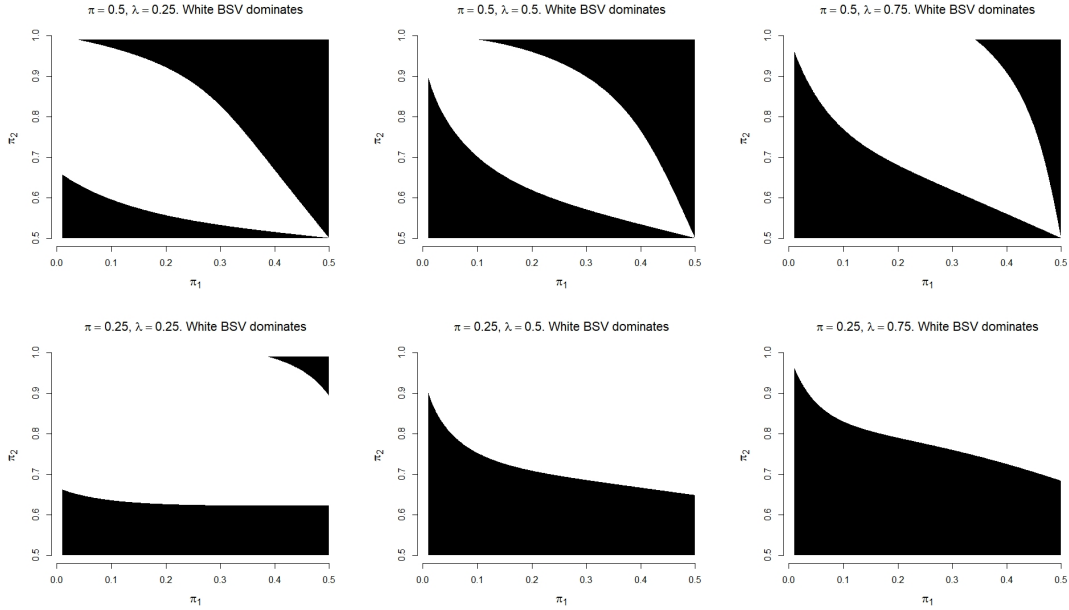


Figure 2: Selection outcomes for several combinations of (π_1, π_2) , for $\pi \in \{0.25, 0.5\}$, and for $\lambda \in \{0.25, 0.5, 0.75\}$. White: BSV dominates ($\Delta < 0$); Black: Bayes dominates ($\Delta > 0$).

of the relative entropy instead of the likelihood is preferred to easy the comparison with the analytical expressions obtained for the other learning models. We set a time horizon T sufficiently large and we draw the sequences of realizations $\sigma_{r,T} = (s_{r,1}, s_{r,2}, \dots, s_{r,T})$ with $r = 1, 2, \dots, R$. For every sequence r we iterate (9) and compute $D_p(p_2, \sigma_{r,t}) \forall t \in \{1, 2, \dots, T\}$. Then, we approximate the average relative entropy of the BSV agent along any sequence truncating the infinite average involved to T . We refine our estimate averaging over R independent sequences to obtain an estimated average relative entropy

$$\hat{D}_p(p_2) = \frac{1}{TR} \sum_{r=1}^R \sum_{t=1}^T D_p(p_2, \sigma_{r,t}). \quad (16)$$

4.3 Survival and dominance of sentiment investor

In this section we perform a series of pairwise comparison on the asymptotic behavior of agents' consumption. We first focus on the case $\lambda_1 + \lambda_2 = 1$, exploiting the closed form solution for the relative entropy of both agents. Consider the difference of the relative entropy of the learning models of the two agents on a given

path σ

$$\Delta = \overline{D}_p(p_2, \sigma) - \overline{D}_p(p_1, \sigma). \quad (17)$$

