Economic Theory
https://doi.org/10.1007/s00199-025-01677-x

RESEARCH ARTICLE

Under-reaction: Irrational behavior or robust response to model misspecification?

Filippo Massari [1]

Received: 20 October 2023 / Accepted: 28 July 2025

Abstract
I find that under-reaction is a robust response to model misspecification, rewarded by financial markets, rather than an "irrational" attitude that leads to extinction. When a Bayesian agent trades with an under-reacting agent who has access to the same information, there are no paths on which the under-reacting agent loses all his wealth to the Bayesian. Conversely, the Bayesian agent loses all his wealth to the under-reacting agent in misspecified learning settings, provided that a combination of parameters is more accurate than any single parameter in the support, and the under-reaction is sufficiently strong.

Keywords Misspecified learning · Under-reaction · Conservatism bias · Positive prior bias · Market selection · Behavioral finance

JEL Classification D53 · D83 · D9 · G1 · G4

1 Introduction

A long-standing conjecture about learning and financial markets is that "rational" Bayesian agents eventually drive "irrational" non-Bayesian agents–who use the same information but make asymptotically different predictions–out of the market.

This conjecture has been theoretically investigated and confirmed in well-specified learning problems (Sandroni 2000; Blume and Easley 2006), i.e., under the assumption that Bayesian agents eventually learn the truth. However, little work has been done

The author thanks Pietro Dindo, Simone Cerreia Vioglio, Fabio Maccheroni, Massimo Marinacci and Jonathan Newton for their guidance. This project has been funded by European Union - Next Generation EU, Mission 4 Component 2, CUP J53D23004180006.

✉ Filippo Massari
massari3141@gmail.com

1 University of Bologna, Bologna, Italy

allowing for misspecified learning environments.[1] In misspecified learning settings, the true model is not in the prior support, and it is never learned because the agent subjectively believes it to be impossible.

The main difference between well-specified and misspecified environments is as follows: in general, the Bayesian approach guarantees predictions that are as accurate as those of the most accurate model in the support (Berk 1966). If the learning problem is correctly specified, the most accurate model in the support is also the true model, and alternative belief-updating rules can only match Bayes' accuracy. Conversely, in misspecified learning environments, the most accurate model in the support may be less accurate than a combination of models in the support. Thus, a non-Bayesian belief-updating rule might combine models in the support to deliver even more accurate predictions than Bayes'.

In this paper, I focus on under-reaction, a heuristic widely documented in empirical finance (Barberis et al. 1998; Giglio and Kelly 2018) and in experimental settings—the formula I adopt is consistent with the extensive experimental evidence in favour of *under-inference, conservatism bias* and *positive prior bias* (for references, see the meta-analysis and review in Benjamin 2019, Ch.2). In the standard general equilibrium model with complete markets of Sandroni (2000) and Blume and Easley (2006), I study the consumption-share dynamics between a Bayesian (B) and a non-Bayesian (NB) agent who under-reacts to information according to a *prospective-acceptive* clumping attitude described in Benjamin et al. (2016) of the learning rule axiomatized in Epstein (2006), Epstein et al. (2008) and discussed in Bohren and Hauser (2023). The two agents have the same discount factor and identical information since they share the same prior support and observe the same path of realizations. In this market setup, an agent vanishes if there is another agent who is more accurate and at least as patient.[2] I show that agent NB has an evolutionary advantage over agent B. Specifically,

Theorem 1 shows that there are no paths, and thus, no true data-generating processes, such that agent NB vanishes against agent B. Agent NB survives even in well-specified learning problems in which agent B learns the truth because he learns the truth at a comparable rate.

Theorem 3 provides sufficient conditions for agent B to vanish against agent NB. If a combination of parameters (or models) in the prior leads to a more accurate model than any parameter (conditions **C2**), there is a level of under-reaction (condition **C3**) such that agent NB dominates. While the Bayesian agent is as accurate as the most accurate model in the support $\pi_{\hat{\theta}}$ (Berk 1966), the under-reacting agent is more accurate than $\pi_{\hat{\theta}}$ because its predictions remain an appropriate non-degenerate combination of models in the prior.

My findings are of particular interest to the portfolio selection literature, where the true process of stock returns is unknown, and the evidence in favor of Bayesian meth-

[1] The work of Sandroni (2005) lies between well-specified and misspecified settings. It focuses on a specific class of learning problems in which the Bayesian agent cannot learn the truth because the true data-generating process changes over time, but a version of Wald (1947)'s complete class theorem holds.

[2] Theorems 1, 2, and 3 apply to any finite-agent economy since agent survival depends solely on relative belief accuracy and is independent of equilibrium prices.

ods is mixed. While there is a vast literature supporting Bayesian methods for portfolio selection problems (Klein and Bawa 1976; Frost and Savarino 1986), an equally rich body of work proposes non-Bayesian approaches for robust portfolio allocation rules (DeMiguel et al. 2007; Garlappi et al. 2006). My result complements and strengthens the arguments in favour of the "robust" methods proposed in the portfolio selection literature. It complements them by identifying model misspecification–rather than parameter estimation error–as the reason for the suboptimal performance of Bayesian methods. It strengthens them due to the known magnification effect that model misspecification has on parameter estimation error. In misspecified learning problems, Bayesian posterior convergence can be significantly slower than in well-specified settings (Grünwald and Van Ommen 2017), warranting special consideration in the use of Bayesian methods for prediction (Grünwald and Langford 2007; Grünwald 2012).

Furthermore, my result offers an intuitive explanation for the difficulties often encountered in exploiting under-reaction anomalies through practical trading strategies: it is hard to take advantage of these anomalies because under-reaction is a robust response to model misspecification, rather than an irrational deviation from a correct learning procedure.

In Appendix A, I discuss the likelihood of encountering situations in which an under-reacting rule outperforms Bayes'. I argue that conditions **C1**-**C2** are likely to hold in many learning models in the portfolio selection literature cited above because they have prior support consisting of a finite set of models. In Appendix B, I prove a series of results on the relative accuracy between Bayes' rule and prediction rules that under-react to information. These findings are instrumental to Theorems 1 - 3, which are proven in Appendix C. Furthermore, the results in Appendix B nest those of Epstein et al. (2010) for well-specified learning environments and generalize them to potentially misspecified environments. In a well-specified learning environment, Epstein et al. (2010) shows that under-reaction is a transient phenomenon because eventually it delivers predictions as accurate as Bayes' — under-reacting rules *weakly merge* to the truth (Kalai and Lehrer 1994). Here, I show that their conclusion critically depends on the assumption that the learning problem is well specified; it does not hold when we allow for misspecification. For every finite prior support, there is a generic set of parameters for the true data-generating process such that agent B and agent NB predictions remain distinct, and agent B is less accurate than agent NB. In these circumstances, under-reaction is a long-lasting phenomenon because the market reflects the beliefs of the under-reacting agent in every distant future.

2 The model

I consider an infinite horizon Arrow-Debreu exchange economy with complete markets. Time is discrete, indexed by t, and begins at date $t = 0$. In each period $t \geq 1$, the economy can be in one of S mutually exclusive states, $\mathcal{S}$. The set of partial histories until t is the Cartesian product $\Sigma^t = \times^t \mathcal{S}$ and the set of all paths is $\Sigma := \times^\infty \mathcal{S}$. $\sigma = (\sigma_1, ...)$ is a representative path, $\sigma^t = (\sigma_1, ..., \sigma_t)$ is a partial history until period t, and $\mathcal{F}_t$ is the σ-algebra generated by the cylinders with base σ^t. By construction $(\mathcal{F}_t)_{t=0}^\infty$ is a filtration and $\mathcal{F}$ is the σ-algebra generated by their union. P denotes the true measure on $(\Sigma, \mathcal{F})$; unless stated, I make no assumptions on P; P_{θ^*} indicates the empirical distribution of states, with θ^* indicating the vector of limit frequencies of each state. For any probability measure ρ on $(\Sigma, \mathcal{F})$, $\rho(\sigma^t) := \rho(\{\sigma_1 \times ... \times \sigma_t \times S \times S \times ...\})$ is the marginal probability of the partial history σ^t, while $\rho_t := \rho(\sigma_t|) := \frac{\rho(\sigma^t)}{\rho(\sigma^{t-1})}$ is the conditional probability of the generic state σ_t given σ^{t-1}.

Next, I introduce a number of economic variables with time index t. All these variables are adapted to the information filtration $(\mathcal{F}_t)_{t=0}^{\infty}$.

The economy has two agents indexed by i: a Bayesian (B) and a non-Bayesian (NB) with common discount factor β.[3] For all paths σ, each agent i is endowed with a stream of the consumption good, $(e_t^i(\sigma))_{t=0}^{\infty}$. Agents' objective is to maximize their stream of discounted subjective expected utility from consumption. Subjective expectations are computed according to agent beliefs p^i, a measure on $(\Sigma, \mathcal{F})$. Naming $q(\sigma^t)$ the time zero price of the asset that delivers one unit of consumption in the period/event σ^t and none otherwise, agent $i = B, NB$ maximization reads:

$$\max_{(c_t^i(\sigma))_{t=0}^{\infty}} E_{p^i}\left[\sum_{t=0}^{\infty} \beta^t u^i(c_t^i(\sigma))\right] \quad s.t. \quad \sum_{t\geq 0}\sum_{\sigma^t \in \Sigma^t} q(\sigma^t)\left(c_t^i(\sigma) - e_t^i(\sigma)\right) \leq 0.$$

A competitive equilibrium is a sequence of prices and, for each agent, a consumption plan that is preference maximal on the budget set, and such that markets clear in every period: $\forall(t, \sigma)$, $\sum_{i=B,NB} e_t^i(\sigma) = \sum_{i=B,NB} c_t^i(\sigma)$. Assumptions **A1-A3** below are standard in the market selection literature to ensure the existence of a competitive equilibrium (Peleg and Yaari 1970).

A1 For all agents $i \in \mathcal{I}$ the utility $u^i : \mathbb{R}_+ \to [-\infty, +\infty]$ is C^1, strictly concave, increasing, and satisfies the Inada condition at 0; that is, $\lim_{c \searrow 0} u^i(c)' = \infty$.

