Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
http://www.econ.upenn.edu/pier

PIER Working Paper 08-017

"Supplementary Appendix for 'Non-Bayesian Updating: A Theoretical Framework'"

by

Larry G. Epstein, Jawwad Noor and Alvaro Sandroni

http://ssrn.com/abstract=1133989

SUPPLEMENTARY APPENDIX FOR "NON-BAYESIAN UPDATING: A THEORETICAL FRAMEWORK"

Larry G. Epstein Jawwad Noor Alvaro Sandroni

January 18, 2008

1. INTRODUCTION

This appendix applies the model in 'Non-Bayesian Updating: A Theoretical Framework' to address the question: *What do non-Bayesian updaters learn?*

A central focus of the literature on Bayesian learning is on what is learned asymptotically and how an agent forecasts as more and more observations are available. Bayesian forecasts are eventually correct with probability 1 under the truth given suitable conditions, the key condition being absolute continuity of the true measure with respect to initial beliefs (see Kalai and Lehrer [3]). Hence, multiple repetitions of Bayes' Rule transforms the historical record into a near perfect guide for the future. We investigate the corresponding question for non-Bayesian updaters who face a statistical inference problem and conform to one of the above noted biases. We show that, like Bayesian updating, multiple repetitions of non-Bayesian updating rules that underreact to observations uncover the true data generating process with probability one. So, non-Bayesian updaters who underreact to the data eventually forecast correctly. Multiple repetitions of non-Bayesian updating rules that overreact to observations uncover the true data generating process with positive probability. However, in some cases, with strictly positive probability, these non-Bayesian updaters become certain that a false parameter is true and thus converge to incorrect forecasts.

Our results cast doubt on the idea that non-Bayesian learning is a transient phenomenon without long-run implications. The agent in our model is self-aware of her updating rules but, regardless of how much data is observed, persistently updates beliefs in a non-Bayesian way. These updating rules do not necessarily lead to a contradiction between forecasts and the data. In the case of under-reaction to observations, the non-Bayesian forecasts are eventually correct and, hence, consistent with the data generating process. This suggests, but does not prove, that in a multi-agent dynamic general equilibrium model, some of our non-Bayesian updaters will not be driven out of the market by Bayesian agents (see Sandroni [5, 6] and Blume and Easley [1] for related results on market selection).

The case in which the agent overreacts to data is quite different. Here our richer hypothesis about updating behavior permits a broader ranger of possible forecasts in the long-run. In particular, with positive probability, these non-Bayesian updaters may permanently forecast incorrectly. Whether these agents can survive in the market is also an open question.

The next section presents our results. Definitions and specifications in Section 3 of the paper are reproduced in this Appendix for the convenience of the reader.

2. LEARNING ABOUT PARAMETERS

This section specializes our model so as to capture the case where the data generating process is unknown up to a parameter $\theta \in \Theta$. In the benchmark Bayesian model, time t beliefs have the form

$$P_t(\cdot) = \int_\Theta \otimes_{t+1}^T \ell(\cdot \mid \theta)\, d\mu_t, \tag{2.1}$$

where: $\ell(\cdot \mid \theta)$ is a likelihood function (measure on S), μ_0 represents prior beliefs on Θ, and μ_t denotes Bayesian posterior beliefs about the parameter at time t and after observations s_1^t. The de Finetti Theorem shows that beliefs admit such a representation if and only if P_0 is exchangeable. We describe (without axiomatic foundations) a generalization of (2.1) that accommodates non-Bayesian updating.

To accommodate parameters, adopt a suitable specification for (p_t, q_t), taking (α_t), δ and u as given. We fix (Θ, ℓ, μ_0) and suppose for now that we are also given a process (ν_t), where each ν_t is a probability measure on Θ. (The σ-algebra associated with Θ is suppressed.) The prior μ_0 on Θ induces time 0 beliefs about S_1 given by

$$p_0(\cdot) = m_0(\cdot) = \int_\Theta \ell(\cdot \mid \theta)\, d\mu_0.$$

Proceed by induction: suppose that μ_t has been constructed and define μ_{t+1} by

$$\mu_{t+1} = \alpha_{t+1} BU(\mu_t; s_{t+1}) + (1 - \alpha_{t+1})\nu_{t+1}, \tag{2.2}$$

where $BU(\mu_t; s_{t+1})(\cdot)$ is the Bayesian update of μ_t. This equation constitutes the *law of motion* for beliefs about parameters. Finally, define (p_{t+1}, q_{t+1}) by

$$p_{t+1}(\cdot) = \int_\Theta \ell(\cdot \mid \theta)\, d\left(BU(\mu_t; s_{t+1})\right) \quad \text{and} \tag{2.3}$$

$$q_{t+1}(\cdot) = \int_\Theta \ell(\cdot \mid \theta)\, d\nu_{t+1}. \tag{2.4}$$

This completes the specification of the model for any given process (ν_t).

Notice that

$$m_{t+1}(\cdot) = \alpha_{t+1} p_{t+1} + (1-\alpha_{t+1}) q_{t+1} = \int_\Theta \ell(\cdot \mid \theta)\, d\mu_{t+1}. \tag{2.5}$$

In light of the discussion in the paper, preferences at $t+1$ are based on the beliefs about parameters represented by μ_{t+1}. If $\alpha_{t+1} \equiv 0$, then (μ_t) is the process of Bayesian posteriors and the above collapses to the exchangeable model (2.1). More generally, differences from the Bayesian model depend on (ν_t), examples of which are given next.[1]