The quantity $\overline{D}_p(p_2, \sigma)$ is as in (15), while, as discussed in Section 4.1, it is $\overline{D}_p(p_1, \sigma) = \min\{\overline{D}_p(M_1, \sigma), \overline{D}_p(M_2, \sigma)\}$. If it is $\Delta > 0$ then agent 1 is more accurate than agent 2 and, by Proposition 4.1, agent 1 dominates and agent 2 vanishes. If, instead, $\Delta < 0$ it is the other way round: agent 2 is more accurate than agent 1 and agent 2 dominates and agent 1 vanishes. Due to the bi-stochastic nature of the transitions matrices defining the underlying Markov models, the relative entropy of both agents do not change if π is replaced with $1 - \pi$ (this is equivalent to relabeling the states of nature). Thus, it is sufficient to investigate the system for $0 < \pi \leq 0.5$. In Figure 2 we report the selection outcomes based on the sign of Δ for several combinations of (π_1, π_2) , for $\pi \in \{0.25, 0.5\}$, and for $\lambda \in \{0.25, 0.5\}$. Generically, there exist regions of the parameter space in which one agent dominates over the other and vice-versa. That is, the particular choice of the models' transition probabilities determines the final fate of the traders. Another implication that clearly emerges is that there are no regions in which both agents survive. This is generally true and directly implied by Proposition 4.1 combined with the analytic solutions in (14) and (15). Indeed, by direct inspection, one notices that, apart from hairline cases, Δ has a definite sign. Focusing on top row, that is, the cases in which $\pi = 0.5$, one notices that agent 2 dominates when the transition probabilities of the models are (more or less) evenly spread around the truth, that is $\pi_2 \simeq 1 - \pi_1$. In the case of less evenly spread transition probabilities, bottom row, the behavioral parameter λ influences more strongly the results, with $\lambda = 0.25$ favoring agent 2 over a larger set of (π_1, π_2) with $\pi_1 < 1 - \pi_2$. The salient feature is that, increasing the value of λ , agent 2 decreases its region of dominance. In the extreme case $\lambda = 1$, agent 2 gives full weight to model 1. This obviously favors agent 1, whose Bayesian learning ensures that it asymptotically converges to the most accurate model.

Next, we investigate the more general case $\lambda_1 + \lambda_2 < 1$. In this case we evaluate selection results observing the sign of

$$\hat{\Delta} = \hat{D}_p(p_2) - \overline{D}_p(p_1, \sigma), \quad (18)$$

where $\hat{D}_p(p_2)$ is as in (16) and $\overline{D}_p(p_1, \sigma)$ is as described in advance. In particular, for our numerical exercises we set $T = 250000$ and $R = 5$.⁵ The selection outcomes based on the sign of $\hat{\Delta}$ are reported in Figure 3. The value of π and

⁵Such a low number of independent replications is due to the extremely low volatility of the estimated average relative entropy across different replicas of 250000 time steps each. In the different replica, we obtain standard errors in order of 10^{-5} . As expected from an ergodic process, a single sufficiently long sequence is basically enough to obtain reliable estimates (Vandin et al., 2020).

