



Market selection and learning under model misspecification

Giulio Bottazzi^a, Daniele Giachini^{a,*}, Matteo Ottaviani^{b,c}

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

^b German Centre for Higher Education Research and Science Studies, Schützenstraße 6a, 10117 Berlin, Germany

^c Scuola Normale Superiore, Piazza dei Cavalieri 7, 56126 Pisa, Italy

ARTICLE INFO

JEL classification:

C60
D53
D81
D83
G11
G12

Keywords:

Learning
Market selection
Model misspecification
Financial markets

ABSTRACT

This paper studies market selection in an Arrow-Debreu economy with complete markets where agents learn over misspecified models. In this setting, standard Bayesian learning loses its formal justification and biased learning processes may provide a selection advantage. Studying two cases of model misspecification and four learning processes, our analysis reveals that, differently from correctly specified settings, the ecology of traders populating the market crucially affects selection dynamics and, thus, long-run asset valuation. In fact, model misspecification implies a general difficulty in ranking learning behaviors with respect to their survival prospects. For instance, prediction averaging shows an advantage when the true data generating process belongs to the same family of models that agents use to learn. This advantage partially disappears when the true model belongs to a more general class, as a trade off emerges between approximating the projection of the true model on the space on which the agents learn and adapting to the part of the true model that cannot be represented in that space. Rules that guarantee survival are possible, but they exploit imitative mechanisms that require information about all the other market participants.

1. Introduction

In pure exchange Arrow-Debreu economies with complete markets and bounded endowments, where traders with homogeneous discount factors update their beliefs learning over a correctly specified set of models that includes the state of nature process, only those who dynamically incorporate evidence into their probabilistic predictions according to Bayes' rule are able to survive and influence assets' long-run evaluation (Blume and Easley, 2006, 2009b). In such a context, a Bayesian agent asymptotically learns the true model, drives those who persistently forecast differently out of the market, and prices assets as in a representative agent model with rational expectations. Hence, the ecology of traders and the selection dynamics generated by their competition do not play any role for asset valuation in the long run.

When, in contrast, traders face model misspecification and the true data generating process does not belong to the set on which they learn, one of the assumptions that underpins Bayes' theorem for computing conditional probabilities is missing and Bayesian learning loses its formal justification (see the discussion in Massari, 2021). In this case, little is known about the consequences on selection outcomes. In fact, previous studies show that Bayesian learning may lose its evolutionary advantage. Massari (2020) proves that an underreacting trader, assigning more weight to the prior than what Bayesian learning would prescribe, does not vanish and,

* Corresponding author.

E-mail address: daniele.giachini@santannapisa.it (D. Giachini).

<https://doi.org/10.1016/j.jedc.2023.104739>

Received 8 May 2023; Received in revised form 2 August 2023; Accepted 6 September 2023

Available online 12 September 2023

0165-1889/© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

under some circumstances, drives a Bayesian agent out of the market. Antico et al. (2023) show that a sentiment investor following the learning model of Barberis et al. (1998) can dominate a Bayesian trader over generic regions of the parameter space. These results suggest that the introduction of an ecology of different learning rules may actually generate nontrivial selection results and, as a consequence, have a nontrivial effect on prices. This observation raises several questions that remain unaddressed in the market selection literature. For example, consider an ecology of traders characterized by heterogeneous learning rules, how are the emerging selection outcomes influenced by model misspecification? Can particular survival learning mechanisms be identified? How general are they?

In recent years, the economic consequences of model misspecification have attracted increasing attention (see, for instance, the blooming contributions in the game and decision theoretic literature, e.g., Hansen, 2014; Esponda and Pouzo, 2016; Fudenberg et al., 2017; Marinacci and Massari, 2019; Cerreia-Vioglio et al., 2020; Hansen and Sargent, 2022). The basic motivation is to try to embed in economic modeling the idea that “*all models are wrong*” (Box, 1976) and that “*the very word “model” implies simplification and idealization*” (Cox, 1995). Real traders face a *large world* situation (Gigerenzer and Gaissmaier, 2011) and usually must deal with misspecified learning problems.

This paper investigates how model misspecification affects selection outcomes in pure-exchange Arrow-Debreu economies with complete markets and bounded endowments. To do that, we consider four different learning processes and two cases of model misspecification.

The learning processes we consider include Bayesian learning, underreaction, moving average over a reference learning process, and limited memory Bayesian learning. The first two processes match those considered by Massari (2020). While Bayesian learning is a natural choice, underreaction was originally introduced and studied assuming correct specification by Epstein et al. (2010). The results of Massari (2020) in a context of model misspecification provide a rationale for our choice. Moreover, the mathematical formulation of underreaction is equivalent to the *Soft-Bayes* algorithm of Orseau et al. (2017), matches the dynamics of prices and wealth in the prediction market model of Bottazzi and Giachini (2017, 2019b), and describes how risk neutral probabilities and consumption shares evolve in the economy analyzed by Dindo and Massari (2020). The other two learning processes represent extreme benchmarks. On one end of the spectrum, the moving average learning process consists in building conditional probabilities by averaging the last predictions of a reference learning process. Referencing underreaction, for instance, one can strengthen the smoothing behavior that underreaction already displays. On the other end of the spectrum, limited memory Bayesian learning deliberately forgets past observations, continuously resetting Bayesian updating and, thus, causing a sort of overreaction to recent observations.

The first case of model misspecification that we consider, parametric misspecification, consists in assuming that the true probability measure belongs to the same class of probabilistic models that agents use to learn, but with different parameter values. This is the smallest possible deviation from correct specification. The second case of model misspecification, structural misspecification, assumes that the process driving the states of nature belongs to a different and more general class of processes than the one on which the agents learn. We do this to investigate whether the intuitions and the insights obtained under parametric misspecification extend to a more general scenario or not.

Our analysis bridges the general equilibrium literature with intertemporal utility maximization and complete markets (see e.g. Sandroni, 2000; Blume and Easley, 2006, 2009a; Jouini and Napp, 2011; Kogan et al., 2006, 2017; Massari, 2017; Dindo and Massari, 2020; Beddock and Jouini, 2021; Bottazzi and Giachini, 2022) with temporary equilibrium models, based on bounded rationality and evolutionary dynamics among investment rules (see e.g. Hens and Schenk-Hoppé, 2005; Evstigneev et al., 2009, 2016; Holtfort, 2019; Bottazzi and Dindo, 2013, 2014; Bottazzi et al., 2018, 2019; Bottazzi and Giachini, 2017, 2019b,a; Elmiger, 2020).¹ Specifically, we combine the complete market Arrow-Debreu economy, characterizing most of the contributions belonging to the first strand of literature, with biased learning schemes, which are closer to the second. We do this to avoid compensation effects between nonoptimality in investment rules and misspecification in beliefs (see the discussion in Bottazzi et al., 2018; Giachini, 2021).

We prove several accuracy results on the learning processes considered and extend previous contributions, connecting long-run outcomes with beliefs' accuracy in a context in which the existence of specific limits cannot be automatically assumed. In general, we show that when model misspecification is considered, the ecology of learning behaviors operating in the market and the nature of their misspecification are of crucial importance in determining the outcome of the selection and, as a consequence, the value of assets in the long run. For example, while prediction smoothing, as prescribed by underreaction and moving average learning, generates a generic selection advantage under parametric misspecification, this advantage may break down when structural misspecification occurs. The latter induces a trade-off between approximating the projection of the true model on the space on which the agents learn and adapting to the part of the true model that cannot be represented in that space. In general, learning models inducing any kind of convergence toward a single, best misspecified model do not operate efficiently in this framework. Finally, we discuss some examples of learning rules that can survive regardless of the form of model misspecification. Their common feature is exploiting information about other agents to asymptotically adapt and imitate the best learning processes present in the market. This further confirms the huge impact that the specific ecology of learning rules and the type of model misspecification have on asset pricing.

¹ The two approaches can lead to identical dynamics. In fact, they can be linked by means of *effective beliefs*, see Bottazzi et al. (2018), Dindo (2019), Giachini (2021).

2. The model

Consider an Arrow-Debreu economy with an infinite horizon and discrete time $t = 0, 1, \dots$. There is a homogeneous consumption good and the market is complete. Let $s_t \in \{1, 2, \dots, S\}$ be the state realized at time $t > 0$, $\sigma = (s_1, s_2, \dots, s_t, \dots)$ a path, and $\sigma_t = (s_1, s_2, \dots, s_t)$ a partial history until time t .

The set of all possible paths is Σ while Σ_t indicates the set of all partial histories until time t . Let $\mathcal{C}(\sigma_t) = \{\sigma \in \Sigma | \sigma = (\sigma_t, \dots)\}$ be the cylinder with base σ_t and \mathcal{F}_t the σ -algebra generated by the cylinders $\mathcal{C}(\sigma_t)$. Then, by construction, $(\mathcal{F}_t)_{t=0}^\infty$ is a filtration and \mathcal{F} is the σ -algebra generated by the union of filtrations. We indicate by p the true probability measure on (Σ, \mathcal{F}, p) is a well-defined probability space. We assume that any partial history has a positive probability of being realized, $p(\sigma_t) > 0, \forall \sigma_t$. Expectation is denoted with E and, when there is no subscript or superscript, it is computed with respect to p .

The economy is populated by N agents indexed by $i = 1, 2, \dots, N$. Every agent i is endowed with a stream of non-zero and uniformly bounded consumption good for any path $\sigma, (e_i(\sigma_t))_{t=0}^\infty$.