A2 The aggregate endowment is uniformly bounded from above and away from 0:

$$\infty > F > \sup_{t,\sigma} \sum_{i=B,NB} e_t^i(\sigma) > \inf_{t,\sigma} \sum_{i=B,NB} e_t^i(\sigma) > f > 0.$$

A3 $\exists \epsilon > 0$ such that for all agents $i = B, NB$ and for all (t, σ), $p^i(\sigma_t|) > \epsilon$.

In our learning environment, a sufficient condition for **A3** to hold is that the (common) prior support of the two agents contains at least one strictly positive measure.

2.1 Beliefs' dynamics

The Bayesian agent and the under-reacting agent share identical information. Adopting Bayesian terminology, at time zero they believe that states are i.i.d. according to multinomial distributions π_θ parametrized by K vectors of parameters $\theta \in \Delta^{|S|-1}$. Both agents have an identical full-support, time-zero prior distribution μ_0 over the set

[3] I assume a common discount factor to guarantee that the market selects for the most accurate agent(s) rather than for those that save the most.

of parameters $\Theta = \{\theta^1, ...\theta^K\}$ and construct posterior probabilities upon observing the same history σ^t.[4]

Agent NB under-reacts to information; his next-period beliefs are recursively revised according to the formula axiomatized by Epstein (2006), Epstein et al. (2008).[5] Consider an agent who is trying to learn the true parameter in a set Θ. Updating of beliefs in response to observations $\{\sigma_1, ..., \sigma_t\}$ leads to the process of posteriors $\{\mu_t\}$, each μ_t is a probability measure on Θ. Bayesian updating leads to the process

$$\mu_{t+1} = BU(\mu_t; \sigma_{t+1})$$

where $BU(\mu_t; \sigma_{t+1}) := \mu^B(\cdot|\sigma^{t+1})$ denotes the Bayesian update of $\mu(\cdot|\sigma^t)$ upon observing state σ_{t+1} (Definition 2). A belief updating rule under-reacts to information (has *prior-bias*) with respect to Bayes' if the weights given by its posterior distribution are a convex combination between the prior weights and the Bayesian posterior weights calculated using the same information

Definition 1 A belief updating rule **under-reacts to information** if its process of posteriors $\{\mu_t\}$ is

$$\mu_{t+1} = (1-\alpha)\mu_t + \alpha BU(\mu_t; \sigma_{t+1})$$

where $\alpha \in (0, 1)$.

A belief updating rule satisfying Definition 1 can be interpreted as attaching too much weight to prior beliefs μ_t and hence under-reacting to observations. The parameter α regulates the amount of under-reaction of agent NB. Lower values of α correspond to stronger under-reaction; with $\alpha = 1$, Definitions 1 is Bayes' rule.

The following definitions characterize the dynamics of agents' beliefs.

Definition 2 The next-period beliefs of agent B evolve according to Bayes' rule:

$$\forall(t, \sigma), \quad \begin{cases} p^B(\sigma_t|) &= \sum_{\theta\in\Theta} \pi_\theta(\sigma_t)\mu^B(\theta|\sigma^{t-1}) \\ \mu^B(\theta|\sigma^t) &= \frac{\pi_\theta(\sigma_t)}{p^B(\sigma_t|)}\mu^B(\theta|\sigma^{t-1}) \end{cases}.$$

Definition 3 Let $\alpha \in (0, 1)$, the next-period beliefs of agent NB evolve as follow:[6]

$$\forall(t, \sigma), \quad \begin{cases} p^{NB}(\sigma_t|) &= \sum_{\theta\in\Theta} \pi_\theta(\sigma_t)\mu^{NB}(\theta|\sigma^{t-1}) \\ \mu^{NB}(\theta|\sigma^t) &= \frac{p_\theta(\sigma_t|)}{p^{NB}(\sigma_t|)}\mu^{NB}(\theta|\sigma^{t-1}) \\ p_\theta(\sigma_t|) &= (1-\alpha)p^{NB}(\sigma_t|) + \alpha\pi_\theta(\sigma_t) \end{cases}.$$

[4] All results remain valid for heterogeneous full-support prior distributions on the same finite support.

[5] The belief updating rule I use is the same as the one axiomatized by Epstein (2006) to rationalize the behaviour of an agent who is self-aware of her biases and fully anticipates her updating behavior when formulating plans. However, agents in my model have standard time-separable utilities and must have time-consistent beliefs to avoid arbitrage. Borrowing from the experimental literature I assume that agents form time-consistent beliefs according to the perspective-acceptive attitude described in Benjamin et al. (2016) (see Section 2.2). Thus, Epstein (2006) axiomatization does not apply to agents in my model.

[6] Inspection of Definition 3 reveals that p^{NB} belong to the class of market probability introduced by Massari (2021); specifically, they correspond to the evolution of the risk neutral probability of an economy in which agents have log-utility and behavioral beliefs described in Dindo and Massari (2020).

Proposition 1, below, verifies that indeed agent NB's beliefs under-react to information according to Definition 1.

Proposition 1 *For all $\alpha \in (0, 1)$, p^{NB} under-reacts to information:*

$$\begin{aligned}\forall \theta \in \Theta, \mu^B(\theta|\sigma^{t-1}) = \mu^{NB}(\theta|\sigma^{t-1}) \Rightarrow \mu^{NB}(\theta|\sigma^t) \\ = (1-\alpha)\mu^B(\theta|\sigma^{t-1}) + \alpha\mu^B(\theta|\sigma^t).\end{aligned}$$

2.2 Unconditional Beliefs

Bayes' rule defines time-consistent belief dynamics. The unconditional beliefs on cylinders of length t, σ^t, obtained from Definition 2, coincide with the product of next-period probabilities computed recursively using Definition 2 on the same partial history. Accordingly, we have that

Definition 4 The unconditional beliefs of agent B are as follows:

$$\forall(t, \sigma), \quad p^B(\sigma^t) = \prod_{\tau=1}^{t} p^B(\sigma_\tau|) = \sum_{\theta\in\Theta} \pi_\theta(\sigma^t)\mu^B(\theta|\emptyset).$$

Conversely, non-Bayesian updating is inherently time-inconsistent. The unconditional beliefs on cylinders of length t, σ^t, obtained from Definition 3, differ from the product of next-period probabilities computed recursively using Definition 3 on the same partial history. Thus, different ways of clumping information result in different unconditional probabilities. Time inconsistency is problematic in a dynamically complete market, as it generates arbitrage opportunities: the time-zero allocation plan and the recursively constructed allocation plans differ.

I assume that our NB agent avoids arbitrage by adopting a *prospective-acceptive* clumping attitude (Benjamin et al. 2016): at any date when the agent makes a decision, he processes signals received before and after that date as belonging to separate groups– and he anticipates this grouping in advance.[7] Specifically, the definition below models the beliefs of an NB agent who anticipates that, in every period, he will observe a new state of the economy and incorporate this information recursively.

Definition 5 The unconditional beliefs of agent NB are as follows:

$$\forall(t, \sigma), \quad p^{NB}(\sigma^t) := \prod_{\tau=1}^{t} p^{NB}(\sigma_\tau|). \tag{1}$$

By construction, equation (1) defines time-consistent beliefs. Since every time-consistent probability admits a Bayesian representation, p^{NB} also admits a Bayesian representation (see Lemma 1). So, are agent NB's beliefs Bayesian or non-Bayesian?

[7] Benjamin (2019) presents experimental evidence supporting agents' having an *acceptive* clumping attitude.

Agent NB's next-period beliefs are non-Bayesian because they do not obey Bayes' rule given the information he has. Experimentally, his beliefs are consistent with *under-inference, conservatism bias*, and *positive prior bias*, which are predominant biases in the population (Benjamin 2019). Conversely, agent NB's beliefs on infinite sequences are Bayesian because they are time-consistent. However, if these beliefs were tested experimentally, they would be considered Bayesian only with respect to a prior support containing models different from those given by the experimenter (i.e., i.i.d. multinomial models). Furthermore, these models are path-dependent and can only be calculated recursively using the under-reacting formula above–a nearly impossible cognitive task. In the emerging economic theory jargon, these types of beliefs are often called Bayesian with biased likelihood. This term indicates a learning procedure that defines time-consistent beliefs with dynamics that differ from the Bayes rule dynamic derived from the given prior beliefs.

2.3 Agents' accuracy and survival

The asymptotic fate of an agent is characterized by their consumption shares as follows.

Definition 6 Agent i **vanishes** on path σ if $\lim_{t\to\infty} c_t^i(\sigma) = 0$ on σ, he **survives** on path σ if $\limsup_{t\to\infty} c_t^i(\sigma) > 0$ on σ.

Agents' accuracy is ranked according to their likelihoods and, more coarsely, by their average (conditional) relative entropies, reflecting the well-known relationship between agents' likelihood and their survival (Sandroni 2000; Massari 2017).

Definition 7
- Agent i is **more accurate** than agent j on σ if $\lim_{t\to\infty} \frac{p^j(\sigma^t)}{p^i(\sigma^t)} = 0$.
- Agent i is **averagely more accurate** than agent j if $\bar{d}(P||p^i) < \bar{d}(P||p^j)$, P-a.s.; where

$$\bar{d}(P||p) := \lim_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^{t} d_\tau(P||p) = \lim_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^{t} E_{P_\tau}\left[\ln \frac{P(\sigma_\tau|)}{p(\sigma_\tau|)}\right],$$

is the average (conditional) relative entropy from p to the true probability P.[8]

The average (conditional) relative entropies accuracy ranking is the standard in the market selection literature, being instrumental to the following sufficient condition for an agent to vanish.[9]

Proposition 2 Sandroni (2000)*: under **A1-A3**, agent i vanishes P-a.s. if agent j is averagely more accurate.*

[8] The relative entropy $d_t(P||p)$ is uniquely minimized at $p_t = P_t$, strictly convex, and $d_t(P||p) = \bar{d}(P||p)$ P-a.s. whenever P and p are i.i.d. measures. Under **A3**, the average relative entropy is an approximation of the average likelihood ratio that holds almost surely according to the true probability.

[9] I use the average (conditional) relative entropies only to identify sufficient conditions for the Bayesian agent to vanish. Following an established standard, I assume that $\lim_{t\to\infty} \frac{1}{t}\sum_{\tau=1}^{t} d_\tau(P||p)$ exists P-a.s., when relying on this quantity.