2.1. Prior-Bias with Parameters

Consider first the case where

$$\nu_{t+1} = (1-\lambda_{t+1})\, BU(\mu_t; s_{t+1}) + \lambda_{t+1}\mu_t, \tag{2.6}$$

where $\lambda_{t+1} \leq 1$. This is readily seen to imply prior-bias (see Section 3.1 in the paper); the bias is positive or negative according to the sign of the λ's. Posterior beliefs about parameters satisfy the law of motion

$$\mu_{t+1} = (1-\lambda_{t+1}(1-\alpha_{t+1}))\ \ BU(\mu_t; s_{t+1}) + \lambda_{t+1}(1-\alpha_{t+1})\ \mu_t. \tag{2.7}$$

The latter equation reveals something of how the inferences of an agent with prior-bias differ from those of a Bayesian updater. Compute that (assuming $\alpha_{t+1} \neq 1$)

$$\tfrac{\mu_{t+1}(\theta)}{\mu_{t+1}(\theta')} < \tfrac{\ell(s_{t+1}|\theta)}{\ell(s_{t+1}|\theta')} \tfrac{\mu_t(\theta)}{\mu_t(\theta')} \quad \text{iff} \quad \lambda_{t+1}\ell(s_{t+1} \mid \theta') < \lambda_{t+1}\ell(s_{t+1} \mid \theta). \tag{2.8}$$

[1] One general point is that, in contrast to the exchangeable Bayesian model, μ_{t+1} depends not only on the set of past observations, but also on the order in which they were realized.

For a concrete example, consider coin tossing, with $S = \{H, T\}$, $\Theta \subset (0, 1)$ and $\ell(H \mid \theta) = \theta$ and consider beliefs after a string of H's. If there is positive prior-bias (positive λ's), then repeated application of (2.8) establishes that the agent underinfers in the sense that

$$\frac{\mu_{t+1}(\theta)}{\mu_{t+1}(\theta')} < \frac{\mu^B_{t+1}(\theta)}{\mu^B_{t+1}(\theta')}, \quad \theta > \theta',$$

where μ^B_{t+1} is the posterior of a Bayesian who has the same prior at time 0. Similarly, negative prior-bias leads to overinference.

Turn to the question of what is learned in the long run. Learning may either signify learning the true parameter or learning to forecast future outcomes.[2] The latter kind of learning is more relevant to choice behavior and thus is our focus. Suppose that $\theta^* \in \Theta$ is the true parameter and thus that the i.i.d. measure $P^* = \otimes_{t=1}^{\infty} \ell(\cdot \mid \theta^*)$ is the probability law describing the process (s_t). Say that forecasts are *eventually correct on a path* s_1^∞ if, along that path,

$$m_t(\cdot) \longrightarrow \ell(\cdot \mid \theta^*) \quad \text{as } t \longrightarrow \infty.$$

Rewrite the law of motion for posteriors (2.7) in the form

$$\mu_{t+1} = \left(1 - \gamma_{t+1}\right) BU\left(\mu_t; s_{t+1}\right) + \gamma_{t+1}\mu_t, \tag{2.9}$$

where $\gamma_{t+1} = \lambda_{t+1}(1 - \alpha_{t+1}) \leq 1$. In general, γ_{t+1} is $\mathcal{S}_{t+1}$-measurable (γ_{t+1} may depend on the entire history s_1^{t+1}, including s_{t+1}), but we will be interested also in the special case where γ_{t+1} is $\mathcal{S}_t$-measurable. In that case, (2.9) can be interpreted not only in terms of positive and negative prior-bias as above, but also in terms of underreaction and overreaction to data. For example, let $\gamma_{t+1} \geq 0$ (corresponding to $\lambda_{t+1} \geq 0$). Then μ_{t+1} is a mixture, with weights that are independent of s_{t+1}, of two terms: (i) the Bayesian update $BU\left(\mu_t; s_{t+1}\right)$, which incorporates the 'correct' response to s_{t+1}, and (ii) the prior μ_t, which does not respond to s_{t+1} at all. In a natural sense, therefore, an agent with $\gamma_{t+1} \geq 0$ *underreacts* to data. Similarly, if $\gamma_{t+1} \leq 0$, then $BU\left(\mu_t; s_{t+1}\right)$ is a mixture of μ_{t+1} and μ_t, which suggests that μ_{t+1} reflects *overreaction*. Clearly, if $\gamma_{t+1} = 0$ then the model reduces to the Bayesian updating rule.

Theorem 2.1. *Assume (2.9) and let Θ be finite and $\mu_0(\theta^*) > 0$.*

(a) Suppose that γ_{t+1} is $\mathcal{S}_t$-measurable and that $\gamma_{t+1} \geq 0$. Then forecasts are eventually correct $P^* - a.s.$

[2] See [4] for the distinction between these two kinds of learning.

(b) Suppose that γ_{t+1} is $\mathcal{S}_t$-measurable and that $\gamma_{t+1} \leq 1-\epsilon$ for some $\epsilon > 0$. Then forecasts are eventually correct with P^*-strictly positive probability.

(c) If one drops either of the assumptions in (a), then there exist (S, Θ, ℓ, μ_0) and $\theta \neq \theta^*$ such that

$$m_t(\cdot) \longrightarrow \ell(\cdot \mid \theta) \quad \text{as } t \longrightarrow \infty,$$

with P^*-strictly positive probability.

Assume that before any data are observed the prior belief puts positive weight on the true parameter, that is, assume that $\mu_0(\theta^*) > 0$. Then multiple repetition of Bayes' Rule leads to near correct forecasts. This result is central in the Bayesian literature because it shows that the mere repetition of Bayes' Rule eventually transforms the historical record into a near perfect guide for the future. Part (a) of the theorem generalizes the Bayesian result to the case of underreaction. This result shows that, if repeated sufficiently many times, all non-Bayesian updating rules in (2.9) with the additional proviso of positive prior-bias and the indicated added measurability assumption, eventually produce good forecasting. Hence, in the case of underreaction, agent's forecasts converge to rational expectations although the available information is not processed according to Bayesian laws of probability.