Denote by p_i the subjective probability of agent i over (Σ, \mathcal{F}) , by $p_i(s_t | \sigma_{t-1})$ the subjective conditional probability of state s_t after a partial history σ_{t-1} , and by $p_i(\sigma_t) = \prod_{\tau=1}^t p_i(s_\tau | \sigma_{\tau-1})$ the subjective likelihood of partial history σ_t . Agent i chooses its consumption plan $(c_i(\sigma_t))_{t=0}^\infty$ solving

$$\max_{\{c_i(\sigma_t), \forall t, \sigma\}} E_{p_i} \left[\sum_{t=0}^\infty \beta_i^t u_i(c_i(\sigma_t)) \right] \text{ s.t. } \sum_{t=0}^\infty \sum_{\sigma_t \in \Sigma_t} q(\sigma_t) (e_i(\sigma_t) - c_i(\sigma_t)) \geq 0,$$

where $\beta_i \in (0, 1)$ is agent i 's discount factor, u_i is the Bernoulli utility of consumption of agent i , and $q(\sigma_t)$ is the price of the Arrow-Debreu security paying one if partial history σ_t is realized and zero otherwise. We also assume that individual probabilities p_i are absolutely continuous with respect to p and that Bernoulli utilities are continuously differentiable, increasing, strictly concave, and satisfy the Inada condition at zero. With these hypotheses, there is a unique competitive equilibrium and $\forall \sigma_t, q(\sigma_t) > 0, \sum_{i=1}^N c_i(\sigma_t) = \sum_{i=1}^N e_i(\sigma_t) = e(\sigma_t)$. From the F.O.C. of the optimal consumption problem, $\forall i, j \in 1, \dots, N$ (Blume and Easley, 2006),

$$\frac{u'_i(c_i(\sigma_t))}{u'_j(c_j(\sigma_t))} = \left(\frac{\beta_j}{\beta_i} \right)^t \frac{p_j(\sigma_t)}{p_i(\sigma_t)} \frac{u'_i(c_i(\sigma_0))}{u'_j(c_j(\sigma_0))},$$

that is

$$\frac{1}{t} \log \frac{u'_i(c_i(\sigma_t))}{u'_j(c_j(\sigma_t))} = \log \frac{\beta_j}{\beta_i} + \frac{1}{t} \log \frac{p(\sigma_t)}{p_i(\sigma_t)} - \frac{1}{t} \log \frac{p(\sigma_t)}{p_j(\sigma_t)} + \frac{1}{t} \log \frac{u'_i(c_i(\sigma_0))}{u'_j(c_j(\sigma_0))}. \tag{1}$$

To describe the selection dynamics taking place in this competitive equilibrium, we introduce the following.

Definition 2.1. An agent i *vanishes* if, p -almost surely, $\lim_{t \rightarrow \infty} c_i(\sigma_t) = 0$. It *survives* if it does not vanish.

The study of the asymptotic dynamics of agents' relative consumption can be reduced to the analysis of their individual probability measures and discount factors. Consider the uniform distance of the logarithm of the subjective conditional probability of agent i from the logarithm of the true conditional probability,

$$\| \log p(\cdot | \sigma_t) / p_i(\cdot | \sigma_t) \|_\infty = \max_{s \in S} | \log p(s | \sigma_t) / p_i(s | \sigma_t) |,$$

and the relative entropy of conditional probabilities and its partial average,

$$D_{p|p_i}(\sigma_t) = \sum_{s=1}^S p(s | \sigma_t) \log \frac{p(s | \sigma_t)}{p_i(s | \sigma_t)} \text{ and } \bar{D}_{p|p_i}(\sigma_t) = \frac{1}{t+1} \sum_{\tau=0}^t D_{p|p_i}(\sigma_\tau). \tag{2}$$

Under the hypothesis that the uniform distance of the logarithm of any individual conditional probability from the logarithm of the true conditional probability is bounded, the following proposition can be used to investigate the asymptotic behavior of (1).

Proposition 2.1. Given two agents i and j , assume that $\exists L > 0$ such that, p -almost surely, $\| \log p(\cdot | \sigma_t) / p_h(\cdot | \sigma_t) \|_\infty < L, h = i, j$. Then, $\forall \alpha < 1/2, p$ -almost surely, for large t ,

$$\frac{1}{t} \log \frac{u'_i(c_i(\sigma_t))}{u'_j(c_j(\sigma_t))} = \left(\log \beta_j - \bar{D}_{p|p_j}(\sigma_{t-1}) \right) - \left(\log \beta_i - \bar{D}_{p|p_i}(\sigma_{t-1}) \right) + o(t^{-\alpha}).$$

Moreover if, p -almost surely,

$$\log \beta_j - \log \beta_i + \liminf_{t \rightarrow \infty} \left(\bar{D}_{p|p_i}(\sigma_t) - \bar{D}_{p|p_j}(\sigma_t) \right) > 0,$$

then agent i vanishes.

Proof. See Section A.1. \square

Proposition 2.1 encompasses and extends previous results in the literature (Sandroni, 2000; Blume and Easley, 2006). It connects the analysis of agents’ relative consumption to the analysis of their beliefs, as expressed by their individual measures. The necessity of a bounded uniform distance is often fulfilled by assuming that the conditional probabilities of both the true and individual measures are uniformly bounded away from zero. Moreover, the literature often assumes models for which the asymptotic limit of the average relative entropy of individual conditional probabilities, $\overline{D}_{p|p_i}(\sigma) = \lim_{t \rightarrow \infty} \overline{D}_{p|p_i}(\sigma_t)$, exists and is p -almost surely constant. This makes the previous proposition stronger and relatively easier to apply. In the following analysis, we will retain the first assumption, using measures that are uniformly bounded away from zero. Instead, the asymptotic convergence of the average relative entropy of individual conditional probabilities cannot generally be assumed when the analysis is extended to a context of model misspecification.

3. Learning processes

This section describes the four learning processes that we will analyze in this paper: *Bayesian learning*, *learning with underreaction*, *moving average learning* of an underlying process, and *limited memory Bayesian learning*. They are designed to show different degrees of sophistication, mechanisms, and asymptotic dynamics. They all share the following common structure.

Consider K i.i.d. measures whose conditional probabilities are the vectors π_1, \dots, π_K that belong to the topological interior of the $(S - 1)$ -simplex, $\pi_k = (\pi_k(1), \pi_k(2), \dots, \pi_k(S)) \in \Delta_+^{S-1}$. These vectors are uniformly bounded away from zero and diverse, that is, $\exists \epsilon, d, \pi > 0$ such that $\pi_k(s) > \epsilon$ and $\|\pi_k - \pi_h\| > d, \forall s, k, h$. To simplify our investigation, we assume the following.

Assumption 1. The individual conditional probabilities of the agents belong to the convex hull H_K generated by the conditional probabilities of the K models, $\forall \sigma_t$

$$(p_i(1 | \sigma_t), \dots, p_i(S | \sigma_t)) \in H_K = \left\{ \sum_{k=1}^K \eta_k \pi_k \mid \sum_{k=1}^K \eta_k = 1, \eta_k \geq 0 \right\} \subseteq \Delta_+^{S-1}.$$

Moreover, $\exists L > 0$ such that, $\forall k$ and $\forall \sigma_t$, $\|\log p(\cdot | \sigma_t) / \pi_k(\cdot)\|_\infty < L$.

Let $w_{i,k}(\sigma_t)$ be the weight agent i attaches to model k after having observed the partial history σ_t . Then, agents’ individual conditional probabilities read, $\forall s$,

$$p_i(s | \sigma_t) = \sum_{k=1}^K w_{i,k}(\sigma_t) \pi_k(s), \text{ with } w_{i,k}(\sigma_t) \geq 0, \forall k, \text{ and } \sum_{k=1}^K w_{i,k}(\sigma_t) = 1. \tag{3}$$

Learning processes differ on how they compute the weights. By Assumption 1, the first requirement of Proposition 2.1 is verified, and the quantities in (2) are bounded.

Bayesian learning. Define $\pi_k(\sigma_t) = \prod_{\tau=1}^t \pi_k(s_\tau)$. Then, in Bayesian learning, weights are updated according to Bayes’ rule,

$$w_{i,k}(\sigma_t) = \frac{\pi_k(s_t) w_{i,k}(\sigma_{t-1})}{p_i(s_t | \sigma_{t-1})} = \frac{\pi_k(\sigma_t)}{p_i(\sigma_t)} w_{i,k}(\sigma_0) \quad \forall k, t, \sigma. \tag{4}$$

The weight $w_{i,k}(\sigma_t)$ is the probability a Bayesian agent attaches to the event “model k is the true one” conditional upon the observation of partial history σ_t . By substituting (4) in $p_i(\sigma_t)$, one obtains $p_i(\sigma_t) = \sum_{k=1}^K w_{i,k}(\sigma_0) \pi_k(\sigma_t)$. Bayesian learning can be considered the cornerstone of online learning. In fact, given a correct model specification, it is guaranteed to converge to the truth. This is generally not true when misspecified models are considered. However, Bayesian learning can retain its role as a benchmark due to the following result.

Proposition 3.1. Define $\pi^*(\sigma_t) = \max_{k \in \{1, \dots, K\}} \{\pi_k(\sigma_t)\}$. For any Bayesian agent i and $\forall \alpha < 1/2$, p -almost surely, for large t ,

$$\overline{D}_{p|p_i}(\sigma_{t-1}) - \overline{D}_{p|\pi^*(\sigma_t)}(\sigma_{t-1}) = o(t^{-\alpha}).$$

Proof. See Section A.2. \square

The model $\pi^*(\sigma_t)$ can be considered the best ex post model among those available to the Bayesian learner. In general, it depends on the specific realization σ_t . Clearly, the knowledge of which model is going to be the best is not available ex ante. However, the Bayesian learning algorithm is able to converge to it in terms of average relative entropy. Proposition 3.1 also provides a lower bound to the speed of convergence.

Learning with underreaction. This learning protocol can be considered a form of “moderate” Bayesian learning obtained by combining Bayes update and prior probability (Epstein et al., 2010; Massari, 2020). The updating rule (4) is replaced by

$$w_{i,k}(\sigma_t) = \lambda_i w_{i,k}(\sigma_{t-1}) + (1 - \lambda_i) \frac{\pi_k(s_t) w_{i,k}(\sigma_{t-1})}{p_i(s_t | \sigma_{t-1})} \quad \forall k, t, \sigma, \tag{5}$$

with $\lambda_i \in [0, 1)$. Conditional probabilities after a partial history σ_t can be seen as the convex combination of the conditional probabilities agent i would build after the partial history σ_{t-1} and the Bayesian conditional probabilities obtained after observing s_t given a prior $w_i(\sigma_{t-1}) = (w_{i,1}(\sigma_{t-1}), \dots, w_{i,K}(\sigma_{t-1}))$ (Epstein et al., 2010; Giachini, 2021). Setting $\lambda_i = 0$, Bayesian learning is recovered. The next result clarifies why underreaction represents a robust learning strategy in the case of model misspecification.