Conditions that rely on average accuracy, however, are unable to deliver the converse of Proposition 2 even in a two-agent economy.[10] To precisely characterize agents' survival the first accuracy notion must be used (Massari 2017). In a two-agent economy, Theorem 1 of Massari (2017) can be refined to show that accuracy and survival are linked as follows.

Proposition 3 *Under **A1-A3**, agent NB survives on path σ with consumption shares uniformly bounded away from zero if and only if $\frac{p^{NB}(\sigma^t)}{p^B(\sigma^t)}$ is uniformly bounded away from zero:*

$$\exists \eta > 0 : \forall t,\ \frac{p^{NB}(\sigma^t)}{p^B(\sigma^t)} > \eta \text{ on path } \sigma \ \Leftrightarrow\ \exists \eta' > 0 : \forall t,\ c_t^{NB} > \eta' \text{ on path } \sigma.$$

In this paper, I make use of both results: Proposition 2 to provide sufficient conditions under which agent B vanishes almost surely in misspecified learning problems (Theorems 2 and 3); Proposition 3 to show that the NB agent survives on every path (Theorem 1), and thus for every true data generating process.

3 Main result

In this section, I demonstrate that agent NB has an advantage over agent B. Theorem 1 tells us that an under-reacting agent cannot lose all his wealth against a Bayesian agent with the same information.

Theorem 1 *Under* **A1-A3***, on every equilibrium path σ and for all $\alpha \in (0, 1]$, agent NB survives with consumption shares uniformly bounded away from zero.*

Theorem 1 shows that there are no paths (and thus no true data-generating process) on which agent NB vanishes against agent B. Therefore, a Bayesian agent cannot drive an under-reacting agent with the same information out of the market. Agent NB survives because, regardless of the path of realizations, agent NB's beliefs are guaranteed to be at least as accurate as those of agent B (Theorem 1*, in Appendix C shows that $\lim_{t\to\infty} \frac{p^{NB}(\sigma^t)}{p^N(\sigma^t)}$ is uniformly bounded away from zero on every path, and the result follows from Proposition 3). If the learning problem is correctly specified, both agents survive because they learn the true model quickly (their beliefs merge with the truth). If the learning problem is misspecified, the Bayesian agent (generically) learns the most accurate model in its support (Berk 1966), while the NB agent's predictions are guaranteed to be at least as accurate as the most accurate model in that support.

Theorem 2 precisely characterizes the sequences on which agent B vanishes

[10] In general, having two agents with identical average accuracy does not imply that both agents survive because the average (conditional) relative entropies are too coarse to discriminate between log-likelihood ratios that diverge at rates slower than t. This problem is particularly salient in learning environments in which agent beliefs converge to the same models because the averaging factor masks differences in the converging rate (Blume and Easley 2006; Massari 2013).

Theorem 2 *Under* **A1-A3***, for all* $\alpha \in (0, 1]$*, agent* B *vanishes on every equilibrium path* σ *on which the prior process of agent* NB *is not concentrated on a single model for a positive fraction of periods.*

By Theorem 1, the NB agent's predictions are guaranteed to be at least as accurate as the most accurate model in its support. Theorem 2 shows that the NB agent's predictions are, on average, more accurate than those of agent B whenever its prior process does not converge to a Dirac measure on any single model in the support. The intuition is as follows. First notice that if NB's prior process converges then also agent B's prior process converges to the same model — because NB under-reacts to information — and the two agents are averagely equally accurate. Second, agent NB prior process does not converge when there are at least two models in the support, that if combined generate a new model $P_{\theta*}$ that is averagely more accurate than the most accurate model in the support and the under-reaction parameter is small enough.[11] Since the prior process does not concentrate, the resulting predictions are closer to $P_{\theta*}$ than the most accurate model in the support, which is averagely as accurate as the beliefs of the Bayesian agent by Berk (1966).

Now I present sufficient conditions on the parameter space, **C1**, the truth, **C2**, and the under-reacting parameter α, **C3**, such that the prior process of agent NB never converges and thus for agent B to vanish against agent NB P-a.s.. While it is easy to verify that **C1** is necessary for **C2**, which is necessary for **C3**, it is useful to present the three conditions independently because they regulate different aspects of the learning problem.

First, the prior support must allow for the existence of a probability distribution $P_{\hat{\theta}*}$ such that a combination of models is more accurate — has lower K-L divergence — than any model in the support if draws were iid with probabilities $P_{\hat{\theta}*}$.

Condition C1

$$\exists P_{\hat{\theta}*} \in \Delta^{|S|-1} : \min_{\theta \in Conv(\Theta)} d(P_{\hat{\theta}*}||\pi_\theta) < d(P_{\hat{\theta}*}||\pi_{\hat{\theta}}),$$

where $Conv(\Theta)$ *denotes the convex hull of* Θ*, and* $\hat{\theta} \in \text{argmin}_{\theta \in \Theta} \bar{d}(P_{\hat{\theta}*}||\pi_\theta)$.

A necessary and sufficient condition for **C1** is that the prior support is not convex, $Conv(\Theta) \neq \Theta$, because $Conv(\Theta) = \Theta \Leftrightarrow \forall P_{\theta*} \in \Delta^{|S|-1}$, $\min_{\theta \in Conv(\Theta)} d(P_{\hat{\theta}*}||\pi_\theta) = d(P_{\hat{\theta}*}||\pi_{\hat{\theta}})$. This requirement might seem artificial, considering that in all classroom examples of Bayesian learning, we assume a convex prior support on the space of parameters. However, **C1** holds in all prediction problems in which the prior support domain is a finite set of parameters (e.g., if $|\Theta| < \infty$, as assumed in this paper); but also if Θ is an index over a finite set of models (e.g., regressions models), which is the typical domain utilized in the portfolio selection literature (DeMiguel et al. 2007; Garlappi et al. 2006).

Second, the true probability, P, must induce P-a.s. sequences on which the empirical distribution $P_{\theta*}$ satisfies **C1**.

[11] Inspection of Definitions 2 and 3 shows that if p^{NB} prior concentrates on a unique parameter θ, then it must be the case that $\lim_{t\to\infty} p_t^{NB} = \lim_{t\to\infty} p_{\theta,t} = \pi_\theta = \lim_{t\to\infty} p_t^B$ — where the last equality is proven in Lemma 3. Thus p^{NB} is averagely as accurate as p^B whenever it concentrates on a unique model.

Condition C2

$$P: \min_{\theta\in Conv(\Theta)} d(P_{\theta^*}||\pi_\theta) < d(P_{\theta^*}||\pi_{\hat\theta}) \quad P\text{-}a.s.$$

Notably, the expected value in the definition of the K-L divergence in **C1** and **C2** is not taken with respect to the true distribution P, which may have a time-varying structure, but rather with respect to the time-invariant distributions $P_{\hat\theta^*}$ and P_{θ^*}, respectively. This substitution can be made without loss of generality (WLOG) because these conditions are instrumental in ensuring that none of the $\pi_{\hat\theta}$ in the prior support can serve as an attractor of the prior dynamics. If any were attractors of the prior dynamics, then the time structure of P would be irrelevant, since the π_θ models are time-invariant, allowing us to substitute P_{θ^*} for P when verifying the drift condition around each model WLOG.

The last condition, **C3**, requires that the NB agent under-reacts to information sufficiently to prevent its prior from ever concentrating on a single model. The minimum level of the under-reaction parameter, $\bar\alpha$, such that the prior never concentrates has an analytical form, which I indicate below. The intuition behind this result requires a characterization of the NB prior dynamics, relying on two intermediate lemmas provided in Appendix C following Theorem 2*.

Condition C3 *Condition C3 is satisfied if* $\alpha \in (0, \bar\alpha)$,

$$\textit{with } \bar\alpha := \max\left\{\alpha\in[0,1] : d(P_{\theta^*}||\pi_{\hat\theta}) \geq \min_{\theta\in\Theta\setminus\hat\theta} d(P_{\theta^*}||(1-\alpha)\pi_{\hat\theta}+\alpha\pi_\theta)\ P\text{-}a.s.\right\}.$$

Notably, if it is impossible to combine models in the support to improve the accuracy of the most accurate model–e.g., in well-specified learning problems–then $\bar\alpha = 0$ and condition **C3** cannot be satisfied. I dismiss this eventuality without loss of generality (WLOG) because I use **C3** in conjunction with **C2**, which is a necessary and sufficient condition for $\bar\alpha > 0$ (see Lemma 6).

Theorem 3 shows that if there exists a combination of models that is more accurate —has lower K-L divergence — than all models in the support (**C2**) and agent NB under-reaction parameter is below $\bar\alpha$ (**C3**) then, agent NB beliefs are averagely more accurate than agent B's (Theorem 3*) and agent B vanishes (Proposition 2).

Theorem 3 *Under* **A1-A3** *and* **C1-C3** *agent B vanishes P-a.s. against agent NB.*

Together, Theorems 1, 2, and 3 imply that financial markets do not favor Bayesian agents over non-Bayesian under-reacting agents with the same information, and that they might favor non-Bayesian under-reacting agents over Bayesian agents with the same information.

An example

In this Section, I present a simple example illustrating the relationship between survival and accuracy and how an under-reacting agent can come to dominate a Bayesian agent.

Let agents B and NB have log-utility and symmetric endowment. In every equilibrium path, it must be the case (by the FOC) that

$$\forall(t,\sigma),\ \frac{c_t^{NB}(\sigma)}{c_t^{B}(\sigma)} = \frac{p^{NB}(\sigma^t)}{p^{B}(\sigma^t)}\frac{c_0^{NB}}{c_0^{B}};$$

so, agent NB survives on path σ with consumption shares uniformly bounded away from zero if and only if, for all large t, $\frac{p^{NB}(\sigma^t)}{p^{B}(\sigma^t)} > 0$ on σ (Proposition 3). Furthermore,

$$\forall(t,\sigma),\ \frac{c_t^{B}(\sigma)}{c_t^{NB}(\sigma)} = \frac{p^{B}(\sigma^t)}{p^{NB}(\sigma^t)} = e^{\ln\frac{P(\sigma^t)}{p^{NB}(\sigma^t)} - \ln\frac{P(\sigma^t)}{p^{B}(\sigma^t)}} \underbrace{\approx}_{\text{by Lemma 3}} e^{t\left(\bar{d}(P||p^{NB}) - \bar{d}(P||p^{B})\right)};$$

so, $\bar{d}(P||p^{NB}) < \bar{d}(P||p^{B}) \Rightarrow$ agent B vanishes (Proposition 2).