Part (b) shows that, with positive probability, forecasts are eventually correct provided that the Bayesian term on the right side of (2.9) receives weight that is bounded away from zero. This applies in the case of negative prior-bias, corresponding to overreaction. In fact, the results holds even if the forecaster sometimes overreacts and sometimes underreacts to new information. However, part (c) shows that convergence to wrong forecasts may occur in the absence of either of the assumptions in (a). This is demonstrated by two examples. In the first example the weight γ_{t+1} is constant, but sufficiently negative, corresponding to a forecaster that sufficiently overreacts to new information. In the second example, the weight γ_{t+1} is positive corresponding to underreaction, but γ_{t+1} depends on the current signal and, therefore, γ_{t+1} is only $\mathcal{S}_{t+1}$-measurable. In both examples, forecasts may eventually converge to an incorrect limit. Moreover, wrong forecasts in the limit are at least as likely to occur as are correct forecasts.

The proof of Theorem 2.1 builds on classic arguments of the Bayesian literature. Consider the probability measure μ_t on the parameter space and let the random variable μ_t^* be the probability that μ_t assigns to the true parameter. It follows that the expected value (according to the true data generating process) of the

Bayesian update of μ_t^* (given new information) is greater than μ_t^* itself. Hence, in the Bayesian case, the weight given to the true parameter tends to grow as new information is observed. This submartingale property ensures that Bayesian forecasts must converge to some value and cannot remain in endless random fluctuations. The submartingale property follows because under the Bayesian paradigm future changes in beliefs that can be predicted are incorporated in current beliefs. It is immediate from the linear structure in (2.9) that this basic submartingale property still holds in our model as long as the weight γ_{t+1} depends upon the history only up to period t. Hence, with this measurability assumption, forecasts in our model must also converge and, as in the Bayesian case, cannot remain in endless random fluctuations.[3] In addition, convergence to the truth holds in both the Bayesian paradigm and in the case of underreaction. However, given sufficiently strong overreaction, it is possible that forecasts will settle on an incorrect limit. This follows because the positive drift of the above mentioned submartingale property on μ_t^* may be compensated by sufficiently strong volatility which permits that, with positive probability, μ_t^* converges to zero.

2.2. Sample-Bias with Parameters

Learning about parameters is consistent also with sample-bias. Take as primitive a process (ψ_{t+1}) of probability measures on Θ that provides a representation for empirical frequency measures Ψ_{t+1} of the form

$$\Psi_{t+1} = \int \ell(\cdot \mid \theta)\, d\psi_{t+1}(\theta). \tag{2.10}$$

Let μ_0 be given and define μ_{t+1} and ν_{t+1} inductively for $t \geq 0$ by (2.2) and

$$\nu_{t+1} = (1 - \lambda_{t+1})\ BU(\mu_t, s_{t+1}) + \lambda_{t+1}\psi_{t+1}, \tag{2.11}$$

for $\lambda_{t+1} \leq 1$. Then one obtains a special case of sample-bias; the bias is positive or negative according to the sign of the λ's. The implied law of motion for posteriors is

$$\mu_{t+1} = (1 - \lambda_{t+1}(1 - \alpha_{t+1}))\, BU(\mu_t; s_{t+1}) + \lambda_{t+1}(1 - \alpha_{t+1})\ \psi_{t+1}. \tag{2.12}$$

To illustrate, suppose that $S = \{s^1, ..., s^K\}$ and that $\ell\left(s^k \mid \theta\right) = \theta_k$ for each $\theta = (\theta_1, ..., \theta_K)$ in Θ, the interior of the K-simplex. Then one can ensure (2.10) by

[3] We conjecture that beliefs μ_t may not converge in some examples when the weight γ_{t+1} is $\mathcal{S}_{t+1}$-measurable. In our example, it does converge, but to an incorrect limit.

taking ψ_0 to be a suitable noninformative prior; subsequently, Bayesian updating leads to the desired process (ψ_{t+1}). For example, the improper Dirichlet prior density

$$\frac{d\psi_0(\theta)}{\Pi_{k=1}^K d\theta_k} \propto \Pi_{k=1}^K \theta_k^{-1}$$

yields the Dirichlet posterior with parameter vector $\left(n_t\left(s^1\right), ..., n_t\left(s^K\right)\right)$, where $n_t\left(s^k\right)$ equals the number of realizations of s^k in the first t periods; that is,

$$\frac{d\psi_t(\theta)}{\Pi_{k=1}^K d\theta_k} \propto \Pi_{k=1}^K \theta_k^{n_t\left(s^k\right)-1}. \tag{2.13}$$

By the property of the Dirichlet distribution,

$$\int \ell\left(s^k \mid \theta\right) d\psi_t(\theta) = \int \theta_k \, d\psi_t(\theta) = \tfrac{n_k(t)}{t},$$

the empirical frequency of s^k, as required by (2.10).

