

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

Filippo Massari\* *Bocconi University and University of East Anglia*

July 10, 2020

## Abstract

Under-reaction is a robust response to model misspecification rewarded by financial markets, rather than an “irrational” attitude that leads to extinction. I show that under-reacting rules guarantee predictions as accurate as Bayes’ in well-specified learning problems and beat Bayes’ in many misspecified learning environments. Therefore, when a Bayesian agent trades with an under-reacting agent with the same information, there are no paths on which the under-reacting agent loses all his wealth against the Bayesian, while there is a large class of misspecified learning settings in which the Bayesian agent loses all his wealth against the under-reacting agent.

*Keywords:* misspecified learning, under-reaction, 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 out of the market “irrational” non-Bayesian agents who use the same information and make asymptotically different predictions.

Theoretically, this conjecture has been investigated and confirmed in well-specified learning problems (Sandroni, 2000; Blume and Easley, 2006), i.e., under the assumption that the Bayesian agents eventually learn the truth. However, little work allowing for misspecified learning environments has been done.<sup>1</sup>

A learning problem is misspecified if the true data generating problem cannot be learned because it is subjectively believed impossible by the agent. While well-specified decision problems are theoretically appealing, we have to recognize that many, if not

---

\*E-mail: massari3141@gmail.com.

<sup>1</sup>The work of (Sandroni, 2005) lies between the well-specified and misspecified settings because it focuses on a specific class of learning problems in which the Bayesian agent cannot learn the truth because the true data generating process is not kept the same over time, but a version of Wald (1947) complete class theorem holds.

most, real-world decision problems are misspecified — “*all models are wrong, but some are useful*” *Box (1976)*; sure enough, the absence of a consensus on what is the “true” model of stock market returns indicates that investment decisions are made in a misspecified learning environment. In a misspecified-learning environment, the Bayesian paradigm is not justified because internal consistency with an incorrect premise (a model support that does not contain the truth) cannot guarantee learning the most useful model in the support (Timmermann, 2006; Grünwald et al., 2017; Massari, 2019; Csaba and Szoke, 2018); so, it is crucial to verify if Bayes’ rule should have the predominant role among all learning rules in potentially misspecified learning problem.

In this paper I focus on under-reaction (Epstein, 2006; Epstein et al., 2008), an heuristic widely documented experimentally and in empirical finance (Barberis et al., 1998; Cutler et al., 1991; Giglio and Kelly, 2018). I show that under-reacting prediction schemes guarantee predictions as accurate as Bayes’ in well-specified learning problems and beat Bayes’ in many misspecified learning environments. Then, I make use of a theoretical model to show that under-reaction reflects a prudent attitude rewarded by financial markets, rather than a “transient and irrational” behaviour. Given a Bayesian and an under-reacting agent with the same information, there are no paths on which the under-reacting agent loses all his wealth against the Bayesian, while there is a large class of probabilities such that the Bayesian agent loses all his wealth against the under-reacting agent almost surely.

In section 3, I prove a series of results about the relative accuracy between Bayes’ rule and the prediction schemes that under-react to information. These results nest those of Epstein et al. (2008, 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 accurate predictions as Bayes’ — under-reacting rules *weakly merge* (Kalai and Lehrer, 1994) to the truth. 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 true parameters for the data generating process such that agent  $B$ ’s and agent  $NB$ ’s 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.

In section 4, I study the consumption-share dynamics between a Bayesian ( $B$ ) and a non-Bayesian ( $NB$ ) agent who under-reacts to information in the standard general equilibrium model with complete markets of Sandroni (2000) and Blume and Easley (2006). The two agents have identical information as 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 (Sandroni, 2000; Massari, 2017). An agent who under-reacts to information never vanishes against a Bayesian agent with the same information because in all cases in which agent  $B$  learns the truth,  $NB$  also learns the truth at a comparable rate. Conversely, a sufficient condition for a Bayesian agent to vanish against an under-reacting agent is that the learning problem is misspecified in such a way that the empirical distribution (true probability, if draws are i.i.d.) is not in the prior support and belongs to its convex hull. While the Bayesian agent is as accurate

as the most accurate model in the support (Berk, 1966), the under-reacting agent is more accurate than the most accurate model in the support because its predictions are a non-degenerate mixture of the most accurate models in the support.

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 methods is mixed. While there is a vast literature supporting Bayesian methods for portfolio selection problems (Klein and Bawa, 1976; Frost and Savarino, 1986), equally rich is the set of non-Bayesian approaches for robust portfolio allocation rules (DeMiguel et al., 2007; Garlappi et al., 2006; Goldfarb and Iyengar, 2003).

My result complements and strengthens the arguments in favour of the “robust” methods proposed in the portfolio selection literature. It complements them by pointing to model misspecification, rather than parameter estimation error, as the reason for the sub-optimal behaviour of Bayesian methods. It strengthens them because of the known magnification effect that model misspecification has on parameter estimation error. In misspecified learning problems, the Bayesian posterior convergence can be significantly slower than in well-specified problems (Grünwald and van Ommen, 2014), demanding of special consideration in the use of model selection criteria (Balasubramanian, 1997) and in the use of Bayesian methods for predictions (Grünwald and Langford, 2007; Grünwald, 2012). Furthermore, my result offers an intuitive explanation of the difficulties often found in exploiting under-reaction anomalies with practical trading strategies. In terms of portfolio returns, it is hard to take advantage of under-reaction to information without a radical change of the underlying statistical model because it is a robust response to model misspecification, rather than an irrational deviation from a correct learning procedure.

## 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  $\Sigma := \times^\infty \mathcal{S}$ .  $\sigma = (\sigma_1, \dots)$  is a representative path,  $\sigma^t = (\sigma_1, \dots, \sigma_t)$  is a partial history until period  $t$ , and  $\mathcal{F}_t$  is the  $\sigma$ -algebra generated by the cylinders with base  $\sigma^t$ . By construction  $(\mathcal{F}_t)_{t=0}^\infty$  is a filtration and  $\mathcal{F}$  is the  $\sigma$ -algebra generated by their union.  $P$  denotes the true measure on  $(\Sigma, \mathcal{F})$ ; unless stated, I make no assumptions on  $P$ . For any probability measure  $\rho$  on  $(\Sigma, \mathcal{F})$ ,  $\rho(\sigma^t) := \rho(\{\sigma_1 \times \dots \times \sigma_t \times S \times S \times \dots\})$  is the marginal probability of the partial history  $\sigma^t$ , while  $\rho_t := \rho(\sigma_t |) := \frac{\rho(\sigma^t)}{\rho(\sigma^{t-1})}$  is the conditional probability of the generic state  $\sigma_t$  given  $\sigma^{t-1}$ , so that  $\rho(\sigma^t) = \prod_{\tau=1}^t \rho(\sigma_\tau |)$ .