Suppose there are two states $S := \{u, d\}$, that states are i.i.d with true probability $P(u) = .5 = 1 - P(d)$, that the two agents have identical time 0 prior μ_0 on the support $\Theta := \{\theta_1, \theta_2\}$ with $[\pi_{\theta_1}(u), \pi_{\theta_2}(u)] = [.4, .7]$, and that $\alpha = .2$.

Because π_{θ_1} is averagely more accurate than π_{θ_2}, the Bayesian agent is averagely as accurate as π_{θ_1}, $\bar{d}(P||p^B) = \bar{d}(P||\pi_{\theta_1})$ (Berk 1966). Conversely, the sufficient conditions of Theorem 3 are satisfied. So, agent B vanishes because agent NB's beliefs are averagely more accurate than agent B, $\bar{d}(P||p^{NB}) < \bar{d}(P||p^B)$.

Theorem 3 requires three conditions that are satisfied by example 1. Condition C1 asks for the existence of a combination of parameters in the support that is more accurate than π_{θ^1} and π_{θ^2}; in example 1, condition **C1** is satisfied by all probabilities $P(u)$ in the interval (.4,.6). Condition C2 requires the true probability to generate P-a.s. an empirical distribution in the interval (.4,.6); this condition is satisfied because, by the strong law of large numbers $P_{\theta^*}(u) = .5 \in (.4, .6)$P-a.s.. Condition C3 requires the NB agent to under-react “enough" to information. Since the parameter α is “small enough" ($\alpha = .2 < 2/3 = \bar{\alpha}$), p^{NB} does not converge to any parameter. Instead, p^{NB} spends most of its time in the interval (.4,.6), so that p^{NB} is averagely more accurate than p^B (see Fig. 1 for an illustrative simulation). Prices eventually reflect agent NB's beliefs, which are not Bayesian.

Notably, since **C2** and **C3** only depend on the empirical distribution, rather than P, the same conclusion above is valid for every P such that $P_{\theta^*} \in (.4, .6)$. For example, if P is a Markov 1 model with $P(u|d) + P(u|u) = P_{\theta^*}(u) = .5$.

4 Conclusion

Under-reaction is not a transient phenomenon in financial markets. Financial markets favor “irrational” non-Bayesian agents that under-react to information with respect to Bayes' rule over “rational” Bayesian agents because under-reacting rules are more robust to model misspecification than Bayes'.”

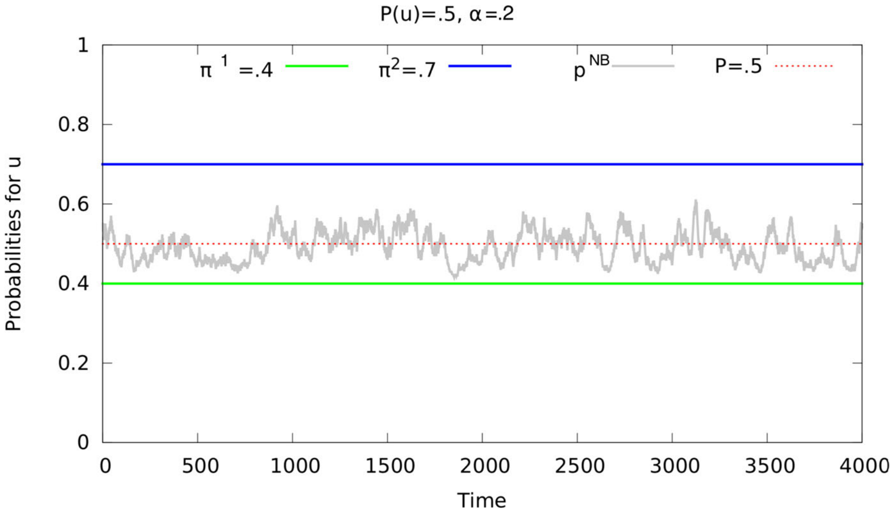

Fig. 1 p^{NB} dynamics with $[\pi_{\theta_1}(u), \pi_{\theta_2}(u)] = [.4, .7]$, $\mu_0(\theta_1) = .5$ and $\alpha = .2$. Notably, the p^{NB} predictions do not converge and remain closer to the truth than the most accurate model in the support, π_{θ_1}, in most periods

Appendices

A discussion about the relevance of under-reacting models

Focusing on conditions **C1**-**C3** help identify environments in which under-reacting non-Bayesian models can/are likely to overperform Baues' rule

- **C1** never holds in parametric inference problems in which the prior support is convex because the K-L divergence is convex.
- **C2** never holds in well-specified learning problems, because the averagely most accurate model is, by the strong law of large numbers, the true model.
- **C2** holds if Θ is finite and $\theta^* \in Conv(\Theta) \setminus \Theta$. That is, **C2** is satisfied if the true data-generating process is such that the empirical distribution has parameters that are a non-degenerate convex combination of parameters in the support. So, if states are i.i.d. and generated according to a parameter selected at random in the simplex, **C2** is satisfied on a generic, and potentially large, subset of parameters — e.g., if the economy has two states $\{H, L\}$, and $\Theta = \{.4, .7\}$, **C2** is satisfied for every $P(H) \in (.4, .7)$.
- Therefore, **C1** holds in the model mixture problems studied in the portfolio selection literature because they consider mixture over a finite set of models of stock market returns. And the evidence of sub-optimal performance of empirical and Bayesian models against robust models strongly suggests that also **C2** holds.
- Finiteness of the prior support is not necessary for **C1** to hold. For example, consider the following non-convex prior support $\Theta = \{\theta \in (0, .4) \cup (.7, 1)\}$. Condition C1 is satisfied for all $\theta^* \in (.4, .7)$.

Remark It could be argued that condition **C1** is practically irrelevant because it is possible to robustify Bayesian predictions by convexifying Θ. There are at least three strong counter-arguments to this objection.

1) To convexify Θ qualitatively slows down the learning rate achieved by Bayes' rule. If Θ is finite, a Bayesian that learns from $Conv(\Theta)$ learns at a rate that is qualitatively slower than that of a Bayesian (or an under-reacting agent) that learns from Θ.[12] In terms of selection outcomes, the following Corollary holds.

Corollary 1 *An under-reacting agent with finite prior support* Θ, *drives out of the market a Bayesian agent with prior support* $conv(\Theta)$ P_θ *almost surely for any* $\theta \in \Theta$.

Proof An implication of Blume and Easley (2006) (Theorem 6) is that a Bayesian agent with finite support Θ drives out of the market a Bayesian agent with prior support $Conv(\Theta)$ almost surely according to any model in Θ — because $Conv(\Theta)$ has positive Lebesgue measure. Conversely, Theorem 1 above shows that an under-reacting agent is never driven out of the market by a Bayesian with the same support. □

2) If Θ is an index over a finite set of models, it might not be possible to construct a proper convexification of Θ — For example, it is impossible to construct the convex hull between a forest forecast algorithm and regression model.

3) A Bayesian agent would never choose to convexify its prior support. By definition, a Bayesian attaches probability 1 to the event that the true model is one of the models in its support. Thus, he believes it impossible to be in a misspecified learning setting and has no incentive to hedge against this risk.

B Relative accuracy of p^B and p^{NB}

Here, I present results characterizing the relative accuracy of beliefs p^B and p^{NB}. Theorem 1* shows that the likelihood ratio between p^{NB} and p^B is universally (i.e., on every path) uniformly bounded away from zero; Theorem 2* shows that non-concentration of the prior process of the NB agent on any model of the support for a positive fraction of periods is a sufficient condition for p^{NB} to be averagely more accurate than p^B, and Theorem 3* shows that **C1**-**C3** are sufficient condition for p^{NB} to be averagely more accurate than p^B, because **C1**-**C3** are sufficient to guarantee the non-concentration of the prior process of the NB agent on any model of the support for a positive fraction of periods.

I start with two lemmas to aid intuition. In Lemma 1 I apply the chain rule to obtain the analytic form of the unconditional probabilities of Definitions 2 and 3. It shows that p^{NB} can be analytically rewritten as a Bayesian mixture model on path-dependent time-varying models p_θ, rather than the time-invariant models π_θ. Therefore, standard Bayesian results, such as the universal bounds on the right, hold for the Bayesian mixture model p^{NB} with support $\{p_{\theta^1}, ..., p_{\theta^K}\}$.

[12] The learning rate is qualitatively slower because the cardinality of Θ is finite, while the cardinality of $Conv(\Theta)$ is the continuum.

Lemma 1 *Let* $p_\theta(\sigma^t) := \prod_{\tau=1}^{t}\big((1-\alpha)p^{NB}(\sigma_\tau|) + \alpha\pi_\theta(\sigma_\tau)\big)$*, then, for all* $\alpha \in (0,1]$,

$$\forall(t,\sigma),\;\; p^{NB}(\sigma^t) = \sum_{\theta\in\Theta} p_\theta(\sigma^t)\mu_0(\theta) \in \left[\max_{\theta\in\Theta} p_\theta(\sigma^t)\mu_0(\theta), \max_{\theta\in\Theta} p_\theta(\sigma^t)\right]; \quad (2)$$

when $\alpha = 1$*, then* $\pi_\theta = p_\theta$ *(Definitions 2 and 3 coincide) and we have:*

$$\forall(t,\sigma),\;\; p^{B}(\sigma^t) = \sum_{\theta\in\Theta} \pi_\theta(\sigma^t)\mu_0(\theta) \in \left[\max_{\theta\in\Theta} \pi_\theta(\sigma^t)\mu_0(\theta), \max_{\theta\in\Theta} \pi_\theta(\sigma^t)\right]. \quad (3)$$

Bounds 2 and 3 highlight that the key step to discussing the relative accuracy between p^{NB} and p^B is to characterize the relative accuracy of models p_θ against models π_θ. Lemma 2 below provides this result. It shows that, on every path, and for every θ, the likelihood ratio between p_θ and π_θ is uniformly bounded away from zero.