Proposition 3.2. For any underreacting agent i and $\forall \alpha < 1/2$, it is p -almost surely, for large t ,

$$\overline{D}_{p|p_i}(\sigma_{t-1}) - \overline{D}_{p|\pi_k}(\sigma_{t-1}) \leq o(t^{-\alpha}), \forall k \in \{1, 2, \dots, K\}.$$

Proof. See Section A.3. \square

It follows that, for t sufficiently large, an underreacting agent is at least as accurate as its best ex post model, $\pi^*(\sigma_t)$, but possibly more accurate. Note that Proposition 3.2, such as Proposition 3.1, does not imply or require the asymptotic convergence to a single i.i.d. model. Comparing Proposition 3.1 and Proposition 3.2, it is clear that, for large t , an underreacting agent cannot have a higher average relative entropy than a Bayesian trader. The slower update of beliefs characterizing underreaction confers it a *specific* advantage over Bayesian learning. In fact, Massari (2020) shows that, under discount factor homogeneity, an underreacting agent maintains a positive consumption share along any path when competing against a Bayesian trader.

Moving average learning. Moving average learning represents a further layer of smoothing over the conditional probabilities of the underlying process. It consists in taking a reference learning process p^* and applying a moving average to the sequence of probabilistic predictions generated for every state. Assume agent i adopts a moving average learning with memory M_i , then $\forall \sigma_t$ and for any s

$$p_i(s | \sigma_t) = \begin{cases} p^*(s | \sigma_t) & \text{if } t < M_i - 1, \\ M_i^{-1} \sum_{m=1}^{M_i} p^*(s | \sigma_{t-m+1}) & \text{if } t \geq M_i - 1. \end{cases} \tag{6}$$

If the underlying learning process follows Assumption 1, the same thing can be restated in terms of weights with $w_{i,k}(\sigma_t) = M_i^{-1} \sum_{m=1}^{M_i} w_k^*(\sigma_{t-m+1})$ if $t \geq M_i - 1$.

Limited memory Bayesian learning. The limited memory Bayesian learning is a version of the standard Bayesian learning process in which the agent deliberately forgets observations in the past. Here, we consider the version with the shortest possible memory, that is, a memory of one. In any period t , agent i is forgetting all the sequence of states occurred until $t - 2$ (included) and restarts its Bayesian learning procedure considering the previous state and the initial prior weights,

$$w_{i,k}(\sigma_t) = w_{i,k}(s_t) = \frac{\pi_k(s_t)w_{i,k}(\sigma_0)}{\sum_{k'=1}^K \pi_{k'}(s_t)w_{i,k'}(\sigma_0)} \quad \forall k, t, \sigma. \tag{7}$$

Because the models on which the agent learns are i.i.d., the limited memory Bayesian learning has a Markov structure. Note that this process strongly depends on the initial assignment of weights.

4. Misspecified models

The learning problem is correctly specified if the true probability p belongs to the set of models on which the agents learn. That is, if the true process p is i.i.d. with conditional probabilities equal to one of the K vectors that define the convex hull in Assumption 1. In this case, based on the discussion in the previous section, the predictions of Bayesian and underreacting learners, as well as any moving-average learning built on top of these, will converge to the true model. If those agents share a homogeneous utility discount factor, they all survive. If instead they have different utility discount factors, only the agent with the highest discount factor survives. The limited memory Bayesian learner does not converge to the truth, and it is at a disadvantage. If not rescued by a higher discount factor, it will vanish. In general, this simple picture is no longer true under model misspecification.

In the following, we study two specific cases of misspecification. We start with the case of an i.i.d. true measure that does not belong to the set of models the agents can learn. In this case, the i.i.d. models on which the agents learn belong to the same class of the true measure, but their parameters are generically not correct. We call this case *parametric misspecification*. In the second case, *structural misspecification*, we assume that the true measure is Markov. Here, the models employed by the agents belong to a different and less general class than the truth. In both cases, we start by providing some clues about the relative performance of the models based on their general behavior, and then present some numerical exercises to further illustrate their properties.

4.1. Parametric misspecification

We assume that states of nature follow an i.i.d. process that is different from those on which the agents learn. Formally,

Assumption 2. The true measure p is an i.i.d. process whose conditional distribution is described by the vector $\pi = (\pi(1), \pi(2), \dots, \pi(S)) \in \Delta_+^{S-1}$, such that $p(s_t | \sigma_{t-1}) = \pi(s_t)$ and, $\forall k \in \{1, 2, \dots, K\}$, $\|\pi_k - \pi\| > 0$.

For each model k used by agents, define $D_{\pi|\pi_k} = \sum_{s=1}^S \pi(s) \log \pi(s) / \pi_k(s) > 0$. If there is a single best model $k^* = \operatorname{argmin}_k \{D_{\pi|\pi_k}\}$, then for a Bayesian agent i , $\lim_{t \rightarrow \infty} w_{i,k^*}(\sigma_t) = 1$ and $\lim_{t \rightarrow \infty} \overline{D}_{p|p_i}(\sigma_t) = D_{\pi|\pi_{k^*}}$. In the non-generic situation in which there are multiple models with minimal relative entropy among those on which agents learn, then the conditional probabilities of a Bayesian trader can actually fluctuate. Inspired by Bottazzi and Giachini (2017) and Dindo and Massari (2020), a general result can be derived for the underreacting agent when the true model belongs to the convex hull of the models on which agents learn. This indicates a generic advantage for the underreacting agent.

Proposition 4.1. If $\pi \in H_K$, then for any underreacting agent i and $\forall \alpha < 1/2$, it is p -almost surely, for large t ,

$$\overline{D}_{p|p_i}(\sigma_t) \leq \frac{1 - \lambda_i}{2(\lambda_i + \epsilon)^2} + \frac{o(t^{-\alpha})}{1 - \lambda_i}.$$

Proof. See Section A.4. \square

In other words, when the truth belongs to H_K , an agent with a high level of underreaction can eventually generate extremely accurate conditional probabilities. As a consequence, it has a survival advantage over other traders. To see it, assume homogeneity in the utility discount factors and that agent 1 is under-reacting with the parameter λ_1 . Thus, any trader $i > 1$ for which it is, p -almost surely and eventually in t , $\overline{D}_{p|p_i}(\sigma_t) > (1 - \lambda_1) / (2(\lambda_1 + \epsilon)^2)$, will vanish. Analogously, if all traders, except the underreacting one, are bounded away from the truth, $\overline{D}_{p|p_i}(\sigma_t) > \delta > 0, \forall i > 1$, then the underreacting agent makes everybody else vanish if $\lambda_1 > (\sqrt{1 + 8\delta} - 1) / (4\delta)$.

Concerning moving average learning, we have the following.

Proposition 4.2. Given a moving average learning process p_i with the reference learning process p^* , consider $\sigma^2(s, \sigma_t) = \sum_{m=0}^{M-1} (p^*(s | \sigma_{t-m}) - p_i(s, \sigma_t))^2 / M$. Then, if p satisfies Assumption 2,

$$\frac{\sigma^2(\sigma_t)}{2(1 - \epsilon)} \leq \frac{1}{M} \sum_{m=0}^{M-1} D_{p|p^*}(\sigma_{t-m}) - D_{p|p_i}(\sigma_t) \leq \frac{\sigma^2(\sigma_t)}{2\epsilon}.$$

Proof. See Section A.5. \square

If the underlying process p^* converges p -almost surely to a constant conditional probability distribution, such as Bayesian learning in the presence of a single best model, then $\lim_{t \rightarrow \infty} \sigma^2(\sigma_t) = 0$ and, in the long run, the performance of the moving average process becomes identical to the one of the underlying process. Conversely, if the underlying process entails some sort of persistent fluctuation in conditionals, moving average brings a definite advantage as its relative entropy is strictly lower than the average of the relative entropy of the underlying process.

Finally, if agent i is a limited memory Bayesian learner, under Assumption 2, p -almost surely, $\overline{D}_{p|p_i}(\sigma) = \lim_{t \rightarrow \infty} \overline{D}_{p|p_i}(\sigma_t) = \sum_{s=1}^S \pi(s) D_{p|p_i}(s)$. Alternating among different convex combinations of models depending on the last realized state, the accuracy of the limited memory Bayesian agent depends upon how accurate those convex combinations are on average.

In summary, even if prediction smoothing appears as a key mechanism, the previous results do not allow us to devise any general ranking among the different learning models. Specifically, Proposition 4.1 does not imply that the accuracy of the underreacting learner increases monotonically with λ_i . At the same time, by persistently resetting the learning process, the limited memory Bayesian learner constantly mixes the misspecified i.i.d. models and never converges to a single one. This might be advantageous.

The numerical exercises proposed in the next section exemplify the difficulties in ranking and provide some new insights about the relative performances of the models matter of study.

4.1.1. Numerical exploration

We consider an economy with two possible states of the world, $S = 2$, driven by an i.i.d. true process with conditional probabilities $\pi = (\pi, 1 - \pi)$, $\pi \in (0, 1)$. Agents learn on two models (i.e. $K = 2$) with respective probabilities $\pi_1 = (\pi_1, 1 - \pi_1)$ and $\pi_2 = (\pi_2, 1 - \pi_2)$; $\pi_1, \pi_2 \in (0, 1)$, $\pi_1 < \pi_2$. Agents' initial prior is uniform, that is, $w_{i,k}(\sigma_0) = 0.5 \forall i, k$. The performances of the different learning processes are expressed in terms of their average relative entropy $\overline{D}_{p|p_i}(\sigma_t)$ and are reported in Fig. 1 as a function of the true probability π .