Finally, compute from (2.12) and (2.13) that (assuming $\alpha_{t+1} \neq 0$)

$$\tfrac{\mu_{t+1}(\theta)}{\mu_{t+1}(\theta')} > \tfrac{\ell(s_{t+1}|\theta)}{\ell(s_{t+1}|\theta')}\tfrac{\mu_t(\theta)}{\mu_t(\theta')} \quad \text{iff} \quad \lambda_{t+1}\tfrac{\psi_t(\theta)}{\psi_t(\theta')} > \lambda_{t+1}\tfrac{\mu_t(\theta)}{\mu_t(\theta')}. \tag{2.14}$$

Suppose that all λ_{t+1}'s are negative (negative sample-bias) and consider the coin-tossing example. As above, we denote by $\left(\mu_t^B\right)$ the Bayesian process of posteriors with initial prior $\mu_0^B = \mu_0$. Then it follows from repeated application of (2.13) and (2.14) that

$$\tfrac{\mu_{t+1}(\theta)}{\mu_{t+1}(\theta')} > \tfrac{\mu_{t+1}^B(\theta)}{\mu_{t+1}^B(\theta')},$$

if $s_1^{t+1} = (H, ..., H)$, $\mid \theta - \frac{1}{2} \mid > \mid \theta' - \frac{1}{2} \mid$ and if the common initial prior μ_0 is uniform.[4] After seeing a string of H's the agent described herein exaggerates (relative to a Bayesian) the relative likelihoods of extremely biased coins. If instead we consider a point at which the history s_1^{t+1} has an equal number of realizations of T and H, then

$$\tfrac{\mu_{t+1}(\theta)}{\mu_{t+1}(1-\theta)} > \tfrac{\theta}{1-\theta}\tfrac{\mu_t(\theta)}{\mu_t(1-\theta)} = \tfrac{BU(\mu_t,H)(\theta)}{BU(\mu_t,H)(1-\theta)},$$

for any θ such that $\mu_t(\theta) > \mu_t(1-\theta)$. If there have been more realizations of H, then the preceding displayed inequality holds if

$$\left(\tfrac{\theta}{1-\theta}\right)^{n_{t+1}(H)-n_{t+1}(T)} < \tfrac{\mu_t(\theta)}{\mu_t(1-\theta)},$$

[4] More generally, the latter two conditions can be replaced by $\frac{\theta'(1-\theta')}{\theta(1-\theta)} > \frac{\mu_0(\theta)}{\mu_0(\theta')}$.

for example, if $\theta < \frac{1}{2}$ and $\mu_t(\theta) \geq \mu_t(1-\theta)$. Note that the bias in this case is towards coins that are less biased ($\theta < \frac{1}{2}$). The opposite biases occur in the case of positive sample-bias.

We conclude with a result regarding learning in the long run. In order to avoid technical issues arising from Θ being a continuum as in the Dirichlet-based model, we consider the following variation: as before $S = \{s^1, ..., s^K\}$ and $\ell\left(s^k \mid \theta\right) = \theta_k$ for each k and θ. But now take Θ to be the set of points $\theta = (\theta_1, ..., \theta_K)$ in the interior of the K-simplex having rational co-ordinates. Define

$$\psi_{t+1}(\theta) = \begin{cases} 1 & \text{if the empirical frequency of } s^k \text{ is } \theta_k,\ 1 \leq k \leq K, \\ 0 & \text{otherwise.} \end{cases}$$

Then (2.10) is evident.[5] The law of motion can be written in the form

$$\mu_{t+1} = \left(1 - \gamma_{t+1}\right) BU\left(\mu_t; s_{t+1}\right) + \gamma_{t+1}\psi_{t+1}, \tag{2.15}$$

where $\gamma_{t+1} = \lambda_{t+1}(1 - \alpha_{t+1}) \leq 1$.

We have the following partial counterpart of Theorem 2.1.

Theorem 2.2. *Let S, (Θ, ℓ) and (ψ_t) be as just defined and suppose that posteriors (μ_t) evolve according to (2.15), where γ_{t+1} is $\mathcal{S}_t$-measurable and $0 < \underline{\gamma} \leq \gamma_{t+1} \leq 1$. Then forecasts are eventually correct $P^* - a.s.$*

The positive lower bound $\underline{\gamma}$ excludes the Bayesian case. The result does hold in the Bayesian case $\gamma_{t+1} = 0$. However, unlike the proof of Theorem 2.1, the proof of Theorem 2.2 is in some ways significantly different from the proof in the Bayesian case. We suspect that the differences in the approach make the lower bound assumption technically convenient but ultimately disposable. We also conjecture (but cannot yet prove) that just as in part (c) of Theorem 2.1, convergence to the truth fails in general if γ_{t+1} is only $\mathcal{S}_{t+1}$-measurable. The other case treated in the earlier theorem - γ_{t+1} is $\mathcal{S}_t$-measurable but possibly negative - (which in the context of that model corresponded to overreaction) is not relevant here because these conditions violate the requirement that each ν_{t+1} in (2.11) be a probability measure and hence non-negatively valued.

[5] If Θ were taken to be finite, then one could not assure (2.10) without admitting signed measures for ψ_{t+1} and hence also for μ_{t+1}. Bayesian updating is not well-defined for signed measures and even if that problem were overcome, the interpretation of such a model is not clear.

3. PROOFS

Proof of Theorem 2.1: *(a)* First we show that $\log \mu_t(\theta^*)$ is a submartingale under P^*. Because

$$\log \mu_{t+1}(\theta^*) - \log \mu_t(\theta^*) = \log \left(\left(1 - \gamma_{t+1}\right) \tfrac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} + \gamma_{t+1} \right), \tag{3.1}$$

it suffices to show that

$$E^* \left[\log \left(\left(1 - \gamma_{t+1}\right) \tfrac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} + \gamma_{t+1} \right) \mid \mathcal{S}_t \right] \geq 0, \tag{3.2}$$

where E^* denotes expectation with respect to P^*. By assumption, γ_{t+1} is constant given $\mathcal{S}_t$. Thus the expectation equals

$$\sum_{s_{t+1}} \ell\left(s_{t+1} \mid \theta^*\right) \log \left(\left(1 - \gamma_{t+1}\right) \tfrac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} + \gamma_{t+1} \right) \geq$$

$$\sum_{s_{t+1}} \ell\left(s_{t+1} \mid \theta^*\right) \left(1 - \gamma_{t+1}\right) \log \left(\tfrac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} \right) =$$

$$\left(1 - \gamma_{t+1}\right) \sum_{s_{t+1}} \ell\left(s_{t+1} \mid \theta^*\right) \log \left(\tfrac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} \right) \geq 0$$

as claimed, where both inequalities are due to concavity of $\log(\cdot)$. (The second is the well-known entropy inequality.)