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  $\beta$ .<sup>2</sup> For all paths  $\sigma$ , each agent  $i$  is endowed with

---

<sup>2</sup>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.

a stream of the consumption good,  $(e_t^i(\sigma))_{t=0}^\infty$ . Each agent's objective is to maximize the stream of discounted expected utility he gets from consumption. Expectations are computed according to agent beliefs  $p^i$ , a measure on  $(\Sigma, \mathcal{F})$ . Naming  $q(\sigma^t)$  the date  $t = 0$  price of the asset that delivers one unit of consumption in event  $\sigma^t$  and none otherwise, agent  $i = B, NB$  maximization reads:

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

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:  $\forall(t, \sigma), \sum_{i=N, NB} e_t^i(\sigma) = \sum_{i=N, NB} c_t^i(\sigma)$ . 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  $u^i : \mathbb{R}_+ \rightarrow [-\infty, +\infty]$  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** (i) For all agents  $i = B, NB$  and for all  $(t, \sigma)$ ,  $p^i(\sigma^t) > 0 \Leftrightarrow P(\sigma^t) > 0$ .  
(ii)  $\exists \epsilon > 0$  such that for all agents  $i = B, NB$  and for all  $(t, \sigma)$ ,  $p^i(\sigma_t) > \epsilon$ .

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

## 2.1 Beliefs

The Bayesian agent and the under-reacting agent share identical information. Adopting the Bayesian terminology, at time zero they believe that states are i.i.d. according to multinomial distributions  $\pi_\theta$  parametrized by  $K$  vectors of parameters  $\theta \in \Delta^{|S|}$ . Both agents have an identical full-support, time-zero prior distribution  $\mu_0$  on the set of these parameters  $\Theta = \{\theta^1, \dots, \theta^K\}$ .<sup>3</sup> Furthermore, both agents observe the same path  $\sigma^t$ .

Agent  $NB$  under-reacts to information: his beliefs are consistent with the axiomatization of Epstein (2006) and Epstein et al. (2008) where an agent is self-aware of her biases and fully anticipates her updating behavior when formulating plans. Consider an agent who is trying to learn the true parameter in a set  $\Theta$ . Updating of beliefs in response to observations  $\{\sigma_1, \dots, \sigma_t\}$  leads to the process of posteriors  $\{\mu_t\}$ , each  $\mu_t$  is a probability measure on  $\Theta$ . Bayesian updating leads to the process

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

---

<sup>3</sup>All results remain true for heterogeneous full-support prior distributions on the same finite support.

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  $\sigma_{t+1}$  (Definition 2). A prediction rule under-reacts to information (has *prior-bias*) with respect to Bayes' if the weights given by its predictive distribution are a convex combination between the prior weights and the Bayesian posterior weights calculated using the same information

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

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

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

An updating rule satisfying Definition 1 can be interpreted as attaching too much weight to prior beliefs  $\mu_t$  and hence under-reacting to observations. So, the parameter  $\alpha$  regulates the amount of under-reaction of agent  $NB$ . To lower values of  $\alpha$  correspond greater under-reaction.

The following Definitions characterize the dynamics of the beliefs of agents  $B$  (Bayes' rule) and agent  $NB$ . Proposition 1, below, verifies that indeed agent  $NB$ 's beliefs under-react to information according to Definition 1.

**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:<sup>4</sup>*

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

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

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

The parameter  $\alpha$  regulates the amount of under-reaction of agent  $NB$ . Lower values of  $\alpha$  correspond to higher under-reaction. Notably, with  $\alpha = 1$  Definitions 1 and 3 coincide with Bayes' rule.

## 2.2 Agents accuracy and survival

In this section, I remind the reader of standard definitions of accuracy in statistics and their implications in term of agents' survival. The asymptotic fate of an agent is characterized by his consumption-shares as follows.

---

<sup>4</sup>Inspection of Definition 3 reveals that  $p^{NB}$  belong to the class of market probability introduced by Massari (2018); specifically, they correspond to the evolution of the risk neutral probability of an economy in which agents have log utility and behavioural beliefs described in Dindo and Massari (2017).

**Definition 4.** Agent  $i$  **vanishes** on path  $\sigma$  if  $\lim_{t \rightarrow \infty} c_t^i(\sigma) = 0$  on  $\sigma$ , he **survives** on path  $\sigma$  if  $\limsup_{t \rightarrow \infty} c_t^i(\sigma) > 0$  on  $\sigma$ .

I rank agents' accuracy according to their likelihood and, more coarsely, according to their average (conditional) relative entropies.

**Definition 5.**

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

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

is the average (conditional) relative entropy from  $p$  to the true probability  $P$ .<sup>5</sup>

Using the average (conditional) relative entropies to rank agents' accuracy is the standard in the market selection literature since Blume and Easley (1992) because it can be readily used to determine a sufficient condition for an agent to vanish.

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

Conditions that rely on average accuracy, however, are unable to deliver the necessary part of Proposition 2 even in a two-agent economy. 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-likelihoods 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).

In order to precisely characterise agents' survival I rely on Massari (2017)'s necessary and sufficient condition for an agent to vanish. In a two-agent economy, Massari (2017)'s analysis implies that an agent survives on a path  $\sigma$  if and only if he is more accurate than the other agent.

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

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

---

<sup>5</sup>The relative entropy  $d(P_t||p_t)$  is uniquely minimized at  $p_t = P_t$ , strictly convex, and  $d(P_t||p_t) = \bar{d}(P||p)$   $P$ -a.s. whenever  $P$  and  $p$  are i.i.d. measures. The average relative entropy is an approximation of the average likelihood ratio that holds almost surely according to the true probability. If  $P$  attaches zero probability to some states, I adopt the convention  $0 \ln 0 = 0$ .

In this paper I make use of both results: Proposition 2 to show that agent  $B$  vanishes almost surely according to many true data generating process in misspecified learning problems (Theorem 4); Proposition 3 to show that the NB agent survives on every equilibrium path (Theorem 3), and thus for every true data generating process.

### 3 Relative accuracy of $p^{NB}$ and $p^B$

In this section, 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 provides sufficient conditions on the true probability such that  $p^{NB}$  is averagely more accurate than  $p^B$ , so that the likelihood ratio between  $p^B$  and  $p^{NB}$  converges to zero  $P^*$ -almost surely (Corollary 2).

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 interpreted as being a Bayesian mixture model on the non-i.i.d. probabilities  $p_\theta$ , rather than  $\pi_\theta$ . Therefore, standard Bayesian results, such as the universal bounds on the right, hold for the Bayesian mixture model  $p^{NB}$  with support  $\{p_{\theta^1}, \dots, p_{\theta^k}\}$ . When,  $\alpha = 1$ ,  $\pi_\theta = p_\theta$  (Definitions 2 and 3 coincide) and the second line holds.

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

$$\forall(t, \sigma), p^{NB}(\sigma^t) = \sum_{\theta \in \Theta} p_\theta(\sigma^t) \mu_0(\theta) \in \left[ \max_{\theta \in \Theta} p_\theta(\sigma^t) \mu_0(\theta), \max_{\theta \in \Theta} p_\theta(\sigma^t) \right]; \quad (1)$$

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

Bounds 1 and 2 highlight that the key step to discussing the relative accuracy between  $p^{NB}$  and  $p^B$  is to characterize the relative accuracy of models  $p_\theta$  against models  $\pi_\theta$ . Lemma 2 below provides this result. It shows that, on every path, and for every  $\theta$ , the likelihood ratio between  $p_\theta$  and  $\pi_\theta$  is uniformly bounded away from zero.

**Lemma 2.**

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

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

**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 represents 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  $p^B$  is only marginally faster than that of  $p^{NB}$ . As intuition suggests, to smaller values of  $\alpha$  correspond 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  $p^B(\sigma^t)$  priors do not converge to a unique model. These are the most interesting paths because  $p^{NB}(\sigma^t)$  is averagely more accurate than  $p^B(\sigma^t)$  on those paths. Next, I provide a sufficient condition on  $P$  that guarantees that the  $p^{NB}$  prior never concentrates on a unique model so that  $p^{NB}$  is averagely more accurate than  $p^B$ .

Theorem 2 shows that as long as the learning problem is misspecified in such a way that the parameters of the empirical distribution (true probability if states are i.i.d.) lie in the convex hull of the support, there is an  $\bar{\alpha}$  such that  $p^{NB}$  is averagely more accurate than  $p^B$  for all  $\alpha \in (0, \bar{\alpha}]$ .

**Theorem 2.** *Assume that the true probability is such that the empirical distribution of states,  $P^*$ , exists, and let  $\theta^*$  be its vector of parameters,<sup>6</sup>*

$$\theta^* \in \text{Conv}(\Theta) \setminus \Theta \Rightarrow \exists \bar{\alpha} : \forall \alpha \in (0, \bar{\alpha}), \bar{d}(P^* || p^B) > \bar{d}(P^* || p^{NB}) \text{ } P^* \text{-a.s.}; \quad (3)$$

specifically, let  $\hat{\theta} \in \text{argmin}_{\theta \in \Theta} \bar{d}(P_{\theta^*} || \pi_{\theta})$ ,

$$\bar{\alpha} := \text{argmax}_{\alpha \in (0, 1)} : \bar{d}(P_{\theta^*} || \pi_{\hat{\theta}}) > \min_{\theta \in \Theta \setminus \hat{\theta}} \bar{d}(P_{\theta^*} || (1 - \alpha)\pi_{\hat{\theta}} + \alpha\pi_{\theta}). \quad (4)$$

In order to have  $\bar{d}(P^* || p^B) > \bar{d}(P^* || p^{NB})$   $P^*$ -a.s., I need a condition that prevents  $p^{NB}$  prior from ever concentrating on a unique model.<sup>7</sup> Condition 4 of Theorem 2 precisely guarantees that concentration does not occur  $P^*$ -a.s.. By standard Bayesian argument (Berk, 1966; Marinacci and Massari, 2019), a Bayesian mixture concentrates on the model on its support with the lowest average accuracy, when this model is unique (See Lemma 7 for a proof that applies to this setting). If  $p^{NB}$  prior were to concentrate on a model  $p_{\theta}$ , then  $p^{NB} \rightarrow p_{\theta} = \pi_{\theta}$  and its average accuracy would be higher than the average accuracy of another model in the support because, by condition (4),  $\bar{d}(P_{\theta^*} || \pi_{\hat{\theta}}) > \min_{\theta \in \Theta \setminus \hat{\theta}} \bar{d}(P_{\theta^*} || (1 - \alpha)\pi_{\hat{\theta}} + \alpha\pi_{\theta})$ ; a contradiction. Condition 3,  $\theta^* \in \text{Conv}(\Theta) \setminus \Theta$ , is easy to verify, and it is sufficient to guarantee the existence of such  $\bar{\alpha}$  and  $\theta$ .

<sup>6</sup>If states are i.i.d. with true probability  $P_{\theta_0}$ , the empirical distribution coincides with the true distribution almost surely,  $P_{\theta^*} = P_{\theta_0}$ ; so,  $P_{\theta_0}$  can be replaced for  $P_{\theta^*}$  in conditions 3 and 4.

<sup>7</sup>Inspection of Definitions 2 and 3 shows that if  $p^{NB}$  prior concentrates on a unique model  $p_{\theta}$ , then it must be the case that  $\lim_{t \rightarrow \infty} p_t^{NB} = \lim_{t \rightarrow \infty} p_{\theta, t} = \pi_{\theta} = \lim_{t \rightarrow \infty} p_t^B$  — where the last equality follows is proven in Lemma 7. So that  $p^{NB}$  is averagely as accurate as  $p^B$  whenever it concentrates on a unique model.

I conclude with two corollaries, Corollary 1 shows that if the learning problem is correctly specified, then  $NB$  merges with the truth. That is,  $NB$  converges to the truth qualitatively as fast as  $B$  does. This result strengthens the *weak merging* result of (Epstein et al., 2010)

**Corollary 1.** *For all  $\alpha \in (0, 1]$ , for all  $\theta \in \Theta$ ,  $p^{NB}$  merges with  $\pi_\theta$ ,  $\pi_\theta$ -almost surely.*

*Proof.*

$$\begin{aligned} \theta \in \Theta &\Rightarrow \pi_\theta \text{ is absolutely continuous with respect to } p^B \\ &\Rightarrow p^B \text{ merges with } \pi_\theta, \pi_\theta\text{-a.s.}; \\ \forall \sigma, \frac{p^{NB}(\sigma^t)}{p^B(\sigma^t)} &\underset{\text{By Theorem 1}}{>} 0 \Rightarrow p^B(\sigma^t) \text{ is absolutely continuous with respect to } p^{NB}, p^B\text{-a.s.}; \\ \text{thus, } \theta \in \Theta &\Rightarrow \pi_\theta \text{ is absolutely continuous with respect to } p^{NB} \\ &\Rightarrow p^{NB} \text{ merges with } \pi_\theta, \pi_\theta\text{-a.s.} \end{aligned}$$

□

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

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

*Proof.* By Theorem 2,  $\bar{d}(P^*||p^B) > \bar{d}(P^*||p^{NB})$   $P^*$ -a.s.  $\Rightarrow \frac{p^B(\sigma^t)}{p^{NB}(\sigma^t)} \rightarrow 0$   $P^*$ -a.s., by the strong law of large number for martingale differences (See the application in Lemma 3). □

## 4 Agents' survival

In this section, I present two results showing that agent  $NB$  has an evolutionary advantage over agent  $B$ . First, Theorem 3 shows that there are no paths (and thus no true data generating process) on which agent  $NB$  vanishes against agent  $B$ . It is impossible for a Bayesian agent to drive out of the market an under-reacting agent with the same information. Furthermore, the consumption share of agent  $NB$  is guaranteed to remain strictly positive.

**Theorem 3.** *Under **A1-A3**, on every path  $\sigma \in \Sigma$  and for all  $\alpha \in (0, 1]$ , agent  $NB$  survives. Furthermore, agent  $NB$ , consumption is uniformly bounded away from 0.*

*Proof.*

$$\text{By Theorem 1, } \forall \sigma \in \Sigma, \forall t, \frac{p^{NB}(\sigma^t)}{p^B} \geq \left( \min_{\theta \in \Theta} \mu_0(\theta) \right)^{\frac{1}{\alpha}};$$

thus, the sufficient condition for an agent to survive with positive consumption-share of Proposition 3 is satisfied by agent  $NB$  on every path. □

Theorem 3 tells us that it is impossible for an under-reacting agent to lose all his wealth against a Bayesian agent with the same information. If the learning problem is correctly specified, both agents learn the true model at a comparable rate and their beliefs become identical. If the learning problem is misspecified, the Bayesian agent (generically) learns what is the most accurate model in its support, while the NB agent's predictions might remain a combination of some models in the support. Irrespective of the path of realizations, agent  $NB$ 's beliefs are guaranteed to be at least as accurate as those of agent  $B$  (Theorem 1).

Next, Theorem 4 shows that there are many data generating processes,  $P$ , such that agent  $B$  vanishes against agent  $NB$   $P$ -a.s.. If the learning problem is misspecified and the truth is such that the empirical distribution it generates has parameters that belong to the convex hull of the models in the support, then there is a level for the under-reaction parameter that, if crossed, guarantees that agent  $NB$  dominates over agent  $B$ . When the empirical distribution is a convex combination of two models in the support, agent  $NB$  makes predictions that are more accurate than agent  $B$ 's because its prior remains non-degenerate due to under-reaction.

**Theorem 4.** *Under A1-A3, let  $P^*$  be the empirical distribution and  $\theta^*$  be its vector of parameters;*

$$\theta^* \in \text{Conv}(\Theta) \setminus \Theta \Rightarrow \exists \bar{\alpha} : \forall \alpha \in (0, \bar{\alpha}) \text{ agent } B \text{ vanishes } P^* \text{-a.s..}$$

*Specifically, let  $\hat{\theta} \in \text{argmin}_{\theta \in \Theta} \bar{d}(P_{\theta^*} || \pi_{\theta})$ ,*

$$\bar{\alpha} := \text{argmax}_{\alpha \in (0,1)} : \bar{d}(P_{\theta^*} || \pi_{\hat{\theta}}) > \min_{\theta \in \Theta \setminus \hat{\theta}} \bar{d}(P_{\theta^*} || (1 - \alpha)\pi_{\hat{\theta}} + \alpha\pi_{\theta}).$$

*Proof.* By Proposition 2, agent  $B$  vanishes if  $\bar{d}(P || p^{NB}) < \bar{d}(P || p^B)$ ,  $P$ -a.s.;

$$\theta^* \in \text{Conv}(\Theta) \setminus \Theta, \exists \bar{\alpha} : \alpha \in (0, \bar{\alpha}) \quad \underbrace{\Rightarrow}_{\text{By Theorem 2}} \quad \bar{d}(P || p^{NB}) < \bar{d}(P || p^B), \text{ } P \text{-a.s..}$$

□

Theorems 4 tells us that the survival chances of agent  $B$  are lower against  $NB$  agents whose underreaction to information is greater. The smaller the  $\alpha$  parameter, the larger is the set of true probabilities (and thus sequences) in which the agent  $B$  vanishes.

## Discussion

Theorem 3 and 4 tell us that an under-reacting agent cannot vanish against a Bayesian with the same information, and that there are generic cases of model misspecification in which the Bayesian agent vanishes against an underreacting agent, respectively. Proposition 4, below, shows that agent  $B$  vanishes against agent  $NB$  precisely in those cases in which the true probability is such that agent  $NB$  prior does not concentrate on a unique model. So a non-Bayesian dynamics of prices is a possible evolutionary outcome of model misspecification, rather than a transient phenomenon.

**Proposition 4.** *Under **A1-A3**, agent  $B$  vanishes if the beliefs of agent  $NB$  never settle on a unique model.*

*Proof.* By Proposition 2, agent  $B$  vanishes if  $\bar{d}(P||p^{NB}) < \bar{d}(P||p^B)$ ,  $P$ -a.s.. I have shown that

$$\bar{d}(P||p^{NB}) \underbrace{\leq}_{\text{By Lemma 3}} \min_{\theta \in \Theta} \bar{d}(P||p_{\theta}) \underbrace{<}_{\text{By Lemma 4}} \min_{\theta \in \Theta} \bar{d}(P||\pi_{\theta}) \underbrace{=}_{\text{By Lemma 7}} \bar{d}(P||p^B).$$

Where the second inequality is strict because, by assumption, “the beliefs of agent  $NB$  never settle on a unique model”.  $\square$

## 4.1 Example

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

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

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

Suppose  $S := \{u, d\}$  with true probability  $P_{\theta_0}(u) = .5$ , and that the two agents have identical prior support  $\Theta := \{\pi_{\theta_1}, \pi_{\theta_2}\}$  with  $\pi_{\theta_1}(u) = .4$  and  $\pi_{\theta_2}(u) = .7$  and uniform time zero prior  $\mu_0(\theta_1) = .5 = \mu_0(\theta_2)$ . Because the first model is averagely more accurate —  $\bar{d}(P_{\theta^*}||\pi_{\theta_1}) < \bar{d}(P_{\theta^*}||\pi_{\theta_2})$  — standard results in Bayesian learning (Berk, 1966) tell us that the Bayesian posterior converges almost surely (exponentially fast) to a Dirac on  $\pi_{\theta_1}$ , so that  $\bar{d}(P_{\theta^*}||p^B) = \bar{d}(P_{\theta^*}||\pi_{\theta_1})$ . Conversely, with  $\alpha = .2 < 2/3 = \bar{\alpha}$  agent  $NB$ ’s beliefs do not converge to any parameter and his predictions for state  $u$ ,  $p^{NB}(u_t|)$  spend most of the periods in the interval  $(.4, .6)$ . So, agent  $NB$  makes predictions that are strictly more accurate than  $\pi_{\theta_1}$  in most periods (see Fig 1 for a simulation). Thus, agent  $NB$ ’s beliefs are averagely more accurate than agent  $B$ ’s —  $\bar{d}(P_{\theta^*}||p^{NB}) < \bar{d}(P_{\theta^*}||\pi_{\theta_1}) = \bar{d}(P_{\theta^*}||p^B)$ — and agent  $B$  vanishes. In the long run, prices reflects agent  $NB$ ’s beliefs, which are not Bayesian.

## 5 Conclusion

Under-reaction is not a transient phenomenon in financial markets. Financial markets favour “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’.

## 6 Appendix

### Proof of Proposition 1

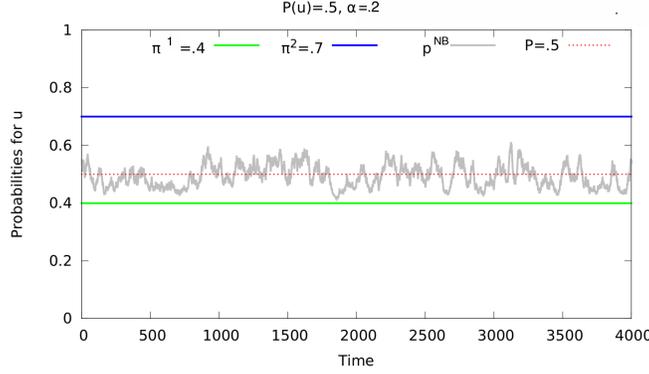


Figure 1:  $p^{NB}$  dynamics with  $[\pi^1(u), \pi^2(u)] = [.4, .7]$ , uniform time zero prior and mixing coefficient  $\alpha = .2$ . The probability weights never find a resting points and the resulting probabilities remains closer to the truth than the most accurate model in the support,  $\pi^1$ , in most periods.

*Proof.* By Definition 3,

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

Letting  $\mu^{NB}(\theta|\sigma^{t-1}) = \mu^B(\theta|\sigma^{t-1})$  on the right hand side we obtain:

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

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

*Proof.* The Lagrangian problem associated with each trader's maximization problem

is

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

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

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

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  $\sigma^t$ ,

$$\frac{u^B(c_t^B(\sigma))'}{u^{NB}(c_t^{NB}(\sigma))'} = \frac{p^{NB}(\sigma^t)}{p^B(\sigma^t)} \frac{u^B(c_0^B)'}{u^{NB}(c_0^{NB})'} \quad (5)$$

Thus,

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

□

### Proof of Lemma 1

*Proof.* Note that

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

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  $\alpha \in (0, 1]$ ,

$$\begin{aligned}
\forall(t, \sigma), p^{NB}(\sigma^t) &:= \prod_{\tau=1}^t p^{NB}(\sigma_\tau | \sigma^{\tau-1}) \\
&= \left( \sum_{\theta \in \Theta} p_\theta(\sigma_t | \sigma^{t-1}) \mu^{NB}(\theta | \sigma^{t-1}) \right) \prod_{\tau=1}^{t-1} p^{NB}(\sigma_\tau | \sigma^{\tau-1}) \\
&\stackrel{\text{By Def.3}}{=} \left( \sum_{\theta \in \Theta} p_\theta(\sigma_t | \sigma^{t-1}) p_\theta(\sigma_{t-1} | \sigma^{t-2}) \mu^{NB}(\theta | \sigma^{t-2}) \right) \frac{1}{p^{NB}(\sigma_{t-1} | \sigma^{t-1})} \prod_{\tau=1}^{t-1} p^{NB}(\sigma_\tau | \sigma^{\tau-1}) \\
&= \left( \sum_{\theta \in \Theta} p_\theta(\sigma_t | \sigma^{t-1}) p_\theta(\sigma_{t-1} | \sigma^{t-2}) \mu^{NB}(\theta | \sigma^{t-2}) \right) \prod_{\tau=1}^{t-2} p^{NB}(\sigma_\tau | \sigma^{\tau-1}) \\
&\quad \vdots \\
&= \sum_{\theta \in \Theta} \prod_{\tau=1}^t p_\theta(\sigma_\tau | \sigma^{\tau-1}) \mu_0(\theta) \\
&= \sum_{\theta \in \Theta} p_\theta(\sigma^t) \mu_0(\theta) \in \left[ \max_{\theta \in \Theta} p_\theta(\sigma^t) \mu_0(\theta), \max_{\theta \in \Theta} p_\theta(\sigma^t) \right]
\end{aligned}$$

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

*Proof.*

$$\begin{aligned}
\forall(t, \sigma), \forall \theta \in \Theta, \quad \ln p_\theta(\sigma^t) &= \ln \prod_{\tau=1}^t ((1-\alpha)p^{NB}(\sigma_\tau) + \alpha\pi_\theta(\sigma_\tau)) \\
&= \sum_{\tau=1}^t \ln ((1-\alpha)p^{NB}(\sigma_\tau) + \alpha\pi_\theta(\sigma_\tau)) \\
&\stackrel{\geq}{\underbrace{\hspace{1cm}}} (1-\alpha) \sum_{\tau=1}^t \ln p^{NB}(\sigma_\tau) + \alpha \sum_{\tau=1}^t \ln \pi_\theta(\sigma_\tau) \\
&\quad \text{By concavity of } \ln(\cdot) \\
&= (1-\alpha) \ln p^{NB}(\sigma^t) + \alpha \ln \pi_\theta(\sigma^t) \\
\Rightarrow \forall(t, \sigma), \forall \theta \in \Theta, \quad \ln \frac{p_\theta(\sigma^t)}{\pi_\theta(\sigma^t)} &\geq \frac{(1-\alpha)}{\alpha} \ln \frac{p^{NB}(\sigma^t)}{p_\theta(\sigma^t)} \\
&\stackrel{=}{\underbrace{\hspace{1cm}}} \frac{(1-\alpha)}{\alpha} \ln \frac{\sum_{\theta \in \Theta} p_\theta(\sigma^t) \mu_0(\theta)}{p_\theta(\sigma^t)} \\
&\quad \text{By Eq. (1)} \\
&\geq \ln \left( \min_{\theta \in \Theta} \mu_0(\theta) \right)^{\frac{(1-\alpha)}{\alpha}} \\
\Rightarrow \forall(t, \sigma), \forall \theta \in \Theta, \quad p_\theta(\sigma^t) &\geq \left( \min_{\theta \in \Theta} \mu_0(\theta) \right)^{\frac{1}{\alpha}-1} \pi_\theta(\sigma^t).
\end{aligned}$$

□

### Proof of Theorem 1

*Proof.* The result follows from Lemmas 1 and 2

$$\begin{aligned}
\forall(t, \sigma), \frac{p^{NB}(\sigma^t)}{p^B(\sigma^t)} &\stackrel{=}{\underbrace{\hspace{1cm}}} \frac{\sum_{\theta \in \Theta} p_\theta(\sigma^t) \mu_0(\theta)}{\sum_{\theta \in \Theta} \pi_\theta(\sigma^t) \mu_0(\theta)} \\
&\quad \text{By Lemma 1} \\
&\geq \frac{\max_{\theta \in \Theta} p_\theta(\sigma^t) \mu_0(\theta)}{\max_{\theta \in \Theta} \pi_\theta(\sigma^t)} \\
&\stackrel{\geq}{\underbrace{\hspace{1cm}}} \frac{\max_{\theta \in \Theta} \pi_\theta(\sigma^t) \mu_0(\theta)}{\max_{\theta \in \Theta} \pi_\theta(\sigma^t)} \left( \min_{\theta \in \Theta} \mu_0(\theta) \right)^{\frac{1}{\alpha}-1} \\
&\quad \text{By Lemma 2} \\
&\geq \left( \min_{\theta \in \Theta} \mu_0(\theta) \right)^{\frac{1}{\alpha}} > 0.
\end{aligned}$$

da farsi

□

### Proof of Theorem 2

*Proof.* By assumption  $\theta^* \in \text{Conv}(\Theta) \setminus \Theta$  so  $\min_{\theta \in \Theta} \bar{d}(P_{\theta^*} || \pi(\theta)) = \eta > 0$ .

$$\text{Thus, } \bar{d}(P_{\theta^*} || p^B) \stackrel{=}{\underbrace{\hspace{1cm}}} \bar{d}(P_{\theta^*} || \pi_{\hat{\theta}}) = \eta > 0.$$

by Lemma 7

What is left to show is that  $\alpha < \bar{\alpha} \Rightarrow \bar{d}(P_{\theta^*} || P^{NB}) < \bar{d}(P_{\theta^*} || \pi_{\hat{\theta}})$ .

This inequality holds because

$$\bar{d}(P_{\theta^*} || P^{NB}) \underbrace{\leq}_{\substack{\text{by} \\ \text{Lemma 3}}} \bar{d}(P_{\theta^*} || p_{\hat{\theta}}) \underbrace{\leq}_{\substack{\text{by} \\ \text{Lemma 4}}} \bar{d}(P_{\theta^*} || \pi_{\hat{\theta}});$$

where Lemma 5 implies that the second inequality is strict for  $\alpha < \bar{\alpha}$  and Lemma 6 shows that  $\theta^* \in \text{Conv}(\Theta) \setminus \Theta$  is a sufficient condition for the existence of such  $\bar{\alpha}$ .  $\square$

**Lemma 3.**

$$\forall \alpha \in (0, 1], \forall \theta \in \Theta, \bar{d}(P_{\theta^*} || P^{NB}) \leq \bar{d}(P_{\theta^*} || p_{\theta}) \quad P_{\theta^*}\text{-a.s.}$$

*Proof.* By Lemma 1, for all  $(t, \sigma)$ ,  $p^{NB}(\sigma^t) = \sum_{\theta \in \Theta} p_{\theta}(\sigma^t) \mu_0(\theta)$ . Thus,

$$\begin{aligned} & \forall (t, \sigma), \forall \theta \in \Theta, \quad \ln p^{NB}(\sigma^t) \geq \ln p_{\theta}(\sigma^t) + \ln \mu_0(\theta), \\ \Rightarrow & \quad \frac{1}{t} \ln \frac{P_{\theta^*}(\sigma^t)}{p^{NB}(\sigma^t)} \leq \frac{1}{t} \ln \frac{P_{\theta^*}(\sigma^t)}{p_{\theta}(\sigma^t)} - \frac{1}{t} \ln \mu_0(\theta) \\ \Rightarrow & \quad \lim_{t \rightarrow \infty} \left[ \frac{1}{t} \left[ \sum_{\tau=1}^t \ln \frac{P_{\theta^*}(\sigma_{\tau})}{p^{NB}(\sigma_{\tau})} - \sum_{\tau=1}^t d(P_{\theta^*} || p_{\tau}^{NB}) \right] + \frac{1}{t} \sum_{\tau=1}^t d(P_{\theta^*} || p_{\tau}^{NB}) \right] \\ & \leq \lim_{t \rightarrow \infty} \left[ \frac{1}{t} \left[ \sum_{\tau=1}^t \ln \frac{P(\sigma_{\tau})}{p_{\theta}(\sigma_{\tau})} - \sum_{\tau=1}^t d(P_{\theta^*} || p_{\theta, \tau}) \right] + \frac{1}{t} \sum_{\tau=1}^t d(P_{\theta^*} || p_{\theta, \tau}) - \frac{1}{t} \ln \mu_0(\theta) \right] \\ \Rightarrow & \quad \bar{d}(P_{\theta^*} || p^{NB}) \leq \bar{d}(P_{\theta^*} || p_{\theta}) \quad P_{\theta^*}\text{-a.s.} \end{aligned}$$

The last implication follows from the strong law of large number for martingale differences (see also Sandroni, 2000) that guarantees that for  $\rho = p_{\theta}, p^{NB}$ ,

$$\lim_{t \rightarrow \infty} \frac{1}{t} \left[ \sum_{\tau=1}^t \ln \frac{P(\sigma_{\tau})}{\rho(\sigma_{\tau})} - \sum_{\tau=1}^t d(P_{\theta^*} || \rho_{\tau}) \right] = 0, P_{\theta^*}\text{-a.s.}$$

$\square$

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

$$\bar{d}(P_{\theta^*} || p_{\theta}) \leq \bar{d}(P_{\theta^*} || \pi_{\theta}),$$

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

*Proof.* For all,  $(t, \sigma)$ , for all  $\theta \in \Theta$ , and for all  $\alpha \in (0, 1)$ ,

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

(a) : if there exists an  $\epsilon > 0$  such that  $\|p_t^{NB} - \pi_\theta\| > \epsilon$  a positive fraction of periods, then inequality (a) is strict a positive fraction of periods because  $d(P_{\theta^*}|\cdot)$  is continuous, strictly convex, which implies that  $\bar{d}(P_{\theta^*}|p_\theta) < \bar{d}(P_{\theta^*}|\pi_\theta)$  by definition.  $\square$

**Lemma 5.**

$\bar{d}(P_{\theta^*}|\pi_{\hat{\theta}}) > \min_{\theta \in \Theta \setminus \hat{\theta}} \bar{d}(P_{\theta^*}|| (1 - \alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta) \Rightarrow \exists \epsilon > 0 : \|p_t^{NB} - \pi_{\hat{\theta}}\| > \epsilon$  a positive fraction of periods.

*Proof.* I prove the contrapositive statement

$$\|p_t^{NB} - \pi_{\hat{\theta}}\| = 0 \text{ in most periods} \Rightarrow \bar{d}(P_{\theta^*}|\pi_{\hat{\theta}}) \leq \min_{\theta \in \Theta \setminus \hat{\theta}} \bar{d}(P_{\theta^*}|| (1 - \alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta)$$

in three steps.

First:

$$\begin{aligned} & \|p_t^{NB} - \pi_{\hat{\theta}}\| = 0 \text{ in most periods} \\ \Rightarrow & |d(P_{\theta^*,t}|\pi_{\hat{\theta}}) - d(P_{\theta^*,t}|p_{\hat{\theta},t})| = 0 \text{ in most periods} \\ \text{By continuity} & \\ \Rightarrow & \bar{d}(P_{\theta^*}|\pi_{\hat{\theta}}) = \bar{d}(P_{\theta^*}|p^{NB}) \end{aligned} \quad (6)$$

Second:

$$\begin{aligned} & \|p_t^{NB} - \pi_{\hat{\theta}}\| = 0 \text{ in most periods} \\ \Rightarrow & \left| \min_{\theta \in \Theta \setminus \hat{\theta}} d(P_{\theta^*,t}|| (1 - \alpha)p_t^{NB} + \alpha\pi_\theta) - \min_{\theta \in \Theta \setminus \hat{\theta}} d(P_{\theta^*,t}|| (1 - \alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta) \right| = 0 \text{ in most periods} \\ \Rightarrow & \left| \min_{\theta \in \Theta \setminus \hat{\theta}} d(P_{\theta^*,t}|p_{\theta,t}) - \min_{\theta \in \Theta \setminus \hat{\theta}} d(P_{\theta^*,t}|| (1 - \alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta) \right| = 0 \text{ in most periods} \\ \text{By definition} & \\ \Rightarrow & \min_{\theta \in \Theta \setminus \hat{\theta}} \bar{d}(P_{\theta^*}|p_{\hat{\theta}}) = \min_{\theta \in \Theta \setminus \hat{\theta}} \bar{d}(P_{\theta^*}|| (1 - \alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta) \end{aligned} \quad (7)$$

Last:

$$\bar{d}(P_{\theta^*}|\pi_{\hat{\theta}}) \underbrace{=}_{\text{Eq.(6)}} \bar{d}(P_{\theta^*}|p^{NB}) \underbrace{\leq}_{\text{By Lem.3}} \min_{\theta \in \Theta} \bar{d}(P_{\theta^*}|p_{\hat{\theta}}) \leq \min_{\theta \in \Theta \setminus \hat{\theta}} \bar{d}(P_{\theta^*}|p_{\hat{\theta}}) \underbrace{=}_{\text{Eq.(7)}} \min_{\theta \in \Theta \setminus \hat{\theta}} \bar{d}(P_{\theta^*}|| (1 - \alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta)$$

$\square$

**Lemma 6.**

$$\theta^* \in \text{Conv}(\Theta) \setminus \Theta \Rightarrow \exists \bar{\alpha} : \forall \alpha \in (0, \bar{\alpha}), \bar{d}(P_{\theta^*}|\pi_{\hat{\theta}}) > \min_{\theta \in \Theta \setminus \hat{\theta}} \bar{d}(P_{\theta^*}|| (1 - \alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta)$$

*Proof.*  $\theta^* \in \text{Conv}(\Theta) \setminus \Theta \Rightarrow \exists w \in \text{int}(\Delta^\Theta) : \sum_k w^k \theta^k = \theta^*$ . Thus, strict convexity of  $d(\cdot|\cdot)$  guarantees that, locally, the accuracy of each parameter can be improved by moving in the direction of at least one other parameter. That is:

$$\forall \hat{\theta} \in \Theta, \exists \theta' \in \Theta \setminus \hat{\theta}, \text{ and } \bar{\alpha} \in (0, 1) : \forall \alpha \in (0, \bar{\alpha}), \bar{d}(P_{\theta^*}|\pi_{\hat{\theta}}) > \min_{\theta \in \Theta \setminus \hat{\theta}} \bar{d}(P_{\theta^*}|| (1 - \alpha)\pi_{\hat{\theta}} + \alpha\pi_\theta).$$

□

**Lemma 7.** Let  $\hat{\theta} \in \operatorname{argmin}_{\theta \in \Theta} \bar{d}(P^* || p_{\hat{\theta}})$

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

So,  $\bar{d}(P^* || p^B) = \bar{d}(P^* || \pi_{\hat{\theta}}) P^* \text{-a.s.}$

*Proof.* By Lemma 3,  $\forall \theta \in \Theta, \bar{d}(P^* || p^B) \leq \bar{d}(P^* || p_{\theta})$ .  
I proceed by ruling out the strict inequality for  $\theta = \hat{\theta}$ .  
By contradiction, suppose that

$$\bar{d}(P^* || p^{NB}) < \bar{d}(P^* || p_{\hat{\theta}}) P_{\theta^*} \text{-a.s.}$$

The above, together with the strong law of large numbers (see lemma 3), would imply that  $\lim_{t \rightarrow \infty} \frac{p_{\hat{\theta}}(\sigma^t)}{p^{NB}(\sigma^t)} = 0 P_{\theta^*} \text{-a.s.}$ ; this implication is absurd because

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{p_{\hat{\theta}}(\sigma^t)}{p^{NB}(\sigma^t)} &\stackrel{\text{By Eq.(1)}}{=} \lim_{t \rightarrow \infty} \frac{p_{\hat{\theta}}(\sigma^t)}{\sum_{\theta \in \Theta} p_{\theta}(\sigma^t) \mu_0(\theta)} \\ &= \lim_{t \rightarrow \infty} \frac{1}{\mu_0(\hat{\theta}) + \sum_{\theta \in \Theta \setminus \hat{\theta}} \frac{p_{\theta}(\sigma^t)}{p_{\hat{\theta}}(\sigma^t)} \mu_0(\theta)} \\ &= \lim_{t \rightarrow \infty} \frac{1}{\mu_0(\hat{\theta}) + \sum_{\theta \in \Theta \setminus \hat{\theta}} e^{\ln \frac{p_{\theta}(\sigma^t)}{p_{\hat{\theta}}(\sigma^t)}} \mu_0(\theta)} \\ &\stackrel{P_{\theta^*} \text{-a.s.; by the SLLNMD}}{=} \lim_{t \rightarrow \infty} \frac{1}{\mu_0(\hat{\theta}) + \sum_{\theta \in \Theta \setminus \hat{\theta}} e^{t(\bar{d}(P_{\theta^*} || p_{\hat{\theta}}) - \bar{d}(P_{\theta^*} || p_{\theta}))} \mu_0(\theta)} \\ &\stackrel{>}{>} 0. \end{aligned}$$

Because  $p_{\hat{\theta}} \in \operatorname{argmin}_{\theta \in \Theta} \bar{d}(P_{\theta^*} || p_{\theta})$ .

Finally, note that  $\alpha = 1 \Rightarrow \forall (t, \sigma), p_t^B = p_t^{NB}$  and  $\forall \theta \in \Theta, p_{\theta,t} = \pi_{\theta}$ .  
So,  $\bar{d}(P^* || p^B) = \bar{d}(P^* || \pi_{\hat{\theta}}) P^* \text{-a.s.}$  □

## References

- Balasubramanian, V. (1997). Statistical inference, occam's razor, and statistical mechanics on the space of probability distributions. *Neural computation*, 9(2):349–368.
- Barberis, N., Shleifer, A., and Vishny, R. (1998). A model of investor sentiment. *Journal of financial economics*, 49(3):307–343.
- Berk, R. H. (1966). Limiting behavior of posterior distributions when the model is incorrect. *The Annals of Mathematical Statistics*, 37(1):51–58.
- Blume, L. and Easley, D. (1992). Evolution and market behavior. *Journal of Economic Theory*, 58:9–40.

- Blume, L. and Easley, D. (2006). If you are so smart why aren't you rich? Belief selection in complete and incomplete markets. *Econometrica*, 74:929–966.
- Box, G. E. (1976). Science and statistics. *Journal of the American Statistical Association*, 71(356):791–799.
- Csaba, D. and Szoke, B. (2018). Learning with misspecified models.
- Cutler, D. M., Poterba, J. M., and Summers, L. H. (1991). Speculative dynamics. *The Review of Economic Studies*, 58(3):529–546.
- DeMiguel, V., Garlappi, L., and Uppal, R. (2007). Optimal versus naive diversification: How inefficient is the  $1/n$  portfolio strategy? *The review of Financial studies*, 22(5):1915–1953.
- Dindo, P. and Massari, F. (2017). The wisdom of the crowd in dynamic economies. *University Ca'Foscari of Venice, Dept. of Economics Research Paper Series No*, 17.
- Epstein, L., Noor, J., and Sandroni, A. (2008). Non-bayesian updating: a theoretical framework. *Theoretical Economics*, 3:193 – 229.
- Epstein, L., Noor, J., and Sandroni, A. (2010). Non-bayesian learning. *The B.E. Journal of Theoretical Economics*, 10:Article 3.
- Epstein, L. G. (2006). An axiomatic model of non-bayesian updating. *The Review of Economic Studies*, 73(2):413–436.
- Frost, P. A. and Savarino, J. E. (1986). An empirical bayes approach to efficient portfolio selection. *Journal of Financial and Quantitative Analysis*, 21(3):293–305.
- Garlappi, L., Uppal, R., and Wang, T. (2006). Portfolio selection with parameter and model uncertainty: A multi-prior approach. *The Review of Financial Studies*, 20(1):41–81.
- Giglio, S. and Kelly, B. (2018). Excess volatility: Beyond discount rates. *The Quarterly Journal of Economics*, 133(1):71–127.
- Goldfarb, D. and Iyengar, G. (2003). Robust portfolio selection problems. *Mathematics of operations research*, 28(1):1–38.
- Grünwald, P. (2012). The safe bayesian. In *Algorithmic Learning Theory*, pages 169–183. Springer.
- Grünwald, P. and Langford, J. (2007). Suboptimal behavior of bayes and mdl in classification under misspecification. *Machine Learning*, 66(2-3):119–149.
- Grünwald, P. and van Ommen, T. (2014). Inconsistency of bayesian inference for misspecified linear models, and a proposal for repairing it. *arXiv preprint arXiv:1412.3730*.
- Grünwald, P., Van Ommen, T., et al. (2017). Inconsistency of bayesian inference for misspecified linear models, and a proposal for repairing it. *Bayesian Analysis*, 12(4):1069–1103.
- Kalai, E. and Lehrer, E. (1994). Weak and strong merging of opinions. *Journal of Mathematical Economics*, 23(1):73–86.
- Klein, R. W. and Bawa, V. S. (1976). The effect of estimation risk on optimal portfolio choice. *Journal of financial economics*, 3(3):215–231.

- Marinacci, M. and Massari, F. (2019). Learning from ambiguous and misspecified models. *Journal of Mathematical Economics*, 84:144–149.
- Massari, F. (2013). Comment on if you're so smart, why aren't you rich? belief selection in complete and incomplete markets. *Econometrica*, 81(2):849–851.
- Massari, F. (2017). Markets with heterogeneous beliefs: A necessary and sufficient condition for a trader to vanish. *Journal of Economic Dynamics and Control*, 78:190–205.
- Massari, F. (2018). Price probabilities: A class of bayesian and non-bayesian prediction rules. Working paper, University of New South Wales, Sydney.
- Massari, F. (2019). Ambiguity, robust statistics, and raiffa's critique. *Available at SSRN 3388410*.
- Peleg, B. and Yaari, M. E. (1970). Markets with countably many commodities. *International Economic Review*, 11(3):369–377.
- Sandroni, A. (2000). Do markets favor agents able to make accurate predictions. *Econometrica*, 68(6):1303–1341.
- Sandroni, A. (2005). Efficient markets and bayes' rule. *Economic Theory*, 26:741–764.
- Timmermann, A. (2006). Forecast combinations. *Handbook of economic forecasting*, 1:135–196.
- Wald, A. (1947). An essentially complete class of admissible decision functions. *The Annals of Mathematical Statistics*, 18(4):549–555.