By Proposition 3.1, the Bayesian process (darker and thicker solid line) always converges to the best model, apart from the single point in which the two models on which the agents learn have the same relative entropy. The average relative entropy of the limited memory Bayesian model converges to the π -average of the relative entropy of its two conditional probability distributions. In both cases, the average relative entropy can be computed analytically. For the underreacting processes with different values of the parameter λ_i and the moving average processes built on them with $M_i = 10$, the reported values of $\overline{D}_{p|p_i}(\sigma_t)$ are computed as the

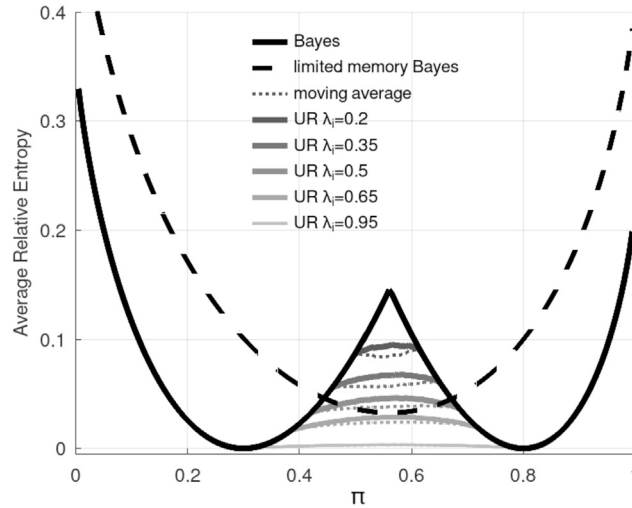


Fig. 1. Average relative entropy of the different learning models as a function of π , i.e. the true probability of the realization of state 1. Parameter settings are $\pi_1 = 0.3$, $\pi_2 = 0.8$, and, for the moving average model, $M_j = 10$. For estimated values, standard errors are on the order of 10^{-4} or smaller.

average over 10^2 independent random partial histories of length $t = 2 \times 10^4$. Over this time length, the quantities appear extremely stable across replicas, with standard errors of the order 10^{-4} , at most.²

As predicted, when the true probability matches one of the two underlying models, i.e. $\pi = \pi_1, \pi_2$, the average relative entropy of the Bayesian and of the underreacting learners are both equal to zero. They remain equal when $\pi < \pi_1$ or $\pi > \pi_2$, irrespective of the value of λ_i . On the contrary, when $\pi \in (\pi_1, \pi_2)$, the average relative entropy of the underreacting learner decreases as its degree of underreaction increases. This is due to the fact that the underreacting agents persistently and effectively use a mixture of models to build their predictions. Moving average learning processes built on top of the respective underreaction processes see improved performances when the latter ones adopt persistent model mixing, $\pi \in (\pi_1, \pi_2)$. They are identical to the underlying processes when those converge to a single i.i.d. model. Note that there is a specific interval for the values of π where the limited memory Bayesian learner outperforms the Bayesian learner. However, it succumbs to an agent showing a sufficiently high level of underreaction.

The fact that the average relative entropy of underreaction monotonically decreases in λ_i , clearly visible in Fig. 1, is a feature revealed by the numerical exercise that was not predicted by the results of the previous sections. However, the reduction in average relative entropy due to underreaction does not seem to come into play immediately when increasing λ_i . To further investigate this point, in Fig. 2 we report the average relative entropy of an underreacting agent i as a function of λ_i , for different values of π . We consider the same number and length of partial histories used in Fig. 1.

For any value of $\pi \in (\pi_1, \pi_2)$, there exists a threshold value $\underline{\lambda}_i$ such that, as λ_i increases beyond it, the monotonically decreasing behavior appears. Intuitively (see also the discussion in Massari, 2020), this should happen when the mixing coefficient λ_i is large enough for the mixture of the two models to start having a lower average entropy than the best model. Assume, without loss of generality, that $D_{\pi|\pi_2} < D_{\pi|\pi_1}$. Thus, the threshold value should solve $D_{\pi|\underline{\lambda}_i, \pi_2 + (1-\underline{\lambda}_i)\pi_1} = D_{\pi|\pi_2}$, that is, $\underline{\lambda}_i = (\bar{\pi}_1 - \pi_1)/(\pi_2 - \pi_1)$, with $\bar{\pi}_1 \in (\pi_1, \pi)$ and such that $D_{\pi|\bar{\pi}_1} = D_{\pi|\pi_2}$. This value is reported as our theoretical prediction in Fig. 2 (dashed line), and it fits the data with high accuracy.

To see how the characteristics of the learning processes shape the dynamics of consumption shares and prices, we consider the above economy populated by four agents: the first is a Bayesian (B); the second underreacts with $\lambda = 0.65$ (UR); the third is a limited memory Bayesian learner (LMB); the fourth uses the moving average learning process with memory $M = 10$ over the predictions of the underreaction learning process of the second agent (MA). We assume, $\forall i \in \mathcal{J} = \{B, UR, LMB, MA\}$, $\beta_i = \beta$, $e_i(\sigma_t) = e/4 > 0 \forall t, \sigma$, and $u_i(c) = (1 - \beta) \log(4c/e)$, so that the market shares evolve according to

$$\frac{c_i(\sigma_{t+1})}{e} = \frac{p_i(s_{t+1} | \sigma_t) c_i(\sigma_t)}{\sum_{j \in \mathcal{J}} p_j(s_{t+1} | \sigma_t) c_j(\sigma_t)}$$

The price of the Arrow-Debreu security relative to σ_t is $q(\sigma_t) = \beta^t \sum_{i \in \mathcal{J}} p_i(\sigma_t)/4$ (Bottazzi and Giachini, 2022) so that the price of a claim traded at σ_t that pays 1 at $t + 1$ if s is realized and zero otherwise is

$$q(s | \sigma_t) = \frac{q(s, \sigma_t)}{q(\sigma_t)} = \frac{\beta}{e} \sum_{i \in \mathcal{J}} p_i(s | \sigma_t) c_i(\sigma_t) \tag{8}$$

Fig. 3, top row, shows the dynamics of consumption and beliefs for $\pi = 0.6$. From Fig. 1, we know that the MA agent is the most accurate trader in this case. As expected, the consumption share of the MA agent converges to 1. Agent B is the fastest to approach

² The first 10^4 steps of each independent replication have been discarded to mitigate the possible initial condition bias.




Fig. 2. Average relative entropy of an underreacting agent as a function of λ_i and for different values of π . Parameter settings: $\pi_1 = 0.3, \pi_2 = 0.8$. The values of π have been chosen such that $\bar{D}_{\pi_1|\pi_2}(\sigma) < \bar{D}_{\pi_1|\pi_1}(\sigma)$ holds. Standard errors are in the order of 10^{-4} or smaller.

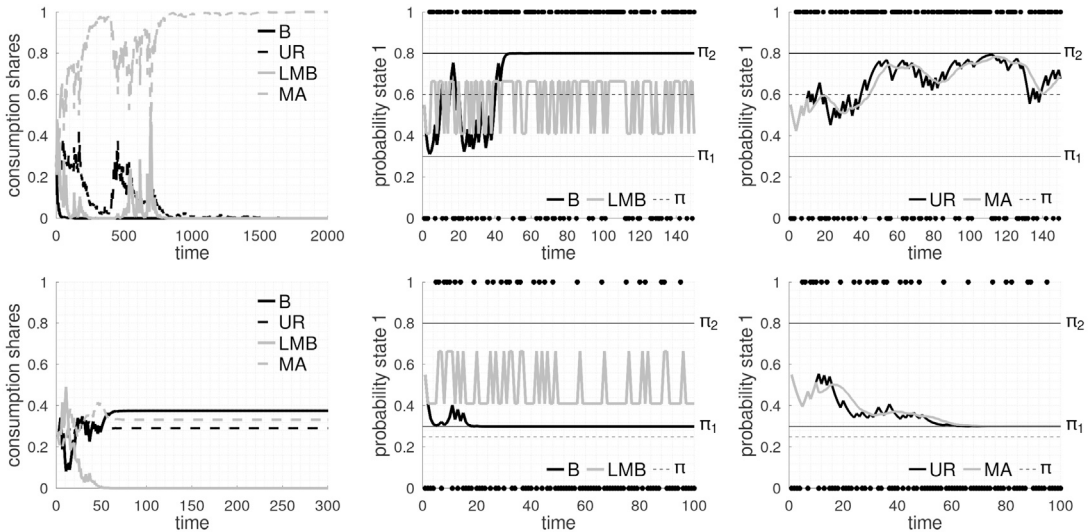


Fig. 3. **Top:** $\pi = 0.6$. **Bottom:** $\pi = 0.25$. **Left:** consumption share dynamics in a market populated by agents B, UR, LMB, MA. **Center:** conditional probability attached to state 1 by B and LMB as a function of time. **Right:** conditional probability attached to state 1 by UR and MA as a function of time. The black dots on 1 represent $s_t = 1$, those on 0 represent $s_t = 2$.

a zero consumption share, while agent UR is the slowest. The speed of convergence to zero is inversely proportional to their average relative entropy. Looking at the subjective probabilities attached to state 1 (Fig. 3, top row, center, and right panels), agent B converges to model 2 quite quickly. Instead, the UR, LMB and MA agents fluctuate persistently. However, while agents UR and MA tend to stay between the truth and the best model, displaying a rather smooth path, agent LMB displays the expected jumpy behavior. For sufficiently large t , the price of the claim in (8) is only influenced by the conditional probability of the MA agent.

Fig. 3, bottom row, shows the dynamics of consumption and beliefs for $\pi = 0.25$. In this case, agents B, UR, and MA have the same level of average relative entropy, whereas agent LMB is less accurate. Consumption shares stabilize quite quickly for the surviving agents and agent LMB vanishes. Looking at subjective probabilities, agents B, UR, and MA converge to model 1, while the predictions of the LMB agent continue to fluctuate far away from the truth, generating on average a higher relative entropy. Concerning the price

of the claim in (8), the convergence of the conditional probabilities of the surviving agents implies that it follows model 1 when t is sufficiently large.

4.2. Structural misspecification

The case of structural misspecification is defined as follows.

Assumption 3. The true measure p follows a discrete-time Markov chain with transition matrix $P: p(s_{t+1} | \sigma_t) = P_{s_t, s_{t+1}} \forall t, \sigma$ and $p(s | \sigma_0) = p_{s,0}$ with $p_{s,0} > 0 \forall s \in \{1, 2, \dots, S\}$. For any $(s, s') \in \{1, 2, \dots, S\} \times \{1, 2, \dots, S\}$, $P_{s, s'} > 0$.

The strict positiveness of the transition matrix's entries implies that the Markov chain defining the true probability measure p is irreducible and, as a consequence, the invariant probability distribution $\pi = (\pi(1), \pi(2), \dots, \pi(S))$, with $\pi(s) > 0 \forall s$, exists unique. To understand how the different learning processes perform in this case, it is useful to compute the average relative entropy of the different i.i.d. models on which the agents learn with respect to the true process.

Proposition 4.3. For any $k = 1, \dots, K$, p -almost surely,

$$\lim_{t \rightarrow \infty} \overline{D}_{p|\pi_k}(\sigma_t) = \overline{D}_{p|\pi_k}(\sigma) = \overline{D}_{\pi|\pi_k}(\sigma) + \overline{D}_{p|\pi}(\sigma), \tag{9}$$

where

$$\overline{D}_{\pi|\pi_k}(\sigma) = \sum_{s=1}^S \pi(s) \log \frac{\pi(s)}{\pi_k(s)} \text{ and } \overline{D}_{p|\pi}(\sigma) = \sum_{s'=1}^S \pi(s') \sum_{s=1}^S P_{s',s} \log \frac{P_{s',s}}{\pi(s)}.$$

Proof. See Section A.6. \square

The average relative entropy of an i.i.d. model π_k with respect to the Markov chain p is the sum of two components: the relative entropy of π_k with respect to the invariant distribution of the chain π and the average relative entropy of the invariant distribution π with respect to the transition probabilities. By Proposition 3.1, a Bayesian agent is asymptotically as accurate as the i.i.d. model with the lowest relative entropy with respect to the invariant distribution. However, it cannot do anything to prevent the information loss due to the second term on the right-hand side of (9). An underreacting agent still maintains a specific advantage over the Bayesian agent (Massari, 2020), but no generic advantage, such as that of Proposition 4.1, is present here. The intuition is that averaging different i.i.d. models may improve the prediction of the invariant distribution, but may be counterproductive when the true probabilities naturally fluctuate. The same limitation affects moving average learning.

The situation for the limited memory Bayesian learning process is the opposite. In this case, the limitation in the number of observations used by the agent makes its conditional probabilities display a Markovian behavior and can be a source of accuracy. Specifically,

$$\overline{D}_{p|p_i}(\sigma) = \sum_{s'=1}^S \pi(s') \sum_{s=1}^S P_{s',s} \log \frac{P_{s',s}}{\sum_{k=1}^K \pi_k(s) w_{i,k}(s')},$$

with $w_{i,k}(s')$ as in (7). Therefore, if the i.i.d. models and initial weights are such that the resulting conditional probabilities are close to the true transition probabilities, $\sum_{k=1}^K \pi_k(s) w_{i,k}(s') \sim P_{s',s} \forall s', s$, the limited memory process can show a high level of accuracy.

In summary, learning processes that do not provide relevant selection advantages in the parameter misspecification case may become effective as structural misspecification occurs. This will be made clearer in the numerical exercises of the next section.

4.2.1. Numerical exploration

For our numerical exercises, we consider the same settings used in subsection 4.1.1 with the exception of the true probability. That is, we set $K = S = 2$ and simplify the notation considering $P_{1,2} = 1 - P_{1,1}$ and $P_{2,2} = 1 - P_{2,1}$. The invariant distribution reads

$$\pi = \left(\frac{P_{2,1}}{1 - P_{1,1} + P_{2,1}}, \frac{1 - P_{1,1}}{1 - P_{1,1} + P_{2,1}} \right).$$

Its average relative entropy with respect to the truth is reported in the left panel of Fig. 4. When $P_{1,1} = P_{2,1}$ the true process is i.i.d. and the average relative entropy of the invariant measure is zero. The right panel of Fig. 4 reports the average relative entropy of a Bayesian agent. It is always strictly positive and progressively grows moving towards the corners $(P_{1,1}, P_{2,1}) = (0, 1)$ and $(P_{1,1}, P_{2,1}) = (1, 0)$. This is due to the second term of (9) and represents the unavoidable loss of accuracy caused by structural misspecification.

Fig. 5 reports the average relative entropy of the four learning processes with respect to the truth, removing the model-independent term $\overline{D}_{p|\pi}(\sigma)$.³ The ‘‘valleys’’ in the Bayesian case, top-left panel, correspond to an invariant distribution that exactly

³ The average relative entropy of the underreacting agent and of the moving average agent are computed numerically. See the caption for details.




Fig. 4. Average relative entropy of the invariant distribution (left) and of a Bayesian agent (right) for different combinations of $P_{1,1}$ and $P_{2,1}$. Parameter settings: $\pi_1 = 0.3, \pi_2 = 0.8$.




Fig. 5. Differences between average relative entropy of the learning process and the average relative entropy of the invariant distribution. **Top-left:** Bayesian learning. **Top-right:** underreaction with $\lambda = 0.65$. **Bottom-left:** moving average agent with $M = 20$ exploiting the underreaction with $\lambda = 0.65$. **Bottom-right:** limited memory Bayesian learning. Parameter settings: $\pi_1 = 0.3, \pi_2 = 0.8$. The average relative entropy of underreaction and moving average have been estimated over 200 independent realizations of 2500 steps each. For estimated values, standard errors are in the order of 10^{-4} or smaller. The plots are rotated of 90° clockwise with respect to Fig. 4 in order to improve the visualization of results.

matches one of the two i.i.d. models on which the agent learns. The underreaction process with $\lambda = 0.65$ and the moving average process with $M = 20$ that leverages the predictions of the underreacting agent, in the upper right and lower left panels, respectively, show an improvement in performance in the region between the two valleys, similar to what was observed for the parametric misspecification case in Fig. 2. However, a trade-off now appears between dampening conditional probability fluctuations to get closer to the best i.i.d. model and keeping changing conditional probabilities to match transition probabilities. In the parameter region where the Markov chain favors switching, $P_{1,1} \approx 0$ and $P_{2,1} \approx 1$, the moving average is more accurate than the underreacting process it exploits. Conversely, in the region where the Markov chain is more persistent, $P_{1,1} \approx 1$ and $P_{2,1} \approx 0$, underreaction can become more accurate than the invariant distribution and outperforms the moving average. Proposition 4.2 ruled out this possibility in the case of parametric misspecification. This is also the parameter region where the limited memory Bayesian process with a uniform prior, in the lower right panel, reaches the best performances in terms of accuracy.

These considerations are summarized in Fig. 6, which reports the most accurate learning models in the $(P_{1,1}, P_{2,1})$ space. Note that in the $P_{1,1} \approx 0.8$ and $P_{2,1} \approx 0.6$ region, underreaction prevails over the other learning processes. In this region, underreaction achieves the best combination between averaging to get close to the invariant and fluctuating to follow transition probabilities. In the large




Fig. 6. Most accurate processes over the $(P_{1,1}, P_{2,1})$ space. **White:** Multiple maximally accurate processes or the difference between the two lowest average relative entropy processes is not significant at $\sim 99\%$ confidence level. **Light gray:** underreaction with $\lambda = 0.65$ is the most accurate. **Dark gray:** moving average with $M = 20$ exploiting underreaction is the most accurate. **Black:** limited memory Bayesian learning is the most accurate.

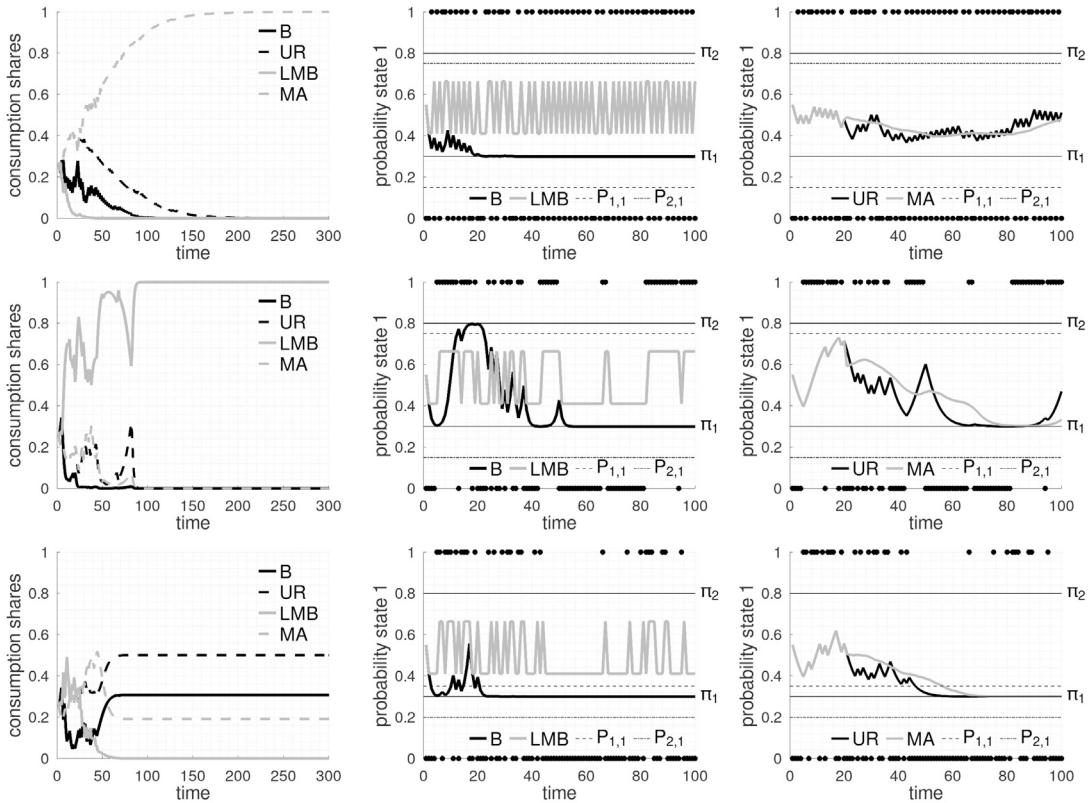


Fig. 7. **Top:** $P_{1,1} = 0.15, P_{2,1} = 0.75$. **Middle:** $P_{1,1} = 0.75, P_{2,1} = 0.15$. **Bottom:** $P_{1,1} = 0.35, P_{2,1} = 0.2$. **Left:** consumption share dynamics in a market populated by agents B, UR, LMB, MA. **Center:** conditional probability attached to state 1 by B and LMB for the first 100 time steps. **Right:** conditional probability attached to state 1 by UR and MA for the first 100 time steps. The black dots on 1 represent $s_t = 1$, those on 0 represent $s_t = 2$.

white areas around $P_{1,1} = P_{2,1} = 0$ and $P_{1,1} = P_{2,1} = 1$, Bayes, underreaction, and moving average achieve the same (maximal) accuracy level. This is analogous to what was observed for large and small values of π in Fig. 1.

We conclude this section with some numerical examples of the dynamics of consumption shares, conditional probabilities, and prices. We populate the economy with the four agents described above, using the same labeling system and making the same

assumptions on discount factors, endowments, and utilities as in Subsection 4.1.1. First, we set $P_{1,1} = 0.15$ and $P_{2,1} = 0.75$, so that the MA agent has the lowest average relative entropy. The share of consumption of the MA agent approaches 1 around time step 200; see the top left panel of Fig. 7. The order in which the consumption of the other agents approaches zero is inversely proportional to the average relative entropy of their conditional probabilities. Concerning the dynamics of the predictions (top-central and top-right panels of Fig. 7), agent B converges to model 1 quite quickly, agent UR shows small fluctuations around a sort of long-run trend, agent LMB strongly fluctuates between its two conditional probabilities, and agent MA captures the long-run trend of agent UR. Concerning the price of the claim in (8), for sufficiently large t , it is completely determined by the conditional probability of agent MA. Next, we set $P_{1,1} = 0.75$ and $P_{2,1} = 0.15$, so that the limited memory Bayesian learning process has the lowest average relative entropy. As shown in the middle row of Fig. 7, the LMB agent achieves a unitary consumption share in less than 100 steps. Looking at how beliefs evolve, agent B settles on model 1 after few fluctuations; agent UR and MA persistently fluctuate, but their smoothing behavior does not allow them to rapidly adapt when a switch occurs; agent LMB, instead, moves and adapts quickly as sequences of equal states alternate. Thus, in this case, it is agent LMB to solely determine the price of the claim in (8) for sufficiently large t . Finally, we set $P_{1,1} = 0.35$ and $P_{2,1} = 0.2$, so that, according to Fig. 6, we are in a situation of multiple maximally accurate agents. As shown in the bottom row of Fig. 7, in less than 100 steps agent B, UR, and MA see their consumption stabilize in positive and heterogeneous shares. Agent LMB's consumption share, instead, goes to zero. Indeed, while agents B, UR, and MA converge to model 1 and become identical, the probability that LMB assigns to state 1 continues to fluctuate between two levels that are outside the range defined by the transition probabilities. As a consequence, the price of the claim in (8) follows model 1 for sufficiently large t .

5. General survival behaviors under model misspecification

Our results show that the kind of model misspecification that agents face crucially affects selection outcomes. In fact, learning rules with low survival prospects under parametric misspecification can dominate when structural misspecification is considered. However, there are belief formation rules that persist in the market no matter the data generating process. For example, survival is guaranteed by any belief structure that causes an agent to always consume its endowment. Another example is provided by Bayesian learning over the learning processes of the other market participants, as the following proposition clarifies.

Proposition 5.1. *Assume that $\forall i, \beta_i = \beta$ and that agent N assigns to the partial history σ_t the weighted arithmetic mean of the likelihoods of the other agents,*

$$p_N(\sigma_t) = \sum_{i=1}^{N-1} v_i p_i(\sigma_t), \forall \sigma_t, \tag{10}$$

with $v_i > 0, \forall i = 1, \dots, N - 1$ and $\sum_{i=1}^{N-1} v_i = 1$. Then, $\limsup_{t \rightarrow \infty} c_N(\sigma_t) > 0$ on all σ .

Proof. See Section A.7. \square

According to (4), agent N derives its conditional probabilities as the averages of the conditional probabilities of the other traders, weighted by their individual likelihood. Therefore, the accuracy of its beliefs asymptotically matches the highest accuracy in the market. To understand the result, consider an economy populated by three agents with the same utility discount factor: agent 1 has nearly correct beliefs, agent 2 has substantially wrong beliefs, and agent 3 behaves as in (10). Then, almost surely, $\lim_{t \rightarrow \infty} p_2(\sigma_t)/p_1(\sigma_t) = 0$ and $\lim_{t \rightarrow \infty} p_3(\sigma_t)/p_1(\sigma_t) = v_1$. Agent 3 manages to maintain its individual likelihood asymptotically proportional to that of agent 1. This is all that matters for survival. In fact, almost surely, $\lim_{t \rightarrow \infty} u'_1(c_1(\sigma_t))/u'_3(c_3(\sigma_t)) = v_1 u'_1(c_1(\sigma_0))/u'_3(c_3(\sigma_0)) > 0$. While agent 2 vanishes, both agents 1 and 3 survive.

This selection result mirrors those obtained with the *Follow the Leader Strategy* (FLS) and the *Follow the Market Strategy* (FMS) by Massari (2017).⁴ However, the learning behavior prescribed by (10), the FLS, and the FMS all share a common feature: to be implemented, they require a trader to know (at least) the learning processes of all the other agents in the market.⁵ This is in contrast with the learning processes investigated in the previous sections, that do not require any information about the market ecology. On the one hand, this indicates that high survival prospects under model misspecification could be more related to the amount of information one trader possesses about market participants than to how sophisticated its learning mechanism is. On the other hand, it confirms the importance of the ecology of traders, as, despite the presence of imitative behaviors and assuming homogeneity in the utility discount factors, prices shall be ultimately dictated by the most accurate learning processes in the market.

6. Conclusions

We study market selection in a complete-market Arrow-Debreu economy considering four learning processes and two cases of model misspecification: parametric and structural. As well as proving new accuracy results on learning processes and extending

⁴ In our framework, the FLS consists in $p_i(s_t | \sigma_{t-1}) = p_j(s_t | \sigma_{t-1})$ with $j : p_j(\sigma_{t-1}) = \arg\max_n \{p_n(\sigma_{t-1})\}$ and $p_i(s_t | \sigma_{t-1}) = \sum_{j \in \mathcal{K}_{t-1}} p_j(s_t | \sigma_{t-1}) / |\mathcal{K}_{t-1}|$ if ties occur (with \mathcal{K}_{t-1} the set of agents whose beliefs have the highest likelihood). Instead, the FMS consists of $p_i(s_t | \sigma_{t-1}) \propto q(s_t) / q(\sigma_{t-1})$.

⁵ Actually, the FMS is more demanding than the other two: it also requires information about preferences, intertemporal discount factors, and endowments.

previous selection results to a context in which some limits may not exist, we show that, in stark contrast with what happens under correct specification, the ecology of traders greatly matters for selection. Deriving a general ranking of learning processes with respect to their survival prospects is difficult, and the type of misspecification considered strongly influences the outcome. Under parametric misspecification, learning processes built upon an averaging approach have a selection advantage over generic regions of the parameter space. Such an advantage partially disappears when structural misspecification occurs, as a trade-off emerges between approximating the best i.i.d. model and capturing the persistent fluctuations in true conditional probabilities.

The examples of learning rules that allow an agent to survive no matter the kind of model misspecification characterizing the economy rely upon imitating the most accurate traders and, as such, require the knowledge of fundamental details concerning all the other market participants. This further confirms that the ecology of traders populating a market and the kind of model misspecification that affects the economy are crucial to understand long-term dynamics.

For our analysis, we have purposely chosen a framework in which selection outcomes are (mostly) driven by belief accuracy. Relaxing or changing part of our assumptions can lead to different conclusions. For example, the assumption of bounded aggregate endowment eliminates the selection effect of risk preferences that has been shown to exist in continuous time and with CRRRA traders by Yan (2008). The analysis in discrete-time economies performed by Bottazzi and Dindo (2022) confirms that an unbounded aggregate endowment can have a nontrivial selection effect in our framework: depending on the assumptions on risk preferences and how the aggregate endowment grows, one may have different scenarios, ranging from risk preferences playing no role to becoming the only things that matter. Along the same lines, investigating how misspecification affects selection under recursive preferences (Easley and Yang, 2015; Dindo, 2019; Borovička, 2020), incomplete markets (Sandroni, 2005; Blume and Easley, 2006), or differential financial constraints (Guerdjikova and Quiggin, 2019) may all be interesting avenues for future contributions.

Appendix A. Proofs of propositions

This section collects the proof of the formal propositions in the paper. We start with a preliminary lemma, used in several proofs, that connects the likelihood ratio of two measures and the average relative entropy of their conditional probabilities.

Lemma A.1. *Let P be a stochastic process on (Σ, \mathcal{F}) adapted to filtration $(\mathcal{F}_t)_{t=0}^\infty$. If $\exists L > 0$ such that, p -almost surely, $\|\log p(\cdot | \sigma_t)/P(\cdot | \sigma_t)\|_\infty < L$, then $\forall \alpha < 1/2$, p -almost surely,*

$$\lim_{t \rightarrow \infty} t^\alpha \left(\frac{1}{t} \log \frac{p(\sigma_t)}{P(\sigma_t)} - \bar{D}_{p|P}(\sigma_{t-1}) \right) = 0.$$

Proof. Define the random variable

$$z(s | \sigma_{\tau-1}) = \log p(s | \sigma_{\tau-1}) / P(s | \sigma_{\tau-1}) - \bar{D}_{p|P}(\sigma_{\tau-1}),$$

so that

$$\log \frac{p(\sigma_t)}{P(\sigma_t)} = \sum_{\tau=1}^t \log \frac{p(s_\tau | \sigma_{\tau-1})}{P(s_\tau | \sigma_{\tau-1})} = \sum_{\tau=1}^t z(s_\tau | \sigma_{\tau-1}) + t \bar{D}_{p|P}(\sigma_{t-1}).$$

For any σ_t , $E[z | \sigma_t] = 0$, and $E[z^2 | \sigma_t] < \|\log p(\cdot | \sigma_t)/P(\cdot | \sigma_t)\|_\infty^2$. If $\alpha < 1/2$,

$$\sum_{t=1}^\infty t^{2\alpha-2} E[z^2 | \sigma_{t-1}] \leq L^2 \sum_{t=1}^\infty t^{2\alpha-2} < +\infty.$$

Thus, by Theorem 3, p. 243, in Feller (1971), p -almost surely,

$$\lim_{t \rightarrow \infty} t^\alpha \left(\frac{1}{t} \log \frac{p(\sigma_t)}{P(\sigma_t)} - \bar{D}_{p|P}(\sigma_{t-1}) \right) = \lim_{t \rightarrow \infty} t^{\alpha-1} \sum_{\tau=1}^t z(s_\tau | \sigma_{\tau-1}) = 0. \quad \square$$

Under the stated condition, when t becomes large, the average likelihood log ratio, $t^{-1} \log p(\sigma_t)/P(\sigma_t)$, and the average relative entropy, $\bar{D}_{p|P}(\sigma_{t-1})$, differ by a term that decreases at a rate that is not slower than $1/\sqrt{t}$.