Clearly $\log \mu_t(\theta^*)$ is bounded above by zero. Therefore, by the martingale convergence theorem, it converges $P^* - a.s.$ From (3.1),

$$\log \mu_{t+1}(\theta^*) - \log \mu_t(\theta^*) = \log \left(\left(1 - \gamma_{t+1}\right) \tfrac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} + \gamma_{t+1} \right) \longrightarrow 0$$

and hence $\frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} \longrightarrow 1 \; P^* - a.s.$ ■

(b) $E^* \left[\left(\left(1 - \gamma_{t+1}\right) \frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} + \gamma_{t+1} \right) \mid \mathcal{S}_t \right] = \left(1 - \gamma_{t+1}\right) E^* \left[\frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} \mid \mathcal{S}_t \right] + \gamma_{t+1} \geq \left(1 - \gamma_{t+1}\right) + \gamma_{t+1} = 1$. (The last inequality is implied by the fact that

$$min_X \left\{ E^* \left[\tfrac{1}{X(s_{t+1})} \mid \mathcal{S}_t \right] : E^* \left[X\left(s_{t+1}\right) \mid \mathcal{S}_t \right] = 1 \right\} = 1.$$

The minimization is over random variable X's, $X : S_{t+1} \longrightarrow \mathbb{R}^1_{++}$, and it is achieved at $X(\cdot) = 1$ because $\frac{1}{x}$ is a convex function on $(0, \infty)$.) Deduce that

$E^* \left[\frac{\mu_{t+1}(\theta^*)}{\mu_t(\theta^*)} \mid \mathcal{S}_t \right] \geq 1$ and hence that $\mu_t(\theta^*)$ is a submartingale. By the martingale convergence theorem,

$$\mu_\infty(\theta^*) \equiv lim \mu_t(\theta^*) \quad \text{exists } P^* - a.s.$$

Claim: $\mu_\infty(\theta^*) > 0$ on a set with positive P^*-probability: By the bounded convergence theorem,

$$E^* \mu_t(\theta^*) \longrightarrow E^* \mu_\infty(\theta^*);$$

and $E^* \mu_t(\theta^*) \nearrow$ because $\mu_t(\theta^*)$ is a submartingale. Thus $\mu_0(\theta^*) > 0$ implies that $E^* \mu_\infty(\theta^*) > 0$, which proves the claim.

It suffices now to show that if $\mu_\infty(\theta^*) > 0$ along a sample path s_1^∞, then forecasts are eventually correct along s_1^∞. But along such a path, $\frac{\mu_{t+1}(\theta^*)}{\mu_t(\theta^*)} \longrightarrow 1$ and hence

$$\left(1 - \gamma_{t+1}\right) \left(\frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} - 1 \right) \longrightarrow 0.$$

By assumption, $\left(1 - \gamma_{t+1}\right)$ is bounded away from zero. Therefore,

$$\left(\frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} - 1 \right) \longrightarrow 0. \qquad \blacksquare$$

Part (c) calls for two examples.

Example 1: Convergence to wrong forecasts may occur with P^*-positive probability when $\gamma_{t+1} < 0$, even where γ_{t+1} is $\mathcal{S}_t$-measurable (overreaction); in fact, we take $(\alpha_{t+1}, \lambda_{t+1}) = (\alpha, \lambda)$ and hence also $\gamma_{t+1} = \gamma$ to be constant over time and states.

Think of repeatedly tossing an unbiased coin that is viewed at time 0 as being either unbiased or having probability of Heads equal to b, $0 < b < \frac{1}{2}$. Thus take $S = \{H, T\}$ and $\ell(H \mid \theta) = \theta$ for $\theta \in \Theta = \{b, \frac{1}{2}\}$. Assume also that

$$1 < -\gamma < \frac{b}{\frac{1}{2} - b}. \tag{3.3}$$

The inequality $\gamma < -1$ indicates a sufficient degree of overreaction.

To explain the reason for the other inequality, note that the model requires that (ν_t) solving (2.6) be a probability measure (hence non-negative valued). This is trivially true if $\lambda_{t+1} \geq 0$ but otherwise requires added restrictions: $\nu_{t+1} \geq 0$ if

$$\frac{\ell\left(s_{t+1} \mid \theta\right)}{m_t\left(s_{t+1}\right)} = \frac{dBU\left(\mu_t; s_{t+1}\right)(\theta)}{d\mu_t} \geq -\frac{\lambda_{t+1}}{1+\lambda_{t+1}}.$$

In the present example $min_{s,\theta} \frac{\ell(s|\theta)}{m_t(s)} \geq 2b$, and thus it suffices to have

$$-\frac{\lambda}{1+\lambda} \leq 2b. \tag{3.4}$$

Because only values for α in $(0,1]$ are admissible, $\gamma = \lambda(1-\alpha)$ is consistent with (3.4) if and only if $-\gamma < b/\left(\frac{1}{2} - b\right)$.

We show that if (3.3), then

$$m_t\left(\cdot\right) \longrightarrow \ell\left(\cdot \mid b\right) \quad \text{as } t \longrightarrow \infty,$$

with probability under P^* at least $\frac{1}{2}$.

Abbreviate $\mu_t\left(\frac{1}{2}\right)$ by μ_t^*.