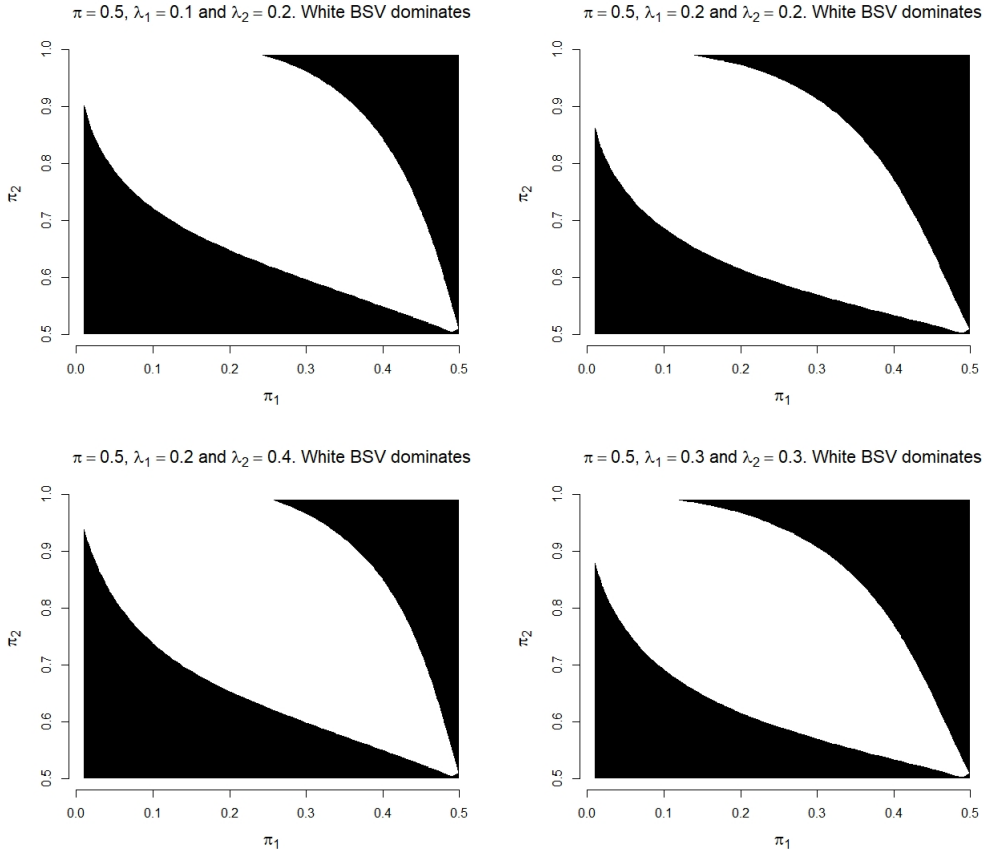


Figure 3: Survival outcomes for several combinations of (π_1, π_2) , for $\pi = 0.5$, and for the same values of λ_1 and λ_2 used in Figure 2 of Barberis et al. (1998). White: BSV dominates ($\hat{\Delta} < 0$), Black: Bayes dominates ($\hat{\Delta} > 0$).

the combinations of λ_1 and λ_2 have been chosen to match those used in Figure 2 of BSV. As one can notice, even with $\lambda_1 + \lambda_2 < 1$, for $\pi = 0.5$ agent 2 tends to dominate over those regions of the parameter space in which $\pi_2 \simeq 1 - \pi_1$. Those regions are broadly consistent (but not identical) with those that BSV indicate as delivering under-reaction and over-reaction.⁶ In Figure 4 we repeat the same exercise of Figure 3 setting $\pi = 0.25$. Even in this case there exist regions of the parameter space where agent 2 dominates over agent 1 and vice-versa. Except

⁶This statement has to be taken with a pinch of salt. The pricing model of BSV is not exactly consistent with the Arrow-Debreu economy we consider here. See Antico et al. (2022) for an investigation of the occurrence of under-reaction and over-reaction in an Arrow-Debreu economy where agents update their beliefs as in (9). Comparing the combinations of values generating both under-reaction and over-reaction the authors report with the selection outcomes presented here, we notice, again, an overall accordance characterized by imperfect superposition.