Lemma 2

$$\forall\alpha \in (0,1], \forall(t,\sigma), \forall\theta \in \Theta,\, p_\theta(\sigma^t) \geq \left(\min_{\theta\in\Theta}\mu_0(\theta)\right)^{\frac{1}{\alpha}-1}\pi_\theta(\sigma^t).$$

The intuition is as follows: in every period p_θ is a convex combination between π_θ and p^{NB}. Taking the logarithm, the concavity of the logarithm implies that in every period, $\ln p_\theta(\sigma_t|) \geq (1-\alpha)\ln p^{NB}(\sigma_t|) + \alpha \ln \pi_\theta(\sigma_t)$. Thus, summing over t, in every path σ, $\ln p_\theta(\sigma^t) \geq (1-\alpha)\ln p^{NB}(\sigma^t) + \alpha\ln\pi_\theta(\sigma^t)$; and the result follows because p^{NB}, is at least as accurate as p_θ by Lemma 1.

To show that there are no paths on which p^B is more accurate than p^{NB} it suffices to combine the universal uniform bounds of Lemmas 1 and 2.

Theorem 1*. *For every value of* $\alpha \in (0,1]$, *the likelihood ratio between* p^{NB} *and* p^B *is uniformly bounded away from zero on every path.*

$$\forall\alpha \in (0,1], \forall(t,\sigma),\, \frac{p^{NB}(\sigma^t)}{p^B(\sigma^t)} \geq \left(\min_{\theta\in\Theta}\mu_0(\theta)\right)^{\frac{1}{\alpha}} > 0.$$

Theorem 1*'s uniform bound holds on all paths $\sigma \in \Sigma$. Thus, it bounds uniformly the maximal accuracy-cost of under-reaction. For example, when the learning problem is correctly specified, both p^B and p^{NB} converge to the truth and the likelihood ratio remains greater than zero because the convergence rate of p^B and that of p^{NB} are both exponentials in t — the difference in the convergence rate of the priors to the Dirac on the true model can be uniformly bounded by a multiplicative constant. As intuition suggests, smaller values of α correspond to higher under-reaction and thus, lower worst-case accuracy against the Bayesian agent. Clearly, the bound of Theorem 1* also holds on those paths on which $p^{NB}(\sigma^t)$ or $p^B(\sigma^t)$ posteriors do not converge. Next, I show that $p^{NB}(\sigma^t)$ is averagely more accurate than $p^B(\sigma^t)$ on all paths on which its prior process does not concentrate on a unique model for a positive fraction of periods. In order to do so, I start by generalizing some standard results of Bayesian statistics to the case of the NB Bayesian mixture model at hand. First, I use Lemma 3 to show that for every data-generating process, a Bayesian mixture is averagely as accurate as the averagely most accurate model(s) in its support.

Lemma 3 *Let* $\hat{\theta} \in \text{argmin}_{\theta\in\Theta}\, \bar{d}(P||p_\theta)$

$$\forall\alpha \in (0,1],\ \bar{d}(P||p^{NB}) = \bar{d}(P||p_{\hat{\theta}})\ \ P\text{-a.s.};$$

and, as a special case, $\bar{d}(P||p^{B}) = \bar{d}(P||\pi_{\hat{\theta}})P$*-a.s.*

Next, I characterize the relative average accuracy of $p_{\hat{\theta}}$ and $\pi_{\hat{\theta}}$.

Lemma 4 *For all* $\theta \in \Theta$ *and for all* $\alpha \in (0,1)$,

$$\bar{d}(P||p_\theta) \le \bar{d}(P||\pi_\theta),\ P\text{-a.s.}$$

with strict inequality if there exists an $\epsilon > 0$ *such that* $||p_t^{NB} - \pi_\theta|| > \epsilon$ *a positive fraction of periods* P*-a.s..*

The intuition is that $p_{\hat{\theta}}$ is averagely at least as accurate as $\pi_{\hat{\theta}}$ because the K-L divergence is strictly convex, and $\pi_{\hat{\theta}}$ is, in every period, a mixture of $\pi_{\hat{\theta}}$ and a model p^{NB}, which is guaranteed to be averagely at least as accurate as $\pi_{\hat{\theta}}$. Moreover, $p_{\hat{\theta}}$ is averagely more accurate than $\pi_{\hat{\theta}}$ on those sequences on which the p^{NB} prior process does not concentrate to any model for a positive fraction of periods. This follows because the strict convexity of the K-L divergence implies that the inequality is strict whenever $p^{NB} \neq \pi_{\hat{\theta}}$, a situation that occurs a positive fraction of periods, by assumption.

We obtain the following result by combining Lemmas 3 and 4.

Theorem 2*. p^{NB} *is averagely more accurate than* p^{B} *on all paths* σ *on which its prior process is not concentrated on a single model for a positive fraction of periods.*

Finally, I show that **C1-C3** are sufficient conditions for p^{NB} to be averagely more accurate than p^{B}.

Theorem 3*. C1-C3 $\Rightarrow \bar{d}(P||p^{B}) > \bar{d}(P||p^{NB})$ P-a.s.

Condition C1 is necessary for condition **C2**, which is necessary for condition **C3**, which is sufficient to prevent the p^{NB}'s prior from ever concentrating on a unique model (Lemma 5 in Appendix). The intuition behind the role played by **C3** is as follows. By standard Bayesian argument, a Bayesian mixture concentrates on the model on its support with the lowest average accuracy, when this model is unique (e.g., Berk 1966; Marinacci and Massari 2019).[13] However, if p^{NB}'s prior were to concentrate on a model $p_{\hat{\theta}}$, then $p^{NB} \to p_{\hat{\theta}} = \pi_{\hat{\theta}}$ and its average accuracy would be higher than the average accuracy of another model in the support (**C2**: $\Rightarrow \bar{d}(P||\pi_{\hat{\theta}}) > \min_{\theta\in\Theta\backslash\hat{\theta}} \bar{d}(P||(1-\alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta)$); a contradiction.

[13] Lemma 3 generalizes this intuition from the i.i.d. draws/time-invariant models in the prior assumptions of Berk (1966) to the arbitrary time structure in the draws/ arbitrary path-dependent models in the prior support needed for this setting.

I conclude with two corollaries. Corollary 2 shows that if the learning problem is correctly specified p^{NB} merges with the truth. That is, p^{NB} converges to the truth qualitatively as fast as p^B does. This result strengthens the *weak merging* result of Epstein et al. (2010)

Corollary 2 *For all* $\alpha \in (0,1]$, *for all* $\theta \in \Theta$, p^{NB} *merges with* π_θ, π_θ*–a.s.*

Corollary 3 verifies that under the conditions of Theorem 2*, the likelihood ratio between p^B and p^{NB} converges to zero P-a.s.

Corollary 3 *Under the conditions of Theorem 2*,* $\frac{p^B(\sigma^t)}{p^{NB}(\sigma^t)} \to 0$ *P-a.s.*

B.1 Proofs of results in Appendix B

Proof of Lemma 1 Note that for all (t,σ)

$$p^{NB}(\sigma_t|) := \sum_{\theta\in\Theta} \pi_\theta(\sigma_t)\mu^{NB}(\theta|\sigma^{t-1})$$

can be equivalently rewritten as

$$p^{NB}(\sigma_t|) = \sum_{\theta\in\Theta} p_\theta(\sigma_t|)\mu^{NB}(\theta|\sigma^{t-1}),$$

and recognize the Bayesian mixture dynamic with respect to the model class $\{p_\theta : \theta \in \Theta\}$. Specifically, for all $\alpha \in (0,1]$,

$$\begin{aligned}
&\forall(t,\sigma),\ p^{NB}(\sigma^t)\\
&\quad := \prod_{\tau=1}^{t} p^{NB}(\sigma_\tau|\sigma^{\tau-1})\\
&\quad = \left(\sum_{\theta\in\Theta} p_\theta(\sigma_t|\sigma^{t-1})\mu^{NB}(\theta|\sigma^{t-1})\right)\prod_{\tau=1}^{t-1} p^{NB}(\sigma_\tau|\sigma^{\tau-1})\\
&\quad \underbrace{=}_{\text{by Def.3}} \left(\sum_{\theta\in\Theta} p_\theta(\sigma_t|\sigma^{t-1})p_\theta(\sigma_{t-1}|\sigma^{t-2})\mu^{NB}(\theta|\sigma^{t-2})\right)\\
&\quad \times \frac{1}{p^{NB}(\sigma_{t-1}|\sigma^{t-2})}\prod_{\tau=1}^{t-1} p^{NB}(\sigma_\tau|\sigma^{\tau-1})\\
&\quad = \left(\sum_{\theta\in\Theta} p_\theta(\sigma_t|\sigma^{t-1})p_\theta(\sigma_{t-1}|\sigma^{t-2})\mu^{NB}(\theta|\sigma^{t-2})\right)\prod_{\tau=1}^{t-2} p^{NB}(\sigma_\tau|\sigma^{\tau-1})\\
&\quad \vdots
\end{aligned}$$

$$
\begin{aligned}
&= \sum_{\theta \in \Theta} \prod_{\tau=1}^{t} p_\theta(\sigma_\tau|\sigma^{\tau-1})\mu_0(\theta) \\
&= \sum_{\theta \in \Theta} p_\theta(\sigma^t)\mu_0(\theta) \in \left[\max_{\theta \in \Theta} p_\theta(\sigma^t)\mu_0(\theta), \max_{\theta \in \Theta} p_\theta(\sigma^t)\right]
\end{aligned}
$$

Finally, note that $\alpha = 1 \Rightarrow \forall(t, \sigma),\ p_t^B = p_t^{NB}$ and $\forall\theta \in \Theta,\ p_{\theta,t} = \pi_\theta$.

$$
\text{So, } \forall(t,\sigma),\ p^B(\sigma^t) = \sum_{\theta \in \Theta} \pi_\theta(\sigma^t)\mu_0(\theta) \in \left[\max_{\theta \in \Theta} \pi_\theta(\sigma^t)\mu_0(\theta), \max_{\theta \in \Theta} \pi_\theta(\sigma^t)\right]
$$