Claim 1: $\mu_\infty^* \equiv lim \mu_t^*$ exists $P^* - a.s.$ and if $\mu_\infty^* > 0$ for some sample realization s_1^∞, then $m_t\left(H\right) \longrightarrow \frac{1}{2}$ and $\mu_t^* \longrightarrow 1$ along s_1^∞. (The proof is analogous to that of part (b).) Deduce that

$$\mu_\infty^* \in \{0,1\} \quad P^* - a.s.$$

Claim 2: $f\left(z\right) \equiv \left[\left(1-\gamma\right)\frac{\frac{1}{2}}{z} + \gamma\right]\left[\left(1-\gamma\right)\frac{1-\frac{1}{2}}{(1-z)} + \gamma\right] \leq 1$, for all $z \in \left[b, \frac{1}{2}\right]$. Argue that $f\left(z\right) \leq 1 \Longleftrightarrow g\left(z\right) \equiv \left[(1-\gamma) + 2\gamma z\right]\left[(1-\gamma) + 2\gamma(1-z)\right] - 4z\left(1-z\right) \leq 0$. Compute that $g\left(\frac{1}{2}\right) = 0$, $g'\left(\frac{1}{2}\right) = 0$ and g is concave because $\gamma < -1$. Thus $g\left(z\right) \leq g\left(0\right) = 0$.

Claim 3: $E^*\left[\log\left(\left(1-\gamma\right)\frac{\ell\left(s_{t+1}|\frac{1}{2}\right)}{m_t(s_{t+1})} + \gamma\right) \mid \mathcal{S}_t\right]$

$= \frac{1}{2}\log\left(\left(1-\gamma\right)\frac{\frac{1}{2}}{b+\left(\frac{1}{2}-b\right)\mu_t^*} + \gamma\right) + \frac{1}{2}\log\left(\left(1-\gamma\right)\frac{1-\frac{1}{2}}{\left(1-b-\left(\frac{1}{2}-b\right)\mu_t^*\right)} + \gamma\right)$

$= \frac{1}{2}\log\left(f\left(b + \left(\frac{1}{2}-b\right)\mu_t\left(\frac{1}{2}\right)\right)\right) \leq 0$, by Claim 2.

By Claim 1, it suffices to prove that $\mu_\infty^* = 1 \;\; P^* - a.s.$ is impossible. Compute that

$$\mu_t^* = \mu_0^*\left[\Pi_{k=0}^{t-1}\left(\left(1-\gamma\right)\frac{\ell\left(s_{k+1} \mid \frac{1}{2}\right)}{m_k\left(s_{k+1}\right)} + \gamma\right)\right],$$

$$log\mu_t^* = log\mu_0^* + \Sigma_{k=0}^{t-1}\, log\left((1-\gamma)\,\frac{\ell\left(s_{k+1}\mid\frac{1}{2}\right)}{m_k\left(s_{k+1}\right)} + \gamma\right)$$

$$= log\mu_0^* + \Sigma_{k=0}^{t-1}\,\left(logz_{k+1} - E\left[logz_{k+1}\mid\mathcal{S}_k\right]\right) + \Sigma_{k=0}^{t-1}\,E\left[logz_{k+1}\mid\mathcal{S}_k\right],$$

where $z_{k+1} = (1-\gamma)\,\frac{\ell\left(s_{k+1}\mid\frac{1}{2}\right)}{m_k(s_{k+1})} + \gamma$. Therefore, $log\mu_t^* \geq \frac{1}{2}log\mu_0^*$ iff

$$\Sigma_{k=0}^{t-1}\,\left(logz_{k+1} - E\left[logz_{k+1}\mid\mathcal{S}_k\right]\right) \geq -\tfrac{1}{2}log\mu_0^* - \Sigma_{k=0}^{t-1}\,E\left[logz_{k+1}\mid\mathcal{S}_k\right] \equiv a_k.$$

By Claim 3, $a_k > 0$. The random variable $logz_{k+1} - E\left[logz_{k+1}\mid\mathcal{S}_k\right]$ takes on two possible values, corresponding to $s_{k+1} = H$ or T, and under the truth they are equally likely and average to zero. Thus

$$P^*\left(logz_{k+1} - E\left[logz_{k+1}\mid\mathcal{S}_k\right] \geq a_k\right) \leq \tfrac{1}{2}.$$

Deduce that

$$P^*\left(log\mu_t^* \geq \tfrac{1}{2}log\mu_0^*\right) \leq \tfrac{1}{2}$$

and hence that

$$P^*\left(log\mu_t^* \longrightarrow 0\right) \leq \tfrac{1}{2}. \quad \blacksquare$$

Example 2: Convergence to wrong forecasts may occur with P^*-positive probability when $\gamma_{t+1} > 0$ (Positive Prior-Bias), if γ_{t+1} is only $\mathcal{S}_{t+1}$-measurable.

The coin is as before - it is unbiased, but the agent does not know that and is modeled via $S = \{H, T\}$ and $\ell(H \mid \theta) = \theta$ for $\theta \in \Theta = \{b, \frac{1}{2}\}$. Assume further that α_{t+1} and λ_{t+1} are such that

$$\gamma_{t+1} \equiv \lambda_{t+1}(1-\alpha_{t+1}) = \begin{cases} w & \text{if } s_{t+1} = H \\ 0 & \text{if } s_{t+1} = T, \end{cases}$$

where $0 < w < 1$. Thus, from (2.9), the agent updates by Bayes' Rule when observing T but attaches only the weight $(1-w)$ to last period's prior when observing H. Assume that

$$w > 1 - 2b.$$

Then

$$m_t(\cdot) \longrightarrow \ell(\cdot \mid b) \quad \text{as } t \longrightarrow \infty,$$

with probability under P^* at least $\frac{1}{2}$.

The proof is similar to that of Example 1. The key is to observe that $E^*\left[\log\left((1-\gamma)\frac{\ell\left(s_{t+1}|\frac{1}{2}\right)}{m_t(s_{t+1})}+\gamma\right) \mid \mathcal{S}_t\right] \leq 0$ under the stated assumptions.