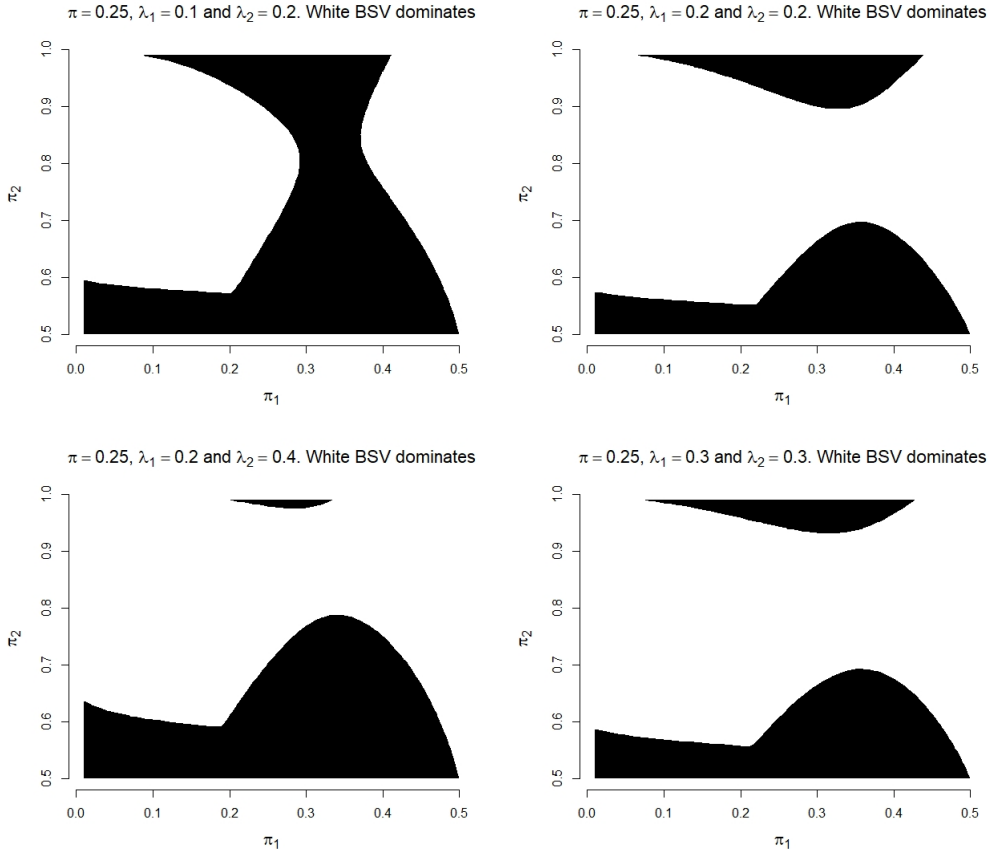


Figure 4: Survival outcomes for several combinations of (π_1, π_2) , for $\pi = 0.25$, and for the same values of λ_1 and λ_2 used in Figure 2 of Barberis et al. (1998). White: BSV dominates ($\hat{\Delta} < 0$), Black: Bayes dominates ($\hat{\Delta} > 0$).

for $\pi_1 \simeq 0.5$, low values of π_2 appear detrimental for agent 2 and beneficial for agent 1. Moreover, values of π_1 roughly included in the interval $(0.2, 0.45)$ seem to favor agent 1 over agent 2. This shows that evenly spread transition probabilities around the truth favor agent 2 only when the two states of nature are (more or less) equally probable. Indeed, what generically emerges is that ranking the level of accuracy of the two agents is difficult because of model misspecification. Both the behavioral parameters λ_1, λ_2 and the true probability π are important in deciding the selection outcome that is eventually reached in the economy.

Intuitively, the BSV learning algorithm has an advantage when models are sufficiently misspecified and by taking a convex combination of them one gets close to the truth. Indeed, when the convex combinations of models provides an advantage in terms of accuracy, the weight updating derived by BSV has an

evolutionary edge with respect to other more selective learning processes, such as the Bayesian one. This intuition is confirmed by observing that BSV dominates in the upper-left corner of every plot we presented ($\pi_1 \simeq 0$ and $\pi_2 \simeq 1$). Along the same lines, Bayesian learning has an evolutionary advantage over BSV when one of the two models is much better than the other and combining the two models one does not achieve an improvement in accuracy. This is true, for instance, in the bottom-left and top-right corners of the plots where $\pi = 0.5$.

5 Conclusions

The limited-rationality learning model proposed by BSV and adapted to an Arrow-Debreu economy by Antico et al. (2022) seems to display a good degree of evolutionary stability when compared to a Bayesian learner in the presence of model misspecification. The reason is that while the agent following the Bayes prescription asymptotically obtains its prediction from a single (wrong) model, the limited-rationality agent persistently uses a convex combination of more models. In this way, it obtains less extreme predictions and, consequently, makes more balanced portfolio choices, ultimately increasing its probability to survive and dominate. In this respect, the learning model derived from BSV is similar to other non-Bayesian learning models explored in the literature (Epstein et al., 2010; Massari, 2020) and shares many of their properties. A question that remains to be investigated is the exact impact that the requirement of being evolutionary stable has on the price misvaluation the original model was actually designed to replicate. This can be done by a thorough analysis of the intersection between the survival areas presented here and the combinations of parameter values that let the misspricing patterns emerge in Antico et al. (2022). Moreover, one might wonder if the same capability is shared by other non-Bayesian learning behaviors and what the aggregate effect of the presence of more than one of those behaviors in the market might be.