A.1. Proof of Proposition 2.1

Applying Lemma A.1 to the individual measures of agents i and j , it is

$$\frac{1}{t} \log \frac{p(\sigma_t)}{p_i(\sigma_t)} - \frac{1}{t} \log \frac{p(\sigma_t)}{p_j(\sigma_t)} - \bar{D}_{p|p_i}(\sigma_{t-1}) + \bar{D}_{p|p_j}(\sigma_{t-1}) = o(t^{-\alpha}),$$

so that the first statement follows from (1). For the second statement, note that the hypothesis implies that, p -almost surely, $\lim_{t \rightarrow \infty} \log u'_i(c_i(\sigma_t))/u'_j(c_j(\sigma_t)) = +\infty$. As the endowment is bounded, $u'_j(c_j(\sigma_t))$ is bounded away from zero. Thus, it must be $\lim_{t \rightarrow \infty} \log u'_i(c_i(\sigma_t)) = +\infty$. According to the Inada condition, this, in turn, implies that $\lim_{t \rightarrow \infty} c_i(\sigma_t) = 0$.

A.2. Proof of Proposition 3.1

Let $\pi^*(\sigma_t) = \pi_{k^*(\sigma_t)}(\sigma_t)$. Iterative substitution of (4) in (3) immediately shows that $p_i(\sigma_t) = \sum_{k=1}^K w_{i,k}(\sigma_0)\pi_k(\sigma_t)$. Thus, $w_{i,k^*(\sigma_t)}(\sigma_0)\pi_{k^*(\sigma_t)}(\sigma_t) \leq p_i(\sigma_t) \leq \pi_{k^*(\sigma_t)}(\sigma_t)$ and taking the logarithm, dividing by t , and rearranging terms,

$$0 \leq \frac{1}{t} \log \frac{p(\sigma_t)}{p_i(\sigma_t)} - \frac{1}{t} \log \frac{p(\sigma_t)}{\pi_{k^*(\sigma_t)}(\sigma_t)} \leq -\frac{1}{t} \log w_{i,k^*(\sigma_t)}(\sigma_0).$$

Consider $\alpha < 1/2$. From the previous inequalities,

$$\lim_{t \rightarrow \infty} t^\alpha \left(\frac{1}{t} \log \frac{p(\sigma_t)}{p_i(\sigma_t)} - \frac{1}{t} \log \frac{p(\sigma_t)}{\pi_{k^*(\sigma_t)}(\sigma_t)} \right) = 0.$$

At the same time, from Lemma A.1, p -almost surely,

$$\lim_{t \rightarrow \infty} t^\alpha \left(\frac{1}{t} \log \frac{p(\sigma_t)}{p_i(\sigma_t)} - \bar{D}_{p|p_i}(\sigma_{t-1}) \right) = \lim_{t \rightarrow \infty} t^\alpha \left(\frac{1}{t} \log \frac{p(\sigma_t)}{\pi_{k^*(\sigma_t)}(\sigma_t)} - \bar{D}_{p|\pi_{k^*(\sigma_t)}}(\sigma_{t-1}) \right) = 0.$$

Hence, p -almost surely $\lim_{t \rightarrow \infty} t^\alpha \left(\bar{D}_{p|p_i}(\sigma_{t-1}) - \bar{D}_{p|\pi_{k^*(\sigma_t)}}(\sigma_{t-1}) \right) = 0$.

A.3. Proof of Proposition 3.2

Define $\rho_{i,k}(s_{t+1}|\sigma_t) = \lambda_i p_i(s_{t+1}|\sigma_t) + (1 - \lambda_i) \pi_k(s_{t+1}) \in H_K$. Note that $w_{i,k}(\sigma_t) = w_{i,k}(\sigma_{t-1})\rho_{i,k}(s_t|\sigma_{t-1})/p_i(s_t|\sigma_{t-1})$. Iterative substitution with the previous equation gives

$$p_i(\sigma_t) = p_i(\sigma_{t-1}) \sum_{k=1}^K \rho_{i,k}(s_t|\sigma_{t-1})w_{i,k}(\sigma_{t-1}) = \dots = \sum_{k=1}^K \rho_{i,k}(\sigma_t) w_{i,k}(\sigma_0),$$

where $\rho_{i,k}(\sigma_t) = \prod_{\tau=1}^{t-1} \rho_{i,k}(s_{\tau+1}|\sigma_\tau)$. For the concavity of the logarithm, $\forall k$,

$$\log p_i(\sigma_t) \geq \log \rho_{i,k}(\sigma_t)w_{i,k}(\sigma_0) \geq \lambda_i \log p_i(\sigma_t) + (1 - \lambda_i) \log \pi_k(\sigma_t) + \log w_{i,k}(\sigma_0),$$

which implies $\log p_i(\sigma_t) \geq \log \pi_k(\sigma_t) + 1/(1 - \lambda_i) \log w_{i,k}(\sigma_0)$. Consider $\alpha < 1/2$. From the previous inequality,

$$\lim_{t \rightarrow \infty} t^\alpha \left(\frac{1}{t} \log \frac{p(\sigma_t)}{p_i(\sigma_t)} - \frac{1}{t} \log \frac{p(\sigma_t)}{\pi_k(\sigma_t)} \right) \leq -\lim_{t \rightarrow \infty} \frac{t^{\alpha-1}}{(1 - \lambda_i)} \log w_{i,k}(\sigma_0) = 0.$$

At the same time, from Lemma A.1, p -almost surely,

$$\lim_{t \rightarrow \infty} t^\alpha \left(\frac{1}{t} \log \frac{p(\sigma_t)}{p_i(\sigma_t)} - \bar{D}_{p|p_i}(\sigma_{t-1}) \right) = \lim_{t \rightarrow \infty} t^\alpha \left(\frac{1}{t} \log \frac{p(\sigma_t)}{\pi_k(\sigma_t)} - \bar{D}_{p|\pi_k}(\sigma_{t-1}) \right) = 0,$$

whence the assertion.

A.4. Proof of Proposition 4.1

Let $\pi = \sum_{k=1}^K \zeta_k \pi_k$, with $\zeta_k \geq 0$ and $\sum_{k=1}^K \zeta_k = 1$. Define $\rho_{i,k}(s_{t+1}|\sigma_t)$ and $\rho_{i,k}(\sigma_t)$ as in Section A.3 and remember that $p_i(\sigma_t) = \sum_{k=1}^K \rho_{i,k}(\sigma_t)w_{i,k}(\sigma_0)$. Using the Taylor expansion with Lagrange remainder in $1 - \lambda_i$, $\forall \sigma_t$, $\exists \eta_{k,s}(\sigma_t) \in [0, 1 - \lambda_i]$ such that

$$D_{p|p_i}(\sigma_t) - \sum_{k=1}^K \zeta_k D_{p|\rho_{i,k}}(\sigma_t) = (1 - \lambda_i) \sum_{s=1}^S \pi(s) \left(\frac{\pi(s)}{p_i(s|\sigma_t)} - 1 \right) - \frac{(1 - \lambda_i)^2}{2} \sum_{k=1}^K \zeta_k \sum_{s=1}^S \pi(s) \frac{(\pi_k(s) - p_i(s|\sigma_t))^2}{(\eta_{k,s}(\sigma_t)\pi_k(s) + (1 - \eta_{k,s}(\sigma_t))p_i(s|\sigma_t))^2}.$$

From Assumption 1,

$$\sum_{k=1}^K \zeta_k \sum_{s=1}^S \pi(s) \frac{(\pi_k(s) - p_i(s|\sigma_t))^2}{(\eta_{k,s}(\sigma_t)\pi_k(s) + (1 - \eta_{k,s}(\sigma_t))p_i(s|\sigma_t))^2} \leq \frac{1}{(1 - \theta + \epsilon)^2},$$

and because $x - 1 \geq \log x$,

$$\sum_{s=1}^S \pi(s) \left(\frac{\pi(s)}{p_i(s|\sigma_t)} - 1 \right) \geq D_{p|p_i}(\sigma_t),$$

so that

$$D_{p|p_i}(\sigma_t) - \sum_{k=1}^K \zeta_k D_{p|\rho_{i,k}}(\sigma_t) \geq (1 - \lambda_i) D_{p|p_i}(\sigma_t) - \frac{(1 - \lambda_i)^2}{2(\lambda_i + \epsilon)^2}.$$

Consider now $\alpha < 1/2$. Sum on τ from 0 to $t - 1$ and divide by t both sides of the previous inequality. For the left-hand side, using Lemma A.1 and simplifying common terms,

$$\lim_{t \rightarrow \infty} t^\alpha \left(\overline{D}_{p|p_i}(\sigma_{t-1}) - \sum_{k=1}^K \zeta_k \overline{D}_{p|p_{i,k}}(\sigma_{t-1}) \right) = \lim_{t \rightarrow \infty} t^{\alpha-1} \sum_{k=1}^K \zeta_k \log \frac{p_{i,k}(\sigma_t)}{p_i(\sigma_t)} \leq - \lim_{t \rightarrow \infty} t^{\alpha-1} \sum_{k=1}^K \zeta_k \log w_{i,k}(\sigma_0) = 0.$$

Thus, for the right-hand side,

$$\lim_{t \rightarrow \infty} t^\alpha \left((1 - \lambda_i) \overline{D}_{p|p_i}(\sigma_{t-1}) - \frac{(1 - \lambda_i)^2}{2(\lambda_i + \epsilon)^2} \right) \leq 0,$$

whence the assertion.

A.5. Proof of Proposition 4.2

From the bounds of the arithmetic and geometric means inequality in Perisastry and Murty (1982) and the bound on probability models in Assumption 1, $\forall s, \sigma_\tau$,

$$\frac{\sigma^2(s, \sigma_t)}{2(1 - \epsilon)} \leq \log p_i(s | \sigma_t) - \frac{1}{M} \sum_{m=0}^{M-1} \log p^*(s | \sigma_{t-m}) \leq \frac{\sigma^2(s, \sigma_t)}{2\epsilon}.$$

Adding and subtracting $\log \pi(s)$ in the middle term,

$$\frac{\sigma^2(s, \sigma_t)}{2(1 - \epsilon)} \leq \frac{1}{M} \sum_{m=0}^{M-1} \log \frac{\pi(s)}{p^*(s | \sigma_{t-m})} - \log \frac{\pi(s)}{p_i(s | \sigma_t)} \leq \frac{\sigma^2(s, \sigma_t)}{2\epsilon}.$$

Multiplying for $\pi(s)$ and summing over s proves the assertion.

A.6. Proof of Proposition 4.3

Using the Law of Large Numbers for variables of bounded variance,

$$\overline{D}_{p|\pi_k}(\sigma) = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \sum_{s=1}^S p(s|\sigma_\tau) \log \frac{p(s|\sigma_\tau)}{\pi_k(s)} = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \sum_{s=1}^S P_{s_\tau, s} \log \frac{P_{s_\tau, s}}{\pi_k(s)} = \sum_{s'=1}^S \pi(s') \sum_{s=1}^S P_{s', s} \log \frac{P_{s', s}}{\pi_k(s)}.$$

The statement follows by adding and subtracting $\sum_{s=1}^S \pi(s) \log \pi(s)$ and exploiting the properties of the invariant distribution, i.e. $\pi(s) = \sum_{s'=1}^S P_{s', s} \pi(s') \forall s$.

A.7. Proof of Proposition 5.1

If $\lim_{t \rightarrow \infty} c_N(\sigma_t) = 0$, then $\lim_{t \rightarrow \infty} \sum_{i=1}^{N-1} c_i(\sigma_t) = e(\sigma_t) > 0$. This implies that $\lim_{t \rightarrow \infty} u'_N(c_N(\sigma_t))^{-1} = 0$ and $\liminf_{t \rightarrow \infty} \sum_{i=1}^{N-1} v_i u'_i(c_i(\sigma_t))^{-1} > 0$. So, if agent N vanishes on σ ,

$$\lim_{t \rightarrow \infty} \frac{u'_N(c_N(\sigma_t))^{-1}}{\sum_{i=1}^{N-1} v_i u'_i(c_i(\sigma_t))^{-1}} = \lim_{t \rightarrow \infty} \frac{p_N(\sigma_t)}{\sum_{j=1}^{N-1} v_j p_j(\sigma_t)} = 0,$$

but this is impossible as $p_N(\sigma_t) / \sum_{i=1}^{N-1} v_i p_i(\sigma_t) = 1 \forall t$.

References

Antico, A., Bottazzi, G., Giachini, D., 2023. On the evolutionary stability of the sentiment investor. In: Bourghelle, D., Grandin, P., Jawadi, F., Rozin, P. (Eds.), Behavioral Finance and Asset Prices: The Influence of Investor's Emotions. Springer International Publishing, Cham, pp. 155–173.

Barberis, N., Shleifer, A., Vishny, R., 1998. A model of investor sentiment. J. Financ. Econ. 49 (3), 307–343.

Beddock, A., Jouini, E., 2021. Live fast, die young: equilibrium and survival in large economies. Econ. Theory 71 (3), 961–996.

Blume, L., Easley, D., 2006. If you are so smart why aren't you rich? Belief selection in complete and incomplete markets. Econometrica 74, 929–966.

Blume, L., Easley, D., 2009a. The market organism: long-run survival in markets with heterogeneous traders. J. Econ. Dyn. Control 33, 1023–1035.

Blume, L., Easley, D., 2009b. Market selection and asset pricing. In: Hens, T., Schenk-Hoppé, K. (Eds.), Handbook of Financial Markets: Dynamics and Evolution. In: Handbooks in Economics Series. North-Holland.

Borovička, J., 2020. Survival and long-run dynamics with heterogeneous beliefs under recursive preferences. J. Polit. Econ. 128 (1), 206–251.

Bottazzi, G., Dindo, P., 2013. Selection in asset markets: the good, the bad, and the unknown. J. Evol. Econ. 23, 641–661.

Bottazzi, G., Dindo, P., 2014. Evolution and market behavior with endogenous investment rules. J. Econ. Dyn. Control 48, 121–146.

Bottazzi, G., Dindo, P., 2022. Drift criteria for persistence of discrete stochastic processes on the line. J. Math. Econ. 101, 102696.

Bottazzi, G., Dindo, P., Giachini, D., 2018. Long-run heterogeneity in an exchange economy with fixed-mix traders. Econ. Theory 66 (2), 407–447.

Bottazzi, G., Dindo, P., Giachini, D., 2019. Momentum and reversal in financial markets with persistent heterogeneity. Ann. Finance 15 (4), 455–487.

Bottazzi, G., Giachini, D., 2017. Wealth and price distribution by diffusive approximation in a repeated prediction market. Physica A 471, 473–479.

Bottazzi, G., Giachini, D., 2019a. Betting, selection, and luck: a long-run analysis of repeated betting markets. Entropy 21 (6), 585.

Bottazzi, G., Giachini, D., 2019b. Far from the madding crowd: collective wisdom in prediction markets. Quant. Finance 19 (9), 1461–1471.

- Bottazzi, G., Giachini, D., 2022. Strategically biased learning in market interactions. *Adv. Complex Syst.* 25 (02n03), 2250004.
- Box, G.E., 1976. Science and statistics. *J. Am. Stat. Assoc.* 71 (356), 791–799.
- Cerreia-Vioglio, S., Hansen, L.P., Maccheroni, F., Marinacci, M., 2020. Making decisions under model misspecification. *arXiv preprint arXiv:2008.01071*.
- Cox, D.R., 1995. Comment on “model uncertainty, data mining and statistical inference”. *J. R. Stat. Soc., Ser. A, Stat. Soc.* 158, 455–456.
- Dindo, P., 2019. Survival in speculative markets. *J. Econ. Theory* 181, 1–43.
- Dindo, P., Massari, F., 2020. The wisdom of the crowd in dynamic economies. *Theor. Econ.* 15 (4), 1627–1668.
- Easley, D., Yang, L., 2015. Loss aversion, survival and asset prices. *J. Econ. Theory* 160, 494–516.
- Elmiger, S., 2020. Market selection and the absence of arbitrage. Available at SSRN 3746988.
- Epstein, L.G., Noor, J., Sandroni, A., 2010. Non-Bayesian learning. *B. E. J. Theor. Econ.* 10 (1), 1–20.
- Esponda, I., Pouzo, D., 2016. Berk–Nash equilibrium: a framework for modeling agents with misspecified models. *Econometrica* 84 (3), 1093–1130.
- Evstigneev, I., Hens, T., Schenk-Hoppé, K., 2009. Evolutionary finance. In: Hens, T., Schenk-Hoppé, K. (Eds.), *Handbook of Financial Markets: Dynamics and Evolution*. In: *Handbooks in Economics Series*. North-Holland.
- Evstigneev, I., Hens, T., Schenk-Hoppé, K.R., 2016. Evolutionary behavioral finance. In: *The Handbook of Post Crisis Financial Modeling*. Springer, pp. 214–234.
- Feller, W., 1971. *An Introduction to Probability Theory and Its Applications*, Volume II. John Wiley and Sons, New York.
- Fudenberg, D., Romanyuk, G., Strack, P., 2017. Active learning with a misspecified prior. *Theor. Econ.* 12 (3), 1155–1189.
- Giachini, D., 2021. Rationality and asset prices under belief heterogeneity. *J. Evol. Econ.* 31 (1), 207–233.
- Gigerenzer, G., Gaissmaier, W., 2011. Heuristic decision making. *Annu. Rev. Psychol.* 62 (1), 451–482.
- Guerdjikova, A., Quiggin, J., 2019. Market selection with differential financial constraints. *Econometrica* 87 (5), 1693–1762.
- Hansen, L.P., 2014. Nobel lecture: uncertainty outside and inside economic models. *J. Polit. Econ.* 122 (5), 945–987.
- Hansen, L.P., Sargent, T.J., 2022. Structured ambiguity and model misspecification. *J. Econ. Theory* 199, 105165.
- Hens, T., Schenk-Hoppé, K., 2005. Evolutionary finance: introduction to the special issue. *J. Math. Econ.* 41, 1–5.
- Holtfort, T., 2019. From standard to evolutionary finance: a literature survey. *Manag. Rev. Q.* 69 (2), 207–232.
- Jouini, E., Napp, C., 2011. Unbiased disagreement in financial markets, waves of pessimism and the risk-return trade-off. *Rev. Finance* 15 (3), 575–601.
- Kogan, L., Ross, S.A., Wang, J., Westerfield, M.M., 2006. The price impact and survival of irrational traders. *J. Finance* 61 (1), 195–229.
- Kogan, L., Ross, S.A., Wang, J., Westerfield, M.M., 2017. Market selection. *J. Econ. Theory* 168, 209–236.
- Marinacci, M., Massari, F., 2019. Learning from ambiguous and misspecified models. *J. Math. Econ.* 84, 144–149.
- Massari, F., 2017. Markets with heterogeneous beliefs: a necessary and sufficient condition for a trader to vanish. *J. Econ. Dyn. Control* 78, 190–205.
- Massari, F., 2021. Price probabilities: a class of Bayesian and non-Bayesian prediction rules. *Econ. Theory* 72, 133–166.
- Massari, F., 2020. Under-reaction: Irrational behavior or robust response to model misspecification? Available at SSRN 3636136.
- Orseau, L., Lattimore, T., Legg, S., 2017. Soft-bayes: prod for mixtures of experts with log-loss. In: *International Conference on Algorithmic Learning Theory*. PMLR, pp. 372–399.
- Perisastry, M., Murty, V.M., 1982. Bounds for the ratio of the arithmetic mean to the geometric mean. *Two-Year Coll. Math. J.* 13 (2), 160–161.
- Sandroni, A., 2000. Do markets favor agents able to make accurate predictions. *Econometrica* 68 (6), 1303–1341.
- Sandroni, A., 2005. Market selection when markets are incomplete. *J. Math. Econ.* 41, 91–104.
- Yan, H., 2008. Natural selection in financial markets: does it work? *Manag. Sci.* 54, 1935–1950.