The proof of Theorem 2.2 requires the following lemmas:

Lemma 3.1. *(Freedman (1975)) Let $\{z_t\}$ be a sequence of uniformly bounded $\mathcal{S}_t$-measurable random variables such that for every $t \geqslant 1$, $E^*(z_{t+1}|\mathcal{S}_t) = 0$. Let $V_t^* \equiv VAR(z_{t+1}|\mathcal{S}_t)$ where VAR is the variance operator associated with P^*. Then,*

$$\sum_{t=1}^{n} z_t \text{ converges to a finite limit as } n \to \infty,\ P^*\text{-a.s. on } \left\{\sum_{t=1}^{\infty} V_t^* < \infty\right\}$$

and

$$\sup_n \sum_{t=1}^{n} z_t = \infty \text{ and } \inf_n \sum_{t=1}^{n} z_t = -\infty,\ P^*\text{-a.s. on } \left\{\sum_{t=1}^{\infty} V_t^* = \infty\right\}.$$

Definition 3.2. *A sequence of $\{x_t\}$ of $\mathcal{S}_t$-measurable random variables is eventually a submartingale if, $P^* - a.s.$, $E^*(x_{t+1}|S_t) - x_t$ is strictly negative at most finitely many times.*

Lemma 3.3. *Let $\{x_t\}$ be uniformly bounded and eventually a submartingale. Then, $P^* - a.s.$, x_t converges to a finite limit as t goes to infinity.*

Proof. Write

$$x_t = \sum_{j=1}^{t}\left(r_j - E^*(r_j|\mathcal{S}_{j-1})\right) + \sum_{j=1}^{t} E^*(r_j|\mathcal{S}_{j-1}) + x_0, \text{ where } r_j \equiv x_j - x_{j-1}.$$

By assumption, $P^* - a.s.$, $E^*(r_j|\mathcal{S}_{j-1})$ is strictly negative at most finitely many times. Hence, $P^* - a.s.$,

$$\inf_t \sum_{j=1}^{t} E^*(r_j|\mathcal{S}_{j-1}) > -\infty.$$

Given that x_t is uniformly bounded, $P^* - a.s.$,

$$\sup_t \sum_{j=1}^{t} z_j < \infty, \text{ where } z_j \equiv r_j - E^*\left(r_j | \mathcal{S}_{j-1}\right).$$

It follows from Freedman's result that $P^* - a.s.$,

$$\sum_{j=1}^{t} z_j \text{ converges to a finite limit as } t \to \infty.$$

It now follows from x_t uniformly bounded that $\sup_t \sum_{j=1}^{t} E^*\left(r_j | \mathcal{S}_{j-1}\right) < \infty$. Because $E^*\left(r_j | \mathcal{S}_{j-1}\right)$ is strictly negative at most finitely many times,

$$\sum_{j=1}^{t} E^*\left(r_j | \mathcal{S}_{j-1}\right) \text{ converges to a finite limit as } t \to \infty.$$

Therefore, $P^* - a.s.$, x_t converges to a finite limit as t goes to infinity. ■

Proof of Theorem 2.2:
Claim 1: Define $f(\theta, m) = \sum_k \theta_k^* \frac{\theta_k}{m_k}$ on the interior of the $2K$-simplex. There exists $\delta' \in \mathbb{R}_{++}^K$ such that

$$\mid \theta_k - \theta_k^* \mid < \delta_k' \text{ for all } k \implies \quad f(\theta, m) - 1 \ \geq\ -\underline{\gamma} K^{-1} \sum_k \mid m_k - \theta_k \mid .$$

Proof: $f(\theta, \theta) = 1$, $f(\theta, \cdot)$ is convex and hence

$$\begin{aligned} f(\theta, m) - 1 \ &\geq\ \sum_{k \neq K} \left(\frac{\partial f(\theta, m)}{\partial m_k} - \frac{\partial f(\theta, m)}{\partial m_K} \right) |_{m=\theta} \ (m_k - \theta_k) \\ &=\ \sum_{k \neq K} \left(-\tfrac{\theta_k^*}{\theta_k} + \tfrac{\theta_K^*}{\theta_K} \right) (m_k - \theta_k). \end{aligned}$$

But the latter sum vanishes at $\theta = \theta^*$. Thus argue by continuity.

Given any $\delta \in \mathbb{R}_{++}^K$, $\delta << \delta'$, define $\Theta^* = (\theta^* - \delta, \theta^* + \delta) \equiv \Pi_{k=1}^K \left(\theta_k^* - \delta_k, \theta_k^* + \delta_k\right)$ and $\mu_t^* = \Sigma_{\theta \in \Theta^*} \mu_t(\theta)$.

Claim 2: Define $m_t^*\left(s^k\right) = \Sigma_{\theta\in\Theta^*}\theta_k\mu_t\left(\theta\right) / \mu_t^*\left(\theta\right)$. Then

$$\mid m_t\left(s^k\right) - m_t^*\left(s^k\right) \mid \leq 1 - \mu_t^*.$$

Proof: $m_t\left(s^k\right) - m_t^*\left(s^k\right) = \frac{\Sigma_{\theta\in\Theta^*}\theta_k\mu_t(\theta)}{\mu_t^*}\left(\mu_t^* - 1\right) + \Sigma_{\theta\notin\Theta^*}\theta_k\mu_t\left(\theta\right).$ Therefore, $\left(\mu_t^* - 1\right) \leq$
$m_t^*\left(s^k\right)\left(\mu_t^* - 1\right) = \frac{\Sigma_{\theta\in\Theta^*}\theta_k\mu_t(\theta)}{\mu_t^*}\left(\mu_t^* - 1\right) \leq m_t\left(s^k\right) - m_t^*\left(s^k\right) \leq \Sigma_{\theta\notin\Theta^*}\theta_k\mu_t\left(\theta\right) \leq 1 - \mu_t^*$.