6 Acknowledgments

We thank the attendance at the “IEAP Meeting: Investor Emotions & Asset Pricing”, IAE Lille University School of Management, February 2022, and in particular Pascal Alphonse for a number of useful comments and suggestions of previous contributions in the literature.

References

- Alchian, A. (1950). Uncertainty, evolution, and economic theory. *Journal of Political Economy* 58, 211–221.
- Antico, A., G. Bottazzi, and D. Giachini (2022). A general equilibrium model of investor sentiment. LEM Working Paper 2022-15, Sant’Anna School for Advanced Studies, Pisa.
- Asness, C. S., T. J. Moskowitz, and L. H. Pedersen (2013). Value and momentum everywhere. *The Journal of Finance* 68(3), 929–985.
- Balvers, R., Y. Wu, and E. Gilliland (2000). Mean reversion across national stock markets and parametric contrarian investment strategies. *The Journal of Finance* 55(2), 745–772.
- Barberis, N., A. Shleifer, and R. Vishny (1998). A model of investor sentiment. *Journal of financial economics* 49(3), 307–343.
- Beddock, A. and E. Jouini (2021). Live fast, die young: equilibrium and survival in large economies. *Economic Theory* 71(3), 961–996.
- Benos, A. V. (1998). Aggressiveness and survival of overconfident traders. *Journal of Financial Markets* 1(3), 353–383.
- Blume, L. and D. Easley (1992). Evolution and market behavior. *Journal of Economic Theory* 58, 9–40.
- Blume, L. and D. Easley (2006). If you are so smart why aren’t you rich? Belief selection in complete and incomplete markets. *Econometrica* 74, 929–966.
- Blume, L. and D. Easley (2009). The market organism: Long-run survival in markets with heterogenous traders. *Journal of Economic Dynamics and Control* 33, 1023–1035.
- Bottazzi, G., P. Dindo, and D. Giachini (2018). Long-run heterogeneity in an exchange economy with fixed-mix traders. *Economic Theory* 66(2), 407–447.
- Bottazzi, G., P. Dindo, and D. Giachini (2019). Momentum and reversal in financial markets with persistent heterogeneity. *Annals of Finance* 15(4), 455–487.
- Bottazzi, G. and D. Giachini (2017). Wealth and price distribution by diffusive approximation in a repeated prediction market. *Physica A: Statistical Mechanics and its Applications* 471, 473–479.

- Bottazzi, G. and D. Giachini (2019a). Betting, selection, and luck: A long-run analysis of repeated betting markets. *Entropy* 21(6), 585.
- Bottazzi, G. and D. Giachini (2019b). Far from the madding crowd: Collective wisdom in prediction markets. *Quantitative Finance* 19(9), 1461–1471.
- Bottazzi, G. and D. Giachini (2022). Strategically biased learning in market interactions. Technical report, Laboratory of Economics and Management (LEM), Sant’Anna School of Advanced Studies.
- Chui, A. C., S. Titman, and K. J. Wei (2010). Individualism and momentum around the world. *The Journal of Finance* 65(1), 361–392.
- Daniel, K. and D. Hirshleifer (2015). Overconfident investors, predictable returns, and excessive trading. *Journal of Economic Perspectives* 29(4), 61–88.
- Daniel, K., D. Hirshleifer, and A. Subrahmanyam (1998). Investor psychology and security market under- and overreactions. *the Journal of Finance* 53(6), 1839–1885.
- De Bondt, W. F. and R. Thaler (1985). Does the stock market overreact? *The Journal of finance* 40(3), 793–805.
- De Bondt, W. F. and R. Thaler (1987). Further evidence on investor overreaction and stock market seasonality. *The Journal of Finance* 42(3), 557–581.
- Dindo, P. and F. Massari (2020). The wisdom of the crowd in dynamic economies. *Theoretical Economics* 15(4), 1627–1668.
- Edwards, W. (1982). Conservatism in human information processing. In D. Kahneman, P. Slovic, and A. Tversky (Eds.), *Judgment under Uncertainty: Heuristics and Biases*, pp. 359–369. Cambridge University Press.
- Epstein, L. G., J. Noor, and A. Sandroni (2010). Non-bayesian learning. *The BE Journal of Theoretical Economics* 10(1), 1–20. Article 3.
- Evstigneev, I., T. Hens, and K. Schenk-Hoppé (2009). Evolutionary finance. In T. Hens and K. Schenk-Hoppé (Eds.), *Handbook of Financial Markets: Dynamics and Evolution*. North-Holland (Handbooks in Economics Series).
- Evstigneev, I., T. Hens, and K. R. Schenk-Hoppé (2016). Evolutionary behavioral finance. In *The handbook of post crisis financial modeling*, pp. 214–234. Springer.
- Fama, E. (1970). Efficient capital markets: a review of theory and empirical work. *Journal of Finance* 25, 383–422.