□

Proof of Lemma 2

$$
\begin{aligned}
\forall(t,\sigma), \forall\theta \in \Theta, \quad \ln p_\theta(\sigma^t) &= \ln \prod_{\tau=1}^{t} \Big((1-\alpha)p^{NB}(\sigma_\tau|) + \alpha\pi_\theta(\sigma_\tau)\Big) \\
&= \sum_{\tau=1}^{t} \ln \Big((1-\alpha)p^{NB}(\sigma_\tau|) + \alpha\pi_\theta(\sigma_\tau)\Big) \\
&\underbrace{\geq}_{\text{by concavity of } \ln(\cdot)} (1-\alpha)\sum_{\tau=1}^{t} \ln p^{NB}(\sigma_\tau|) \\
&\quad + \alpha \sum_{\tau=1}^{t} \ln \pi_\theta(\sigma_\tau) \\
&= (1-\alpha)\ln p^{NB}(\sigma^t) + \alpha \ln \pi_\theta(\sigma^t) \\
\Rightarrow \forall(t,\sigma), \forall\theta \in \Theta,\ \ln \frac{p_\theta(\sigma^t)}{\pi_\theta(\sigma^t)} &\geq \frac{(1-\alpha)}{\alpha} \ln \frac{p^{NB}(\sigma^t)}{p_\theta(\sigma^t)} \\
&\underbrace{=}_{\text{by Eq. (2)}} \frac{(1-\alpha)}{\alpha} \ln \frac{\sum_{\theta' \in \Theta} p_{\theta'}(\sigma^t)\mu_0(\theta')}{p_\theta(\sigma^t)} \\
&\geq \ln \left(\min_{\theta \in \Theta} \mu_0(\theta)\right)^{\frac{(1-\alpha)}{\alpha}} \\
\Rightarrow \forall(t,\sigma), \forall\theta \in \Theta, \quad p_\theta(\sigma^t) &\geq \left(\min_{\theta \in \Theta} \mu_0(\theta)\right)^{\frac{1}{\alpha}-1} \pi_\theta(\sigma^t).
\end{aligned}
$$

□

Proof of Theorem 1* The result follows from Lemmas 1 and 2

$$
\forall(t,\sigma),\ \frac{p^{NB}(\sigma^t)}{p^B(\sigma^t)} \underbrace{=}_{\text{by Lemma 1}} \frac{\sum_{\theta \in \Theta} p_\theta(\sigma^t)\mu_0(\theta)}{\sum_{\theta' \in \Theta} \pi_{\theta'}(\sigma^t)\mu_0(\theta')}
$$

$$
\begin{aligned}
&\geq \quad \frac{\max_{\theta\in\Theta} p_\theta(\sigma^t)\mu_0(\theta)}{\max_{\theta'\in\Theta}\pi_{\theta'}(\sigma^t)}\\
&\underbrace{\geq}_{\text{by Lemma 2}} \quad \frac{\max_{\theta\in\Theta}\pi_\theta(\sigma^t)\mu_0(\theta)}{\max_{\theta'\in\Theta}\pi_{\theta'}(\sigma^t)}\left(\min_{\theta''\in\Theta}\mu_0(\theta'')\right)^{\frac{1}{\alpha}-1}\\
&\geq \quad \frac{\max_{\theta\in\Theta}\pi_\theta(\sigma^t)\min_{\theta'''\in\Theta}\mu_0(\theta''')}{\max_{\theta'\in\Theta}\pi_{\theta'}(\sigma^t)}\left(\min_{\theta''\in\Theta}\mu_0(\theta'')\right)^{\frac{1}{\alpha}-1}\\
&\geq \quad \left(\min_{\theta\in\Theta}\mu_0(\theta)\right)^{\frac{1}{\alpha}} > 0.
\end{aligned}
$$

Proof of Lemma 3 Let $\hat{\Theta}(\sigma^t)$ be the $\operatorname{argmax}_{\theta^*\in\Theta} p_{\theta^*}(\sigma^t)$, and assume WLOG that is a singleton: $\hat{\Theta}(\sigma^t)=\hat{\theta}(\sigma^t)$.[14]

$$
\begin{aligned}
\forall(t,\sigma),\quad \frac{p_{\hat{\theta}(\sigma^t)}(\sigma^t)}{p^{NB}(\sigma^t)} &\underbrace{=}_{\text{By Eq.(2)}} \frac{p_{\hat{\theta}(\sigma^t)}(\sigma^t)}{\sum_{\theta\in\Theta} p_\theta(\sigma^t)\mu_0(\theta)}\\
&= \frac{1}{\mu_0(\hat{\theta})+\sum_{\theta\in\Theta\backslash\hat{\theta}}\frac{p_\theta(\sigma^t)}{p_{\hat{\theta}(\sigma^t)}(\sigma^t)}\mu_0(\theta)}\\
&= \frac{1}{\mu_0(\hat{\theta}(\sigma^t))+\sum_{\theta\in\Theta\backslash\hat{\theta}} e^{\ln\frac{p_\theta(\sigma^t)}{p_{\hat{\theta}(\sigma^t)}(\sigma^t)}}\mu_0(\theta)}\\
&\Rightarrow \begin{cases}\limsup_{(\sigma\in\Sigma,t)}\frac{p_{\hat{\theta}(\sigma^t)}(\sigma^t)}{p^{NB}(\sigma^t)} \leq \frac{1}{\min_{\theta\in\Theta}\mu_0(\theta)}\\ \liminf_{(\sigma\in\Sigma,t)}\frac{p_{\hat{\theta}(\sigma^t)}(\sigma^t)}{p^{NB}(\sigma^t)} \geq 1\end{cases}\\
&\Rightarrow \begin{cases}\limsup_{(\sigma\in\Sigma,t)}\ln\frac{p_{\hat{\theta}(\sigma^t)}(\sigma^t)}{p^{NB}(\sigma^t)} \leq \ln\frac{1}{\min_{\theta\in\Theta}\mu_0(\theta)}\\ \liminf_{(\sigma\in\Sigma,t)}\ln\frac{p_{\hat{\theta}(\sigma^t)}(\sigma^t)}{p^{NB}(\sigma^t)} \geq 0\end{cases}
\end{aligned}
$$

$$
\Rightarrow \forall(t,\sigma),\ \lim_{t\to\infty}\frac{1}{t}\left(\ln\frac{p_{\hat{\theta}(\sigma^t)}(\sigma^t)}{p^{NB}(\sigma^t)}\right)=0
$$

$$
\Rightarrow \forall(t,\sigma),\ \lim_{t\to\infty}\frac{1}{t}\left(\ln\frac{P(\sigma^t)}{p^{NB}(\sigma^t)}-\ln\frac{P(\sigma^t)}{p_{\hat{\theta}(\sigma^t)}(\sigma^t)}\right)=0
$$

$$
\begin{aligned}
\Rightarrow \lim_{t\to\infty}\left[\frac{1}{t}\left[\sum_{\tau=1}^t\ln\frac{P(\sigma_\tau)}{p^{NB}(\sigma_\tau|)}-\sum_{\tau=1}^t d_\tau(P||p^{NB})\right]+\frac{1}{t}\sum_{\tau=1}^t d_\tau(P||p^{NB})\right]\\
= \lim_{t\to\infty}\left[\frac{1}{t}\left[\sum_{\tau=1}^t\ln\frac{P(\sigma_\tau)}{p_{\hat{\theta}(\sigma^t)}(\sigma_\tau|)}-\sum_{\tau=1}^t d_\tau(P||p_{\hat{\theta}(\sigma^t)})\right]\right.\\
\left.+\frac{1}{t}\sum_{\tau=1}^t d_\tau(P||p_{\hat{\theta}(\sigma^t)})\right]
\end{aligned}
$$

[14] This assumption can be disposed of by re-defininig $\mu_0(\hat{\theta}) := \sum_{\hat{\hat{\theta}}\in\hat{\Theta}}\mu_0(\hat{\hat{\theta}})$

$$\Rightarrow \qquad \bar{d}(P||p^{NB}) \;= \bar{d}(P||p_{\hat{\theta}}) \quad P\text{-a.s.}.$$

The last implication follows from the strong law of large number for martingale differences (see also Sandroni 2000) that, in our setting, guarantees that for $\rho = p^{NB}$ (for the LHS), and $\rho = p_{\hat{\theta}(\sigma^t)}$ (for the RHS)

$$\lim_{t\to\infty} \frac{1}{t}\left[\sum_{\tau=1}^{t} \ln \frac{P(\sigma_\tau)}{\rho(\sigma_\tau|)} - \sum_{\tau=1}^{t} d_\tau(P||\rho)\right] = 0,\ P\text{-a.s.}$$

I suppressed the dependence on σ^t, because $\bar{d}(\cdot||\cdot)$ exists by assumption and the equality holds for all models in $\{\hat{\theta} \in \theta\, \text{argmin}_{\theta^*\in\Theta}\, \bar{d}(P||p_{\theta^*})\} \subset \{\theta : p_\theta \in \text{argmax}_{\theta^*\in\Theta}\, p_{\theta^*}(\sigma^t) \text{ infinitely often}\}$. □

Proof of Lemma 4 • Weak inequality: for all, (t, σ), for all $\theta \in \Theta$,and for all $\alpha \in (0, 1)$,

$$\begin{aligned}
d_t(P||p_\theta) &= d_t(P||(1-\alpha)p^{NB} + \alpha\pi_\theta)) \\
&\leq^{(a)} (1-\alpha)d_t(P||p^{NB}) + \alpha d_t(P||\pi_\theta) \qquad \text{; by strict convexity of } d_t(P||\cdot) \\
\Rightarrow \quad \bar{d}(P||p_\theta) &\leq (1-\alpha)\bar{d}(P||p^{NB}) + \alpha\bar{d}(P||\pi_\theta) \qquad \text{; summing and averaging over } t \\
\Rightarrow \quad \bar{d}(P||p_\theta) &\leq \bar{d}(P||\pi_\theta)\ P\text{-a.s.} \qquad \text{; because } \bar{d}(P||p^{NB}) \underbrace{\leq}_{\text{by Lemma 3}} \bar{d}(P||p_\theta)
\end{aligned}$$

- Strict inequality:
  By continuity and strict convexity of $d(P||\cdot)$

$$\begin{aligned}
||p_t^{NB} - \pi_\theta|| > \epsilon \Rightarrow \exists\delta > 0 : d_t(P||(1-\alpha)p^{NB} + \alpha\pi_\theta)) \\
< (1-\alpha)d_t(P||p^{NB}) + \alpha d_t(P||\pi_\theta) - \delta.
\end{aligned}$$

  Thus, if there exists an $\epsilon > 0$ and an $\eta > 0 : \liminf_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} I_{||p_t^{NB}-\pi_\theta||>\epsilon} > \eta$ then the inequalities (a) are strict for a positive fraction of periods η, which, by definition, implies that $\bar{d}(P||p_\theta) \leq \bar{d}(P||\pi_\theta) - \eta\delta < \bar{d}(P||\pi_\theta)$.

□

Proof of Theorem 2*

$$\bar{d}(P||p^{NB}) \underbrace{=}_{By\ Lemma\ 3} \bar{d}(P||p_{\hat{\theta}}) \underbrace{<}_{By\ Lemma\ 4} \bar{d}(P||\pi_{\hat{\theta}}) \underbrace{=}_{By\ Lemma\ 3} \bar{d}(P||p^B)$$

□

Lemma 5

Condition C3 $\Rightarrow \exists\epsilon > 0 : ||p_t^{NB} - \pi_{\hat{\theta}}|| > \epsilon$ *a positive fraction of periods, P-a.s.*

Proof

For $\alpha < \bar{\alpha}$, **C3** implies $d(P_{\theta*}||\pi_{\hat{\theta}}) > \min_{\theta \in \Theta \backslash \hat{\theta}} d(P_{\theta*}||(1-\alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta)$ P-a.s.

I prove the contrapositive statement

$$\begin{aligned}||p_t^{NB} - \pi_{\hat{\theta}}|| = 0 \text{ in most periods } P\text{-a.s.} \Rightarrow d(P_{\theta*}||\pi_{\hat{\theta}}) \\ \leq \min_{\theta \in \Theta \backslash \hat{\theta}} d(P_{\theta*}||(1-\alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta) \ P\text{-a.s.}\end{aligned}$$

in four steps.
First:

$$\begin{aligned}&||p_t^{NB} - \pi_{\hat{\theta}}|| = 0 \text{ in most periods } P\text{-a.s.} \\ &\underbrace{\Rightarrow}_{\text{by continuity}} \quad |d_t(P||\pi_{\hat{\theta}}) - d_t(P||p^{NB})| = 0 \text{ in most periods} \\ &\Rightarrow \quad \bar{d}(P||\pi_{\hat{\theta}}) = \bar{d}(P||p^{NB}) \end{aligned} \tag{4}$$

Second:

$$\begin{aligned}&||p_t^{NB} - \pi_{\hat{\theta}}|| = 0 \text{ in most periods } P\text{-a.s.} \\ &\Rightarrow \quad |\min_{\theta \in \Theta \backslash \hat{\theta}} d_t(P||(1-\alpha)p^{NB} + \alpha\pi_\theta) - \min_{\theta \in \Theta \backslash \hat{\theta}} d_t(P||(1-\alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta)| = 0 \text{ in most periods } P\text{-a.s.} \\ &\underbrace{\Rightarrow}_{\text{by definition}} \quad |\min_{\theta \in \Theta \backslash \hat{\theta}} d_t(P||p_\theta) - \min_{\theta \in \Theta \backslash \hat{\theta}} d_t(P||(1-\alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta)| = 0 \text{ in most periods } P\text{-a.s.} \\ &\Rightarrow \quad \min_{\theta \in \Theta \backslash \hat{\theta}} \bar{d}(P||p_\theta) = \min_{\theta \in \Theta \backslash \hat{\theta}} \bar{d}(P||(1-\alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta)\end{aligned} \tag{5}$$

Third:

$$\begin{aligned}\bar{d}(P||\pi_{\hat{\theta}}) &\underbrace{=}_{\text{Eq.(4)}} \bar{d}(P||p^{NB}) \underbrace{=}_{\text{by Lem.3}} \min_{\theta \in \Theta} \bar{d}(P||p_\theta) \leq \min_{\theta \in \Theta \backslash \hat{\theta}} \bar{d}(P||p_\theta) \\ &\underbrace{=}_{\text{Eq.(5)}} \min_{\theta \in \Theta \backslash \hat{\theta}} \bar{d}(P||(1-\alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta)\end{aligned} \tag{6}$$

Fourth, I have to move from the true probability to the empirical frequencies (i.e. from P to $P_{\theta*}$) and from an empirical measure of accuracy $\bar{d}(\cdot||\cdot)$ to a static one $d(\cdot||\cdot)$. Note that all $\pi_\theta \in \Theta$ are all independent of t by assumption and $P_{\theta*}$ is independent of t by construction, since, for all $s_k \in S$, $P_{\theta*}(s_k) := \lim_{t\to\infty} \frac{1}{t}\sum_{\tau=1}^{t} I_{\sigma_\tau = s_k}$. Thus, for all $\sigma : ||p_t^{NB} - \pi_{\hat{\theta}}|| = 0$ in most periods Equation 6 implies

$$\bar{d}(P||\pi_{\hat{\theta}}) \leq \bar{d}(P||(1-\alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta)$$

$$\Leftrightarrow \lim_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^{t} \ln \frac{P(\sigma_\tau)}{\pi_{\hat{\theta}}(\sigma_\tau)} \leq \lim_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^{t} \ln \frac{P(\sigma_\tau)}{(1-\alpha)\pi_{\hat{\theta}}(\sigma_\tau) + \alpha\pi_\theta(\sigma_\tau)}$$

$$\Leftrightarrow \lim_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^{t} \ln \frac{(1-\alpha)\pi_{\hat{\theta}}(\sigma_\tau) + \alpha\pi_\theta(\sigma_\tau)}{\pi_{\hat{\theta}}(\sigma_\tau)} \leq 0$$

$$\Leftrightarrow \lim_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^{t} \sum_{s_j \in S} I_{\{\sigma_t = s_j\}} \ln \frac{(1-\alpha)\pi_{\hat{\theta}}(s_j) + \alpha\pi_\theta(s_j)}{\pi_{\hat{\theta}}(s_j)} \leq 0$$

$$\underbrace{\Rightarrow}_{\text{By the SLLN}} E_{P_{\theta*}} \ln \frac{(1-\alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta}{\pi_{\hat{\theta}}} \leq 0$$

$$\Leftrightarrow E_{P_{\theta*}} \ln \frac{P_{\theta*}}{\pi_{\hat{\theta}}} \leq E_{P_{\theta*}} \ln \frac{P_{\theta*}}{(1-\alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta}$$

$$\Leftrightarrow d(P_{\theta*}||\pi_{\hat{\theta}}) \leq d(P_{\theta*}||(1-\alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta)$$

where all the double implications are by definition. □

Proof of Theorem 3* By Lemma 5 **C3** $\Rightarrow \exists\epsilon > 0 : ||p_t^{NB} - \pi_{\hat{\theta}}|| > \epsilon$ a positive fraction of periods, P-a.s. and the result follows from Theorem 2*. □

Lemma 6

$$\textit{Condition C2 holds} \Leftrightarrow \underset{\alpha\in[0,1]}{argmax} \min_{\theta\in\Theta\backslash\hat{\theta}} d(P_{\theta*}||(1-\alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta) := \bar{\alpha} > 0$$

Proof Strict convexity of $d(\cdot||\cdot)$ guarantees that

$$\min_{\theta\in Conv(\Theta)} d(P_{\theta*}||\pi_\theta) < d(P_{\theta*}||\pi_{\hat{\theta}}) \Rightarrow \exists w \in int(\Delta^{\Theta}) :$$
$$\sum_k w^k\theta^k = \tilde{\theta} \text{ and } d(P_{\theta*}||\pi_{\tilde{\theta}}) < d(P_{\theta*}||\pi_{\hat{\theta}}).$$

Furthermore, it also guarantees that, locally (i.e., for small α), the accuracy of each parameter ($\hat{\theta}$ in our case) can be strictly improved by moving in the direction of one other parameter.

The reverse implication is obtained fixing an $\bar{\alpha}' \in (0, \bar{\alpha})$ giving weights only to $\hat{\theta}$ and models satisfying the RHS condition with $\bar{\alpha}'$. □

Proof of Corollary 2

$$\theta \in \Theta \Rightarrow \pi_\theta \text{ is absolutely continuous with respect to } p^B$$
$$\Rightarrow p^B \text{ merges with } \pi_\theta, \pi_\theta\text{-a.s.;}$$

$$\forall\sigma, \frac{p^{NB}(\sigma^t)}{p^B(\sigma^t)} \underbrace{>}_{\text{by Theorem 1*}} 0 \Rightarrow p^B(\sigma^t) \text{ is absolutely continuous with respect to } p^{NB}, p^B\text{-a.s.;}$$

$$\text{thus, } \theta \in \Theta \Rightarrow \pi_\theta \text{ is absolutely continuous with respect to } p^{NB}$$

$$\Rightarrow p^{NB} \text{merges with } \pi_\theta,\ \pi_\theta\text{-a.s.}$$

□

Proof of Corollary 3 By Theorem 2*, $\bar{d}(P||p^B) > \bar{d}(P||p^{NB})$ P-a.s..
By the strong law of large number for martingale differences (see the application in Lemma 3)

$$\bar{d}(P||p^B) > \bar{d}(P||p^{NB}) \ P\text{-a.s.} \Rightarrow \frac{p^B(\sigma^t)}{p^{NB}(\sigma^t)} \to 0 \ P\text{-a.s.}$$

□

C Proofs

Proof of Proposition 1 By Definition 3,

$$\begin{aligned}
\forall(\sigma,t), \forall\theta\in\Theta, \mu^{NB}(\theta|\sigma^t) &= \mu^{NB}(\theta|\sigma^{t-1})\frac{p_\theta(\sigma_t|)}{p^{NB}(\sigma_t|)} \\
&= \mu^{NB}(\theta|\sigma^{t-1})\left(\frac{(1-\alpha)p^{NB}(\sigma_t|)+\alpha\pi_\theta(\sigma_t)}{p^{NB}(\sigma_t|)}\right) \\
&= (1-\alpha)\mu^{NB}(\theta|\sigma^{t-1})+\alpha\mu^{NB}(\theta|\sigma^{t-1})\left(\frac{\pi_\theta(\sigma_t)}{p^{NB}(\sigma_t|)}\right) \\
&= (1-\alpha)\mu^{NB}(\theta|\sigma^{t-1})+\alpha\mu^{NB}(\theta|\sigma^{t-1})\left(\frac{\pi_\theta(\sigma_t)}{\sum_\theta \pi_\theta(\sigma_t)\mu^{NB}(\theta|\sigma^{t-1})}\right).
\end{aligned}$$

By assumption, $\mu^{NB}(\theta|\sigma^{t-1}) = \mu^B(\theta|\sigma^{t-1})$; substituting on the right hand side of the equation, we obtain:

$$\begin{aligned}
\forall(\sigma,t), \forall\theta\in\Theta,\ \mu^{NB}(\theta|\sigma^t) &= (1-\alpha)\mu^B(\theta|\sigma^{t-1})+\alpha\mu^B(\theta|\sigma^{t-1}) \\
&\quad\times\left(\frac{\pi_\theta(\sigma_t)}{\sum_\theta \pi_\theta(\sigma_t)\mu^B(\theta|\sigma^{t-1})}\right) \\
&= (1-\alpha)\mu^B(\theta|\sigma^{t-1})+\alpha\mu^B(\theta|\sigma^t).
\end{aligned}$$

Where the last equality holds by definition of Bayes' rule (Definition 2):

$$\mu^B(\theta|\sigma^t) := \frac{\pi_\theta(\sigma_t)}{\sum_\theta \pi_\theta(\sigma_t)\mu(\theta|\sigma^{t-1})}\mu^B(\theta|\sigma^{t-1}).$$

□

Proof of Proposition 3 The Lagrangian problem associated with each trader's maximization problem is

$$L_i = E_{p^i} \sum_{t=0}^{\infty} \beta^t u^i(c_t^i(\sigma)) + \lambda_i \left(\sum_{t=0} \sum_{\sigma^t \in S^t} q(\sigma^t) \left(c_t^i(\sigma) - e_t^i(\sigma) \right) \right).$$

By equating the derivatives of this Lagrangian to 0, I get, for all (t, σ),

$$\frac{\partial L_i}{\partial c_t^i(\sigma)} = 0 \Rightarrow \beta^t p^i(\sigma^t) u^i(c_t^i(\sigma))' = \lambda_i q(\sigma^t)$$

Letting $q_0 = 1$ (the price of one unit of consumption at t=0 equals 1) I find that $\lambda_i = u^i(c_0^i)'$. Thus, on every equilibrium path σ^t,

$$\frac{u^B(c_t^B(\sigma))'}{u^{NB}(c_t^{NB}(\sigma))'} = \frac{p^{NB}(\sigma^t)}{p^B(\sigma^t)} \frac{u^B(c_0^B)'}{u^{NB}(c_0^{NB})'} \tag{7}$$

- Now I show that

$$\exists \eta > 0 : \forall t, \ \frac{p^{NB}(\sigma^t)}{p^B(\sigma^t)} > \eta \text{ on path } \sigma \ \Rightarrow \exists \eta' > 0 : \forall t, \ c_t^{NB} > \eta' \text{ on path } \sigma$$

$$\begin{aligned}
\forall t, \ \frac{p^{NB}(\sigma^t)}{p^B(\sigma^t)} > \eta \text{ on path } \sigma \ &\Rightarrow \ \forall t, \ \frac{u^B(c_t^B(\sigma))'}{u^{NB}(c_t^{NB}(\sigma))'} > 0 \ \text{ on path } \sigma \\
&\Rightarrow \ \forall t, \ u^{NB}(c_t^{NB}(\sigma))' < \infty \text{ on path } \sigma \\
&\underbrace{\Rightarrow}_{\text{by } \mathbf{A1}} \exists \eta' > 0 : \forall t, \ c_t^{NB} > \eta' \text{ on path } \sigma.
\end{aligned}$$

- and its reverse implication

$$\exists \eta > 0 : \forall t, \ \frac{p^{NB}(\sigma^t)}{p^B(\sigma^t)} > \eta \text{ on path } \sigma \ \Leftarrow \exists \eta' > 0 : \forall t, \ c_t^{NB} > \eta' \text{ on path } \sigma$$

$$\begin{aligned}
\forall t, \ c_t^{NB} > \eta' \text{ on path } \sigma \ &\underbrace{\Rightarrow}_{\text{by } \mathbf{A1}} \forall t, \ u^{NB}(c_t^{NB}(\sigma))' < \infty \text{ on path } \sigma \\
&\underbrace{\Rightarrow}_{\text{by } \mathbf{A2}} \forall t, \ \frac{u^B(c_t^B(\sigma))'}{u^{NB}(c_t^{NB}(\sigma))'} > 0 \ \text{ on path } \sigma \\
&\Rightarrow \ \exists \eta > 0 : \forall t, \ \frac{p^{NB}(\sigma^t)}{p^B(\sigma^t)} > \eta \text{ on path } \sigma.
\end{aligned}$$

□

Proof of Theorem 1 By Theorem 1*, $\forall(t,\sigma)$, $\frac{p^{NB}(\sigma^t)}{p^B(\sigma^t)} \geq (\min_{\theta\in\Theta}\mu_0(\theta))^{\frac{1}{\alpha}}$. Thus, the sufficient condition for an agent to survive with positive consumption-share of Proposition 3 is satisfied by agent NB on every path. □

Proof of Theorem 3 By Proposition 2, under **A1**- **A3** agent B vanishes if $\bar{d}(P||p^{NB}) < \bar{d}(P||p^B)$, P-a.s.;
By Theorem 2*, **C1**- **C3** $\Rightarrow \bar{d}(P||p^{NB}) < \bar{d}(P||p^B)$, P-a.s.. □

Funding Open access funding provided by Alma Mater Studiorum - Università di Bologna within the CRUI-CARE Agreement.

Declarations

Conflict Of Interest The author, Filippo Massari, certifies that he has no affiliation with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript.

References

Barberis, N., Shleifer, A., Vishny, R.: A model of investor sentiment. J. Financ. Econ. **49**(3), 307–343 (1998)
Benjamin, D.J.: Errors in probabilistic reasoning and judgment biases. Handbook of Behavioral Economics: Applications and Foundations **1**(2), 69–186 (2019)
Benjamin, D.J., Rabin, M., Raymond, C.: A model of nonbelief in the law of large numbers. J. Eur. Econ. Assoc. **14**(2), 515–544 (2016)
Berk, R.H.: Limiting behavior of posterior distributions when the model is incorrect. Ann. Math. Stat. **37**(1), 51–58 (1966)
Blume, L., Easley, D.: If you are so smart why aren't you rich? Belief selection in complete and incomplete markets. Econometrica **74**, 929–966 (2006)
Bohren, J. A., Hauser, D. N.: The behavioral foundations of model misspecification: A decomposition. Penn Institute for Economic Research, Department of Economics, University of Pennsylvania. (2023)
DeMiguel, V., Garlappi, L., Uppal, R.: Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? The Review of Financial Studies **22**(5), 1915–1953 (2007)
Dindo, P., Massari, F.: The wisdom of the crowd in dynamic economies. Theor. Econ. **15**(4), 1627–1668 (2020)
Epstein, L., Noor, J., Sandroni, A.: Non-bayesian updating: a theoretical framework. Theor. Econ. **3**, 193–229 (2008)
Epstein, L., Noor, J., Sandroni, A.: Non-bayesian learning. The B.E. Journal of Theoretical Economics, 10:Article 3, (2010)
Epstein, L.G.: An axiomatic model of non-bayesian updating. Rev. Econ. Stud. **73**(2), 413–436 (2006)
Frost, P.A., Savarino, J.E.: An empirical bayes approach to efficient portfolio selection. Journal of Financial and Quantitative Analysis **21**(3), 293–305 (1986)
Garlappi, L., Uppal, R., Wang, T.: Portfolio selection with parameter and model uncertainty: A multi-prior approach. The Review of Financial Studies **20**(1), 41–81 (2006)

Giglio, S., Kelly, B.: Excess volatility: Beyond discount rates. Q. J. Econ. **133**(1), 71–127 (2018)
Grünwald, P.: The safe bayesian. In Algorithmic Learning Theory, pages 169–183. Springer, (2012).
Grünwald, P., Langford, J.: Suboptimal behavior of bayes and mdl in classification under misspecification. Mach. Learn. **66**(2–3), 119–149 (2007)
Grünwald, P., Van Ommen, T., et al.: Inconsistency of bayesian inference for misspecified linear models, and a proposal for repairing it. Bayesian Anal. **12**(4), 1069–1103 (2017)
Kalai, E., Lehrer, E.: Weak and strong merging of opinions. J. Math. Econ. **23**(1), 73–86 (1994)
Klein, R.W., Bawa, V.S.: The effect of estimation risk on optimal portfolio choice. J. Financ. Econ. **3**(3), 215–231 (1976)
Marinacci, M., Massari, F.: Learning from ambiguous and misspecified models. J. Math. Econ. **84**, 144–149 (2019)
Massari, F.: Comment on if you're so smart, why aren't you rich? belief selection in complete and incomplete markets. Econometrica **81**(2), 849–851 (2013)
Massari, F.: Markets with heterogeneous beliefs: A necessary and sufficient condition for a trader to vanish. J. Econ. Dyn. Control **78**, 190–205 (2017)
Massari, F.: Price probabilities: A class of bayesian and non-bayesian prediction rules. Econ. Theor. **72**(1), 133–166 (2021)
Peleg, B., Yaari, M.E.: Markets with countably many commodities. Int. Econ. Rev. **11**(3), 369–377 (1970)
Sandroni, A.: Do markets favor agents able to make accurate predictions. Econometrica **68**(6), 1303–1341 (2000)
Sandroni, A.: Efficient markets and bayes' rule. Econ. Theor. **26**, 741–764 (2005)
Wald, A.: An essentially complete class of admissible decision functions. Ann. Math. Stat. **18**(4), 549–555 (1947)