Claim 3: For any $\delta << \delta'$ as above,

$$\sum_k \theta_k^*\frac{m_t^*\left(s^k\right)}{m_t\left(s^k\right)} - 1 \geq -\underline{\gamma}\left(1 - \mu_t^*\right).$$

Proof: Because $\mid m_t^*\left(s^k\right) - \theta_k^* \mid < \delta_k < \delta_k^{'}$, we have that

$$\sum_k \theta_k^*\frac{m_t^*\left(s^k\right)}{m_t\left(s^k\right)} - 1 \geq -\underline{\gamma}K^{-1}\sum_k \mid m_t\left(s^k\right) - m_t^*\left(s^k\right) \mid .$$

Now Claim 3 follows from Claim 2.

Compute that

$$E^*\left[\mu_{t+1}\left(\theta\right) \mid \mathcal{S}_t\right] = \left(1 - \gamma_{t+1}\right)\left[\sum_k \theta_k^*\frac{\theta_k}{m_t\left(s^k\right)}\right]\mu_t\left(\theta\right) + \gamma_{t+1}E^*\left[\psi_{t+1}\left(\theta\right) \mid \mathcal{S}_t\right], \tag{3.5}$$

where use has been made of the assumption that γ_{t+1} is $\mathcal{S}_t$-measurable. Therefore,

$$E^*\left[\mu_{t+1}^*\left(\theta\right) \mid \mathcal{S}_t\right] - \mu_t^* = \left(1 - \gamma_{t+1}\right)\sum_k\left(\theta_k^*\frac{m_t^*\left(s^k\right)}{m_t\left(s^k\right)}\right)\mu_t^* + \gamma_{t+1}\Sigma_{\theta\in\Theta^*}E^*\left[\psi_{t+1}\left(\theta\right) \mid \mathcal{S}_t\right] - \mu_t^*$$

$$= \left(1 - \gamma_{t+1}\right)\left[\sum_k\left(\theta_k^*\frac{m_t^*\left(s^k\right)}{m_t\left(s^k\right)}\right) - 1\right]\mu_t^* + \gamma_{t+1}\Sigma_{\theta\in\Theta^*}E^*\left[\psi_{t+1}\left(\theta\right) \mid \mathcal{S}_t\right] - \gamma_{t+1}\mu_t^*.$$

By the LLN, $P^* - a.s.$ for large enough t the frequency of s^k will eventually be θ_k^* and

$$\Sigma_{\theta\in\Theta^*}E^*\left[\psi_{t+1}\left(\theta\right) \mid \mathcal{S}_t\right] = 1.$$

Eventually along any such path,

$$E^*\left[\mu^*_{t+1}(\theta)\mid \mathcal{S}_t\right]-\mu^*_t=\left(1-\gamma_{t+1}\right)\left[\sum_k\left(\theta^*_k\frac{m^*_t\left(s^k\right)}{m_t\left(s^k\right)}\right)-1\right]\mu^*_t+\gamma_{t+1}\left(1-\mu^*_t\right)$$

$$\geq\left[-\underline{\gamma}\left(1-\gamma_{t+1}\right)\mu^*_t+\gamma_{t+1}\right]\left(1-\mu^*_t\right)\geq 0,$$

where the last two inequalities follow from Claim 3 and the hypothesis $\underline{\gamma}\leq\gamma_{t+1}$.

Hence (μ^*_t) is eventually a P^*-submartingale. By Lemma 3.3, $\mu^*_\infty\equiv lim\mu^*_t$ exists $P^*-a.s.$ Consequently, $E^*\left[\mu^*_{t+1}(\theta)\mid\mathcal{S}_t\right]-\mu^*_t\longrightarrow 0\ P^*-a.s.$ and from the last displayed equation, $\left[-\underline{\gamma}\left(1-\gamma_{t+1}\right)\mu^*_t+\gamma_{t+1}\right]\left(1-\mu^*_t\right)\longrightarrow 0\ P^*-a.s.$ It follows that $\mu^*_\infty=1$. Finally, $m_t(\cdot)=\int\ell(\cdot\mid\theta)\,d\mu_t$ eventually remains in $\Theta^*=(\theta^*-\delta,\theta^*+\delta)$.

Above δ is arbitrary. Apply the preceding to $\delta=\frac{1}{n}$ to derive a set Ω_n such that $P^*(\Omega_n)=1$ and such that for all paths in Ω_n, m_t eventually remains in $\left(\theta^*-\frac{1}{n},\theta^*+\frac{1}{n}\right)$. Let $\Omega\equiv\cap_{n=1}^\infty\Omega_n$. Then, $P^*(\Omega)=1$ and for all paths in Ω, m_t converges to θ^*. ■

References

[1] L. Blume and D. Easley, If you're so smart, why aren't you rich? Belief selection in complete and incomplete markets, *Econometrica* 74 (2006), 929-966.

[2] D. Freedman, On tail probabilities for martingales, *Annals of Probability* 3 (1975), 100-118.

[3] E. Kalai and E. Lehrer, Rational Learning Leads to Nash equilibrium, *Econometrica* 61 (1993), 1021-1045.

[4] E. Lehrer and R. Smorodinsky, Merging and learning, *Statistics, Probability and Game Theory, IMS Lecture Notes* 30 (1996), 147-168.

[5] A. Sandroni, Do markets favor agents able to make accurate predictions? *Econometrica* 68 (2000), 1303-1341.

[6] A. Sandroni, Efficient markets and Bayes' Rule, *Econ. Theory 26* (2005), 741-764.