- Friedman, M. (1953). *Essays in Positive Economics*. Univ. Chicago Press.
- Griffin, J. M., X. Ji, and J. S. Martin (2005). Global momentum strategies. *The Journal of Portfolio Management* 31(2), 23–39.
- Gropp, J. (2004). Mean reversion of industry stock returns in the us, 1926–1998. *Journal of Empirical Finance* 11(4), 537–551.
- Hirshleifer, D. (2015). Behavioral finance. *Annual Review of Financial Economics* 7, 133–159.
- Hong, H. and J. C. Stein (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *The Journal of Finance* 54(6), 2143–2184.
- Jegadeesh, N. and S. Titman (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance* 48(1), 65–91.
- Jegadeesh, N. and S. Titman (2001). Profitability of momentum strategies: An evaluation of alternative explanations. *The Journal of Finance* 56(2), 699–720.
- Jouini, E. and C. Napp (2011). Unbiased disagreement in financial markets, waves of pessimism and the risk-return trade-off. *Review of Finance* 15(3), 575–601.
- Kogan, L., S. A. Ross, J. Wang, and M. M. Westerfield (2006). The price impact and survival of irrational traders. *The Journal of Finance* 61(1), 195–229.
- Kogan, L., S. A. Ross, J. Wang, and M. M. Westerfield (2017). Market selection. *Journal of Economic Theory* 168, 209–236.
- Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, 1315–1335.
- Kyle, A. S. and F. A. Wang (1997). Speculation duopoly with agreement to disagree: Can overconfidence survive the market test? *The Journal of Finance* 52(5), 2073–2090.
- Luo, J., A. Subrahmanyam, and S. Titman (2021). Momentum and reversals when overconfident investors underestimate their competition. *The Review of Financial Studies* 34(1), 351–393.
- Massari, F. (2020). Under-reaction: Irrational behavior or robust response to model misspecification? Technical report, Bocconi University.
- Moskowitz, T. J. and M. Grinblatt (1999). Do industries explain momentum? *The Journal of Finance* 54(4), 1249–1290.

- Mukherji, S. (2011). Are stock returns still mean-reverting? *Review of Financial Economics* 20(1), 22–27.
- Rouwenhorst, K. G. (1998). International momentum strategies. *The Journal of Finance* 53(1), 267–284.
- Sandroni, A. (2000). Do markets favor agents able to make accurate predictions. *Econometrica* 68(6), 1303–1341.
- Sandroni, A. (2005). Efficient markets and bayes' rule. *Economic Theory* 26(4), 741–764.
- Tversky, A. and D. Kahneman (1974). Judgment under uncertainty: Heuristics and biases. *Science* 185, 1124–1131.
- Vandin, A., D. Giachini, F. Lamperti, and F. Chiaromonte (2020). Automated and distributed statistical analysis of economic agent-based models. Technical report, Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies.