

The B.E. Journal of Theoretical Economics

Advances

Volume 10, Issue 1

2010

Article 3

Non-Bayesian Learning

Larry G. Epstein* Jawwad Noor[†]

Alvaro Sandroni[‡]

*Boston University, lepstein@bu.edu

[†]Boston University, jnoor@bu.edu

[‡]University of Pennsylvania, sandroni@kellogg.northwestern.edu

Recommended Citation

Larry G. Epstein, Jawwad Noor, and Alvaro Sandroni (2010) "Non-Bayesian Learning," *The B.E. Journal of Theoretical Economics*: Vol. 10: Iss. 1 (Advances), Article 3.

Non-Bayesian Learning*

Larry G. Epstein, Jawwad Noor, and Alvaro Sandroni

Abstract

A series of experiments suggest that, compared to the Bayesian benchmark, people may either underreact or overreact to new information. We consider a setting where agents repeatedly process new data. Our main result shows a basic distinction between the long-run beliefs of agents who underreact to information and agents who overreact to information. Like Bayesian learners, non-Bayesian updaters who underreact to observations eventually forecast accurately. Hence, underreaction may be a transient phenomenon. Non-Bayesian updaters who overreact to observations eventually forecast accurately with positive probability but may also, with positive probability, converge to incorrect forecasts. Hence, overreaction may have long-run consequences.

KEYWORDS: non-Bayesian learning

*Epstein and Sandroni gratefully acknowledge the financial support of the National Science Foundation (awards SES-0611456 and SES-0820472 and 0922404, respectively).

Erratum

Corrigendum in “Non-Bayesian Learning”

by Larry Epstein, Jawwad Noor and Alvaro Sandroni

Theorem 2.2 in “Non-Bayesian Learning” published in The B.E. Journal of Theoretical Economics: 10-1 (Advances), by Larry Epstein, Jawwad Noor and Alvaro Sandroni is correct, but there is an error in the proof of part *c*). We found this error due to an observation made by Kim-Sau Chung. We gratefully thank him for pointing this difficulty to us. Below we provide a correct argument.

Part *c*) in Theorem 2.2 claims that under suitable conditions multiple repetitions of overreacting updating rules may lead to convergence to a false parameter with strictly positive probability. The mistake in the argument is in the equation

$$\mathbb{P}^* \left(\sum_{k=0}^{t-1} \log z_{k+1} - E^* [\log z_{k+1} \mid \mathcal{S}_k] \geq a_t \right) \leq \frac{1}{2}.$$

that appears at the end of page 10. This equation holds true only if the two values of $y_{k+1} = \log z_{k+1} - E^* [\log z_{k+1} \mid \mathcal{S}_k]$ depended only on the period $k + 1$, but these values are, in fact, history-dependent.

The alternative argument for part *c*) of Theorem 2.2 (below) shows convergence to a false parameter with strictly positive probability as stated in Theorem 2.2 and not convergence to a false parameter with probability greater than 0.5 as incorrectly argued in the original proof.

We maintain the notation and claims 1 – 3 of the original argument. Let

$$\begin{aligned} z_{k+1} &\equiv (1 - \gamma) \frac{\ell(s_{k+1} \mid \frac{1}{2})}{m_k(s_{k+1})} + \gamma; \quad y_{k+1} = \log z_{k+1} - E^* [\log z_{k+1} \mid \mathcal{S}_k] \\ e_{k+1} &= E^* [\log z_{k+1} \mid \mathcal{S}_k] \quad \text{and} \quad v_{k+1} = E^* [(y_{k+1})^2 \mid \mathcal{S}_k]. \end{aligned}$$

Claim 4 : For some real number $x > 0$, $\frac{v_{k+1}}{-e_{k+1}} \leq \frac{1}{x}$.

By claim 3, $-e_{k+1}$ is non-negative (in fact, check that $-e_{k+1}$ is positive if μ_k^* is less than one and zero if μ_k^* is one). In addition, v_{k+1} is limited above because y_{k+1} is uniformly bounded. Let $c > 0$ be such that $|y_{k+1}| \leq c$. Note that e_{k+1} and v_{k+1} converge to zero as μ_k^* goes to one. Using L'Hopital theorem, check that $\frac{v_{k+1}}{-e_{k+1}}$ converges to a finite value as μ_k^* goes to one.

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

$$\begin{aligned} \mu_t^* &= \mu_0^* \left[\prod_{k=0}^{t-1} \left((1 - \gamma) \frac{\ell(s_{k+1} | \frac{1}{2})}{m_k(s_{k+1})} + \gamma \right) \right], \\ \log \mu_t^* &= \log \mu_0^* + \sum_{k=0}^{t-1} \log \left((1 - \gamma) \frac{\ell(s_{k+1} | \frac{1}{2})}{m_k(s_{k+1})} + \gamma \right) \\ &= \log \mu_0^* + \sum_{k=0}^{t-1} (\log z_{k+1} - E^*[\log z_{k+1} | \mathcal{S}_k]) + \sum_{k=0}^{t-1} E^*[\log z_{k+1} | \mathcal{S}_k] \\ &= \log \mu_0^* + \sum_{k=0}^{t-1} y_{k+1} + \sum_{k=0}^{t-1} e_{k+1}. \end{aligned}$$

Therefore, by claim 4, if $\log \mu_t^* \geq \log \mu_0^*$ then

$$\sum_{k=0}^{t-1} y_{k+1} \geq -\sum_{k=0}^{t-1} e_{k+1} \geq x \sum_{k=0}^{t-1} v_{k+1}.$$

Let $y > 0$ be large enough such that $v_1 \geq \frac{1}{y}$ (note that y depends on μ_0^*). By De la Peña (1999) extension of Prokhorov's inequality (Theorem 1.2B equation 1.5),

$$\mathbb{P}^* \left(\frac{\sum_{k=0}^{t-1} y_{k+1}}{\sum_{k=0}^{t-1} v_{k+1}} \geq x \right) \leq \exp\left(-\frac{x^2}{2y(1+cx)}\right).$$

Thus,

$$\mathbb{P}^* (\log \mu_t^* \geq \log \mu_0^*)$$

is bounded away from one and hence

$$\mathbb{P}^* (\log \mu_t^* \rightarrow 0) < 1. \quad \blacksquare$$

Proposition (de la Pena (1999)) Let $\{d_k\}$ be a sequence of uniformly bounded \mathcal{S}_k -measurable random variables such that for every $k \geq 1$, $E(d_k | \mathcal{S}_{k-1}) = 0$ and $|d_k| \leq c$, $0 < c < \infty$. Let $v_k \equiv E((d_k)^2 | \mathcal{S}_{k-1})$. Then,

$$\mathbb{P} \left(\frac{\sum_{k=1}^t d_k}{\sum_{k=1}^t v_k} \geq x, \sum_{k=1}^t v_k \geq \frac{1}{y} \text{ for some } t \right) \leq \exp\left(-\frac{x^2}{2y(1+cx)}\right).$$

References

- [1] de la Peña, V., A General Class of Exponential Inequalities for Martingales and Ratios, *Ann. Probab.* 27 (1999), 537-564.

1. Introduction

A central result in learning theory is that Bayesian forecasts eventually become accurate under suitable conditions, such as absolute continuity of the data generating process with respect to the agent's beliefs (see Kalai and Lehrer, 1993). Hence, multiple repetitions of Bayes' Rule may transform the historical record into a near perfect guide for the future. However, a series of experiments suggest that people may repeatedly process information using non-Bayesian heuristics (see Kahneman and Tversky, 1974, and surveys by Camerer, 1995 and Rabin, 1998). These experiments contributed to a growing interest in the properties of non-Bayesian learning (see, for example, Golub and Jackson, forthcoming, and Gilboa et al., 2008, 2009).¹

Departures from Bayesian updating can occur either because subjects tend to ignore the prior and overreact to the data (we refer to this bias as *overreaction*), or alternatively because subjects place excessive weight on prior beliefs and underreact to new observations (we refer to this bias as *underreaction*).

We investigate a non-Bayesian updater who faces a statistical inference problem and may either overreact or underreact to new data. Consider an agent who is trying to learn the true parameter in a set Θ . Updating of beliefs in response to observations s_1, \dots, s_t , leads to posterior beliefs $\{\mu_t\}$ where each μ_t is a probability measure on Θ . Bayesian updating leads to the process

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

where $BU(\mu_t; s_{t+1})$ denotes the Bayesian update of μ_t given the new observation s_{t+1} . A more general model is the process

$$\mu_{t+1} = (1 - \gamma_{t+1}) BU(\mu_t; s_{t+1}) + \gamma_{t+1} \mu_t, \quad (1.1)$$

where $\gamma_{t+1} \leq 1$. If $\gamma_{t+1} = 0$, the model reduces to standard Bayesian model. If $\gamma_{t+1} > 0$, then the updating rule can be interpreted as attaching too much weight to the prior μ_t and hence underreacting to observations. Conversely, if $\gamma_{t+1} < 0$ then the updating rule can be interpreted as overreacting to observations.²

¹The complexities of Bayesian procedures may make Bayesian updating rules excessively costly to implement in many practical applications. So, even agents who would prefer to use Bayes' Rule often rely on simpler, non-Bayesian heuristics for updating beliefs (see, among others, Bala and Goyal, 1998). Thus, there exists a normative motivation for analyzing non-Bayesian updating rules in addition to the positive motivation coming from laboratory and field experiments (see also Gilboa et al., forthcoming, and Aragonés et al. (2005) for normative motivations on non-Bayesian updating rules).

²Naturally, γ_{t+1} must not be excessively negative so that μ_{t+1} remains a probability measure.

While there may be more than one way to provide a rationale for the non-Bayesian updating rules in (1.1), choice-theoretic foundations of related models were provided in Epstein (2006) and Epstein et al. (2008) in an axiomatic framework where an agent is self-aware of her biases and fully anticipates her updating behavior when formulating plans.³ Foundations take the form of a representation theorem for suitably defined preferences such that both the prior and the way in which it is updated are subjective. In this paper, we describe the asymptotic properties of the process of beliefs defined by (1.1).

We show that, like Bayesian updating, multiple repetitions of non-Bayesian updating rules that underreact to observations eventually lead to accurate forecasts (e.g., forecasts close to the actual data generating process).⁴ Thus, non-Bayesian updaters who underreact to the data eventually forecast accurately. The case in which the agent overreacts to the data is quite different. Multiple repetitions of non-Bayesian updating rules that overreact to the observations eventually lead to accurate forecast with positive probability. In some cases, however, with strictly positive probability, non-Bayesian updaters become certain that a false parameter is true and thus converge to incorrect forecasts. Hence, overreaction may not be a transient phenomenon - it may have long-run implications.

Our results suggest a fundamental difference between underreacting and overreacting to new data. Bayesian and underreacting agents eventually forecast as if they have uncovered the data generating process. However, there is a broader range of possible long-run forecasts for agents who overreact to new observations. These agents may eventually forecast accurately, but they may also permanently forecast incorrectly. Unlike Bayesian and underreacting agents, the ultimate fate of overreacting agents is not entirely pre-determined by the data generating process itself. It also depends on the historical record.

The paper proceeds as follows: In section 2, we present the main concepts and the basic results on overreaction and underreaction. In addition, in section 2, we also consider a bias akin to the hot hand fallacy and show that, like the underreaction bias, agents who update beliefs consistently with this bias also eventually uncover the data generating process. Section 3 concludes and suggests directions for future work. Proofs are in the appendix.

³Hence, an agent may revise her beliefs in a non-Bayesian way even if she is aware of her biases. See also Ali (2009) for a model of learning about self-control.

⁴This result requires an important qualifier. It is also needed that the weights γ_{t+1} for prior and Bayesian posterior depend only on the observations realized until period t and not upon the observation s_{t+1} realized at period $t + 1$.

2. Basic Model and Results

At each period $t = 0, 1, 2, \dots$, one element s_t of (a finite set) S is realized. Thus, the complete uncertainty is given by the state space $\Omega = \prod_{t=1}^{\infty} S_t$, where $S_t = S$ for all $t > 0$. Let $\Delta(S)$ be the set of probability measures on S . A stochastic process (X_t) on Ω is *adapted* if X_t is measurable with respect to the σ -algebra \mathcal{S}_t that is generated by all sets of the form $\{s_1\} \times \dots \times \{s_t\} \times \prod_{t+1}^{\infty} S_t$. Unless otherwise noted, stochastic processes (X_t) on $\prod_{t=1}^{\infty} S_t$ are adapted.

Let Θ denote a countable set of possible parameters. The prior belief over Θ is $\mu_0 \in \Delta(\Theta)$, where $\Delta(\Theta)$ is the set of probability measures over Θ . The σ -algebra associated with Θ is suppressed.

Conditional on parameter θ , at each period $t \geq 0$, an observation $s_t \in S$ is independently generated with a likelihood function $\ell(s_t | \theta)$. Let $\theta^* \in \Theta$ be the parameter determining the data generating process. We define a probability triple $(\Omega, \mathcal{S}, \mathbb{P}^*)$, where \mathcal{S} is the smallest σ -field containing all \mathcal{S}_t for $t > 0$ and \mathbb{P}^* is the probability measure induced over sample paths in Ω by parameter θ^* . That is, $\mathbb{P}^* = \otimes_{t=1}^{\infty} \ell(\cdot | \theta^*)$. Let $\mathbb{E}^*[\cdot]$ be the expectation operator associated with \mathbb{P}^* .

We now define the measures μ_t on the parameter set Θ by induction. The prior μ_0 on Θ is given. Suppose that μ_t is constructed and define μ_{t+1} by

$$\mu_{t+1} = (1 - \gamma_{t+1}) BU(\mu_t; s_{t+1}) + \gamma_{t+1} \mu_t, \quad (2.1)$$

where $BU(\mu_t; s_{t+1})(\cdot)$ is the Bayesian update of μ_t given the new observation s_{t+1} at period $t+1$, and γ_{t+1} is \mathcal{S}_{t+1} -measurable process such that $\gamma_{t+1} \leq 1$. As mentioned in the introduction, γ_{t+1} determine weights given to the Bayesian update of μ_t and the prior belief μ_t at period $t+1$. So, if γ_{t+1} is positive then the posterior belief μ_{t+1} is a mixture of the Bayesian update $BU(\mu_t; s_{t+1})$, which incorporates the Bayesian response to the new observation s_{t+1} , and the prior μ_t , which does not respond to the new observation s_{t+1} at all. In a natural sense, therefore, an agent with weight $\gamma_{t+1} \geq 0$ *underreacts* to data. Similarly, if $\gamma_{t+1} \leq 0$, then the Bayesian update $BU(\mu_t; s_{t+1})$ is a mixture of the posterior belief μ_{t+1} and the prior μ_t , which suggests *overreaction* to new data. Clearly, if $\gamma_{t+1} = 0$ then the model reduces to the Bayesian updating rule. We refer to equation (2.1) as the *law of motion* for beliefs about parameters. Finally, define

$$m_t(\cdot) = \int_{\Theta} \ell(\cdot | \theta) d\mu_t.$$

as the belief at period t over observations at period $t+1$ given measure μ_t over the parameters in Θ .

2.1. Learning with underreaction and overreaction

We now 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.⁵ The latter kind of learning is more relevant to choice behavior and thus is our focus.

Definition 2.1. *Forecasts are eventually accurate on a path $s^\infty \in \Omega$ if, along that path,*

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

That is, if forecasts are eventually accurate then, in the long-run, agents' beliefs converge to the data generating process.⁶

Theorem 2.2. *Assume (2.1). Also assume that $\gamma_t \leq 1 - \epsilon$ for some $\epsilon > 0$ (this assumption rules out the case in which $\gamma_t = 1$ and no learning occurs because prior and posterior always coincide). Let $\mu_0(\theta^*) > 0$ (so, there is a "grain of truth" in prior beliefs).*

(a) *Suppose that $\gamma_{t+1} \geq 0$ for all periods t (i.e., underreaction) and that γ_{t+1} is \mathcal{S}_t -measurable for all periods $t \geq 0$. Then, forecasts are eventually accurate \mathbb{P}^* - a.s.*

(b) *Suppose that γ_{t+1} is \mathcal{S}_t -measurable (both underreaction and overreaction are allowed in different periods). Then forecasts are eventually accurate with \mathbb{P}^* -strictly positive probability.*

(c) *There exists a model (S, Θ, ℓ, μ_0) and weights $\gamma_t = \gamma < 0$ (hence overreaction) and a false parameter $\theta \neq \theta^*$ such that*

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

with \mathbb{P}^ -strictly positive probability. In these cases, the forecast are eventually based on a wrong parameter.*

(d) *Assume that the weights $\gamma_{t+1} \geq 0$ are still nonnegative (i.e., underreaction), but that they may depend upon observations at period $t + 1$ (i.e., γ_{t+1} is \mathcal{S}_{t+1} -measurable). Then, there exist a model (S, Θ, ℓ, μ_0) and a false parameter $\theta \neq \theta^*$ such that*

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

with \mathbb{P}^ -strictly positive probability.*

⁵See Lehrer and Smorodinsky (1996) for the distinction between these two kinds of learning.

⁶This definition is also known in the merging literature as weak merging. (See, among others, Blackwell and Dubins, 1962, Jordan, 1992, Kalai and Lehrer, 1993, 1993b, 1994, Lehrer and Smorodinsky, 1996, 1996b, 1997, Nyarko, 1994, Sandroni, 1998, and Vives, 1993.)

Assume that before any data are observed, prior belief puts positive weight on the true parameter, that is, $\mu_0(\theta^*) > 0$. Then, the basic result in Bayesian updating holds: multiple repetition of Bayes' Rule leads to accurate 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 Theorem 1 generalizes the Bayesian result to underreaction. Multiple repetitions of non-Bayesian updating rules in (2.1) that underreact to the new observations (and the measurability assumption on the weights), eventually produce good forecasting. So, in the case of underreaction, agent's forecasts converge to rational expectations although the available information is not processed by the Bayesian laws of probability.

Part (b) shows that, with positive probability, non-Bayesian forecasts are eventually accurate. This applies to both underreaction and overreaction. Perhaps surprisingly, the results hold even if the forecaster sometimes overreacts and sometimes underreacts to new information.

Parts (c) and (d) are based on examples. The example in part (c) shows that convergence to wrong forecasts may occur for overreactors. The weight γ_t is constant, but negative, corresponding to a forecaster that sufficiently overreacts to new information. In the example, the forecasts converge, but not necessarily to the data generating process. The forecasts may be eventually accurate, but they may also be eventually incorrect (i.e., correspond to a wrong parameter). Hence, whether overreacting updating rules eventually converge to the data generating process may not be pre-determined - it depends upon the realized historical record. In addition, the probability that multiple overreacting updating rules may lead to incorrect forecasts may be greater than 0.5.

The assumption that γ_{t+1} is \mathcal{S}_t -measurable made in parts (a) and (b) of Theorem 1 may seem merely technical, but it is restrictive. The assumption excludes updating rules where the weight on the Bayesian update depends on the current signal. Our example in part (d) shows that this proviso cannot be dropped. In the example, the weight γ_{t+1} is positive corresponding to underreaction, but it depends on the current signal and, therefore, γ_{t+1} is only \mathcal{S}_{t+1} -measurable. As in the case of overreaction, forecasts may eventually converge to an incorrect limit. Moreover, wrong long-run forecasts are at least as likely to occur as are accurate forecasts. Thus, even non-Bayesian updaters who underreact to the new observations may eventually forecast based on a wrong parameter if the weights depend upon the new observation.

2.2. Intuition behind Theorem 2.2

Let $\mu_t(\theta^*)$ be the probability that μ_t assigns to the true parameter θ^* . The expected value (according to the data generating process) of the Bayesian update of $\mu_t(\theta^*)$ (given new information) is greater than $\mu_t(\theta^*)$ itself. This submartingale property ensures that, in the Bayesian case, $\mu_t(\theta^*)$ 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 (2.1) that this basic submartingale property still holds in our model as long as the weight between prior and Bayesian posterior depends upon the history only up to period t . Hence, with this assumption, $\mu_t(\theta^*)$ must also converge and, as in the Bayesian case, cannot remain in endless random fluctuations.⁷ This convergence result holds even if overreaction and underreaction occur in different periods. In the case of underreaction, $\mu_t(\theta^*)$ tends to grow and forecasts are eventually accurate. In the case of 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 $\mu_t(\theta^*)$ may be compensated by sufficiently strong volatility which permits that, with positive probability, $\mu_t(\theta^*)$ converges to zero.

2.3. Sample-Bias

In this section, we consider a bias akin to the *hot-hand fallacy* - the tendency to over-predict the continuation of recent observations (see Kahneman and Tversky, 1971, Camerer, 1989, and Rabin, 2002). Suppose that there are K possible states in each period, $S = \{s^1, \dots, s^K\}$ and that $\ell(s^k | \theta) = \theta_k$ for each parameter $\theta = (\theta_1, \dots, \theta_K)$ in Θ , where Θ is the set of points $\theta = (\theta_1, \dots, \theta_K)$ in the interior of the K -simplex having rational coordinates. Define

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

The law of motion now takes the form

$$\mu_{t+1} = (1 - \gamma_t) BU(\mu_t; s_{t+1}) + \gamma_t \mu_{t+1}, \quad 0 \leq \gamma_t \leq 1. \quad (2.2)$$

If $\gamma_t = 1$, then probability one is assigned to the parameter that coincides with the observed past frequencies, and if $\gamma_t = 0$, the model reduces to the

⁷We conjecture that beliefs $\mu_t(\theta^*)$ 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.

Bayesian model. If $0 < \gamma_t < 1$ posterior beliefs are linear combinations of the Bayesian posterior and the observed frequencies. We have the following partial counterpart of part (a) of Theorem 2.2.

Theorem 2.3. *Suppose that (μ_t) evolve by (2.2), where $0 < \underline{\gamma} \leq \gamma_t \leq 1$. Then forecasts are eventually accurate $\mathbb{P}^* - a.s.$*

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 2.2, the proof of Theorem 2.3 is in some ways significantly different from the standard proof used 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 (d) of Theorem 2.2, convergence to the truth fails in general if the weights γ_{t+1} are allowed to be \mathcal{S}_{t+1} -measurable, instead of being \mathcal{S}_t -measurable as in Theorem 2.3.

3. Conclusion and Directions for Future Work

Our results show an asymmetry between long-run implications of overreaction and underreaction. Multiple revision of beliefs with biases such as underreaction and the sample bias may, like multiple revisions of Bayes' Rule, eventually transform the historical record into a near perfect guide for the future. In the case of overreaction, we show examples such that beliefs may remain incorrect indefinitely. In future work, it is valuable to determine how prevalent is the case of overreaction leading to incorrect forecasts in the long run. In our example, the weight γ_t on the Bayesian update is constant. However, it follows from the proof that, in this example, if the weights γ_t are below a given threshold ($b / (0.5 - b)$), then forecasts may be (i.e., they will be with positive probability) inaccurate in the long-run. A natural conjecture is that, in general, there exists a threshold such that if the weights on the Bayesian update remain below this threshold (i.e., if overreaction is sufficiently severe) then forecasts may be inaccurate in the long-run. However, if the weights on the Bayesian update are always above this threshold (i.e., overreaction, if it occurs, is not too severe) then forecasts will necessarily become accurate in the long-run. The extent to which this conjecture holds is beyond the scope of this paper.

A. Appendix: Proofs

Proof of Theorem 2.2: (a) Given our measurability assumption, we can replace the weights γ_{t+1} with γ_t in parts (a) and (b). First we show that $\log \mu_t(\theta^*)$ is a submartingale under \mathbb{P}^* . Because

$$\log \mu_{t+1}(\theta^*) - \log \mu_t(\theta^*) = \log \left((1 - \gamma_t) \frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} + \gamma_t \right), \quad (\text{A.1})$$

it suffices to show that

$$E^* \left[\log \left((1 - \gamma_t) \frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} + \gamma_t \right) \mid \mathcal{S}_t \right] \geq 0, \quad (\text{A.2})$$

where E^* denotes expectation with respect to \mathbb{P}^* . By assumption, γ_t is constant given \mathcal{S}_t . Thus the expectation equals

$$\begin{aligned} \sum_{s_{t+1}} \ell(s_{t+1} \mid \theta^*) \log \left((1 - \gamma_t) \frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} + \gamma_t \right) &\geq \\ \sum_{s_{t+1}} \ell(s_{t+1} \mid \theta^*) (1 - \gamma_t) \log \left(\frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} \right) &= \\ (1 - \gamma_t) \sum_{s_{t+1}} \ell(s_{t+1} \mid \theta^*) \log \left(\frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} \right) &\geq 0 \end{aligned}$$

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 (A.1),

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

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

Part (b)

$$\begin{aligned} E^* \left[\left((1 - \gamma_t) \frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} + \gamma_t \right) \mid \mathcal{S}_t \right] &= (1 - \gamma_t) E^* \left[\frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} \mid \mathcal{S}_t \right] + \gamma_t \\ &\geq (1 - \gamma_t) + \gamma_t = 1. \end{aligned}$$

The last inequality is implied by the fact that

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

The minimization is over random variable X 's, $X : S_{t+1} \rightarrow \mathbb{R}_{++}^1$, 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,

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

Claim: $\mu_\infty(\theta^*) > 0$ on a set with positive \mathbb{P}^* -probability.

By the bounded convergence theorem,

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

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 $s^\infty \in \Omega$, then forecasts are eventually accurate along s^∞ . But along such a path, $\frac{\mu_{t+1}(\theta^*)}{\mu_t(\theta^*)} \rightarrow 1$ and hence

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

By assumption, $(1 - \gamma_t)$ is bounded away from zero. Therefore,

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

Part (c): Convergence to wrong forecasts may occur with \mathbb{P}^* -positive probability when $\gamma_{t+1} < 0$, even where γ_{t+1} is \mathcal{S}_t -measurable (overreaction); in fact, we take the weight $\gamma_t = \gamma < 0$ to be constant.

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 $\ell(H | \theta) = \theta$ for $\theta \in \Theta = \{b, \frac{1}{2}\}$. Assume also that

$$1 < -\gamma < \frac{b}{\frac{1}{2} - b}. \quad (\text{A.3})$$

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

The other inequality is motivated by the need of having measures non-negative valued in the choice-theoretic model of Epstein (2006) and Epstein et al. (2008) that underlies these laws of motion.

We now show that if (A.3), then

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

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

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

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

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

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

$$\begin{aligned} \text{Claim 3: } E^* \left[\log \left((1-\gamma) \frac{\ell(s_{t+1} | \frac{1}{2})}{m_t(s_{t+1})} + \gamma \right) \mid \mathcal{S}_t \right] \\ = \frac{1}{2} \log \left((1-\gamma) \frac{\frac{1}{2}}{b + (\frac{1}{2}-b)\mu_t^*} + \gamma \right) + \frac{1}{2} \log \left((1-\gamma) \frac{1-\frac{1}{2}}{(1-b - (\frac{1}{2}-b)\mu_t^*)} + \gamma \right) \\ = \frac{1}{2} \log \left(f \left(b + (\frac{1}{2}-b)\mu_t^* \right) \right) \leq 0, \text{ by Claim 2.} \end{aligned}$$

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

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

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

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

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

$$\sum_{k=0}^{t-1} (\log z_{k+1} - E^* [\log z_{k+1} \mid \mathcal{S}_k]) \geq -\frac{1}{2} \log \mu_0^* - \sum_{k=0}^{t-1} E^* [\log z_{k+1} \mid \mathcal{S}_k] \equiv a_t.$$

Let

$$a_t \equiv -\frac{1}{2} \log \mu_0^* - \sum_{k=0}^{t-1} E^* [\log z_{k+1} \mid \mathcal{S}_k].$$

By Claim 3, $a_t > 0$. The random variable $\log z_{k+1} - E^* [\log z_{k+1} \mid \mathcal{S}_k]$ 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

$$\mathbb{P}^* \left(\sum_{k=0}^{t-1} \log z_{k+1} - E^* [\log z_{k+1} \mid \mathcal{S}_k] \geq a_t \right) \leq \frac{1}{2}.$$

Deduce that

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

and hence that

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

Part (d): Convergence to wrong forecasts may occur with \mathbb{P}^* -positive probability when $\gamma_{t+1} > 0$, 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 $\ell(H | \theta) = \theta$ 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 (2.1), 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

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

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

The proof is similar to that of Example in part (c). The key is to observe that $E^* \left[\log \left((1 - \gamma) \frac{\ell(s_{t+1} | \frac{1}{2})}{m_t(s_{t+1})} + \gamma \right) | \mathcal{S}_t \right] \leq 0$ under the stated assumptions (e.g., those concerning λ_{t+1} and α_{t+1}). \blacksquare

The proof of Theorem 2.3 requires the following lemmas:

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

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

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 A.2. A sequence of $\{x_t\}$ of \mathcal{S}_t -measurable random variables is eventually a submartingale if, $\mathbb{P}^* - a.s.$, $E^*(x_{t+1}|S_t) - x_t$ is strictly negative at most finitely many times.

Lemma A.3. Let $\{x_t\}$ be uniformly bounded and eventually a submartingale. Then, $\mathbb{P}^* - a.s.$, x_t converges to a finite limit as t goes to infinity.

Proof. Write

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

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

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

Given that x_t is uniformly bounded, $\mathbb{P}^* - a.s.$,

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

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

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

It now follows from x_t uniformly bounded that $\sup_t \sum_{j=1}^t E^*(r_j|\mathcal{S}_{j-1}) < \infty$.

Because $E^*(r_j|\mathcal{S}_{j-1})$ is strictly negative at most finitely many times,

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

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

Proof of Theorem 2.3:

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

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

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

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

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

Given any $\delta \in \mathbb{R}_{++}^K$, $\delta \ll \delta'$, define

$$\Theta^* = (\theta^* - \delta, \theta^* + \delta) \equiv \Pi_{k=1}^K (\theta_k^* - \delta_k, \theta_k^* + \delta_k)$$

and $\mu_t^* = \sum_{\theta \in \Theta^*} \mu_t(\theta)$.

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

$$|m_t(s^k) - m_t^*(s^k)| \leq 1 - \mu_t^*.$$

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

$$\begin{aligned} (\mu_t^* - 1) &\leq m_t^*(s^k) (\mu_t^* - 1) = \\ \frac{\sum_{\theta \in \Theta^*} \theta_k \mu_t(\theta)}{\mu_t^*} (\mu_t^* - 1) &\leq m_t(s^k) - m_t^*(s^k) \leq \\ \sum_{\theta \notin \Theta^*} \theta_k \mu_t(\theta) &\leq 1 - \mu_t^*. \end{aligned}$$

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

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

Proof: Because $|m_t^*(s^k) - \theta_k^*| < \delta_k < \delta'_k$, we have that

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

Now Claim 3 follows from Claim 2.

Compute that $E^* [\mu_{t+1}(\theta) | \mathcal{S}_t] =$

$$(1 - \gamma_{t+1}) \left[\sum_k \theta_k^* \frac{\theta_k}{m_t(s^k)} \right] \mu_t(\theta) + \gamma_{t+1} E^* [\psi_{t+1}(\theta) | \mathcal{S}_t], \quad (\text{A.4})$$

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

$$\begin{aligned} & E^* [\mu_{t+1}^* | \mathcal{S}_t] - \mu_t^* = \\ & (1 - \gamma_{t+1}) \sum_k \left(\theta_k^* \frac{m_t^*(s^k)}{m_t(s^k)} \right) \mu_t^* + \gamma_{t+1} \sum_{\theta \in \Theta^*} E^* [\psi_{t+1}(\theta) | \mathcal{S}_t] - \mu_t^* \\ & = (1 - \gamma_{t+1}) \left[\sum_k \left(\theta_k^* \frac{m_t^*(s^k)}{m_t(s^k)} \right) - 1 \right] \mu_t^* + \gamma_{t+1} \sum_{\theta \in \Theta^*} E^* [\psi_{t+1}(\theta) | \mathcal{S}_t] - \gamma_{t+1} \mu_t^*. \end{aligned}$$

By the law of large numbers, $\mathbb{P}^* - a.s.$ for large enough t the frequency of s^k will eventually be θ_k^* and

$$\sum_{\theta \in \Theta^*} E^* [\psi_{t+1}(\theta) | \mathcal{S}_t] = 1.$$

Eventually along any such path,

$$\begin{aligned} E^* [\mu_{t+1}^* | \mathcal{S}_t] - \mu_t^* & = (1 - \gamma_{t+1}) \left[\sum_k \left(\theta_k^* \frac{m_t^*(s^k)}{m_t(s^k)} \right) - 1 \right] \mu_t^* + \gamma_{t+1} (1 - \mu_t^*) \\ & \geq [-\underline{\gamma} (1 - \gamma_{t+1}) \mu_t^* + \gamma_{t+1}] (1 - \mu_t^*) \geq 0, \end{aligned}$$

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 A.3, $\mu_\infty^* \equiv \lim \mu_t^*$ exists $\mathbb{P}^* - a.s.$ Consequently, $E^* [\mu_{t+1}^* | \mathcal{S}_t] - \mu_t^* \rightarrow 0$ $\mathbb{P}^* - a.s.$ and from the last displayed equation, $[-\underline{\gamma} (1 - \gamma_{t+1}) \mu_t^* + \gamma_{t+1}] (1 - \mu_t^*) \rightarrow 0$ $\mathbb{P}^* - a.s.$ It follows that $\mu_\infty^* = 1$. Finally, $m_t(\cdot) = \int \ell(\cdot | \theta) d\mu_t$ 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 $\mathbb{P}^*(\Omega_n) = 1$ and such that for all paths in Ω_n , m_t eventually remains in $(\theta^* - \frac{1}{n}, \theta^* + \frac{1}{n})$. Let $\Omega \equiv \cap_{n=1}^\infty \Omega_n$. Then, $\mathbb{P}^*(\Omega) = 1$ and for all paths in Ω , m_t converges to θ^* . ■

References

- [1] Acemoglu D, Nedic A and Ozdaglar A, Convergence of rule-of-thumb learning rules in social networks, *Proceedings of Conference in Decision and Control*, 2007.
- [2] Ali S, Learning self-control, UCSD, 2009.
- [3] Aragonés E, Gilboa I, Postlewaite A and Schmeidler D, Fact-free learning, *Amer. Econ. Rev.* 95 (2005), 1355-1368.
- [4] Bala V, and Goyal S, Learning from neighbours, *Rev. Econ. Stud.* 65 (1998), 595–621.
- [5] Blackwell D, and Dubins L, Merging of opinions with increasing information, *Ann. Math. Stats.* 38 (1962), 882-886.
- [6] Camerer C, Does the basketball market believe in the ‘hot hand’?,” *Amer. Econ. Rev.* 79 (1989), 1257-1261.
- [7] Camerer C, Individual decision-making,” in *Handbook of Experimental Economics*, J. Kagel and A. Roth eds., Princeton U. Press, 1995.
- [8] Epstein LG, An axiomatic model of non-Bayesian updating, *Rev. Econ. Stud.* 73 (2006), 413-436.
- [9] Epstein LG, Noor J, and Sandroni A, Non-Bayesian updating: a theoretical framework, *Theor. Econ.* 3 (2008), 193-229.
- [10] Freedman D, On tail probabilities for martingales, *Ann. Probab.* 3 (1975), 100-118.
- [11] Jordan J, The exponential convergence of Bayesian learning in normal form games, *Games Econ. Behav.* 4 (1992), 202-217.
- [12] Gilboa I, Postlewaite A, and Schmeidler D, Probability and uncertainty in economic modelling, *J. Econ. Perspectives* 22 (2008), 173-188.
- [13] Gilboa I, Postlewaite A, and Schmeidler D, Is it always rational to satisfy Savage’s axioms?,” *Econ. and Phil.* 25 (2009), 285-296.
- [14] Gilboa I, Postlewaite A, and Schmeidler D, Rationality of belief or: why Savage’s axioms are neither necessary nor sufficient for rationality, *Synthese*, forthcoming.

- [15] Golub B, and Jackson M, Naive learning in social networks: convergence, influence, and the wisdom of crowds, *Amer. Econ. J. Microeconomics*, forthcoming.
- [16] Kahneman D, and Tversky A, Belief in the law of small numbers, *Psych. Bull.* 76 (1971), 105-110.
- [17] Kahneman D and Tversky A, Judgement under uncertainty: heuristics and biases, *Science* 185 (1974), 1124-31.
- [18] Kalai E, and Lehrer E, Rational learning leads to Nash equilibrium, *Econometrica* 61 (1993), 1021-1045.
- [19] Kalai E, and Lehrer E, Subjective equilibrium in repeated games, *Econometrica* 61 (1993b), 1231-1240.
- [20] Kalai E, and Lehrer E, Weak and strong merging of opinions, *J. Math. Econ.* 23(1994), 73-100.
- [21] Lehrer E, and Smorodinsky R, Merging and learning,” *Statistics, Probability and Game Theory, IMS Lecture Notes* 30 (1996), 147-168.
- [22] Lehrer E, and Smorodinsky R, Compatible measures and merging, *Math. Oper. Research* 21(1996b), 697-706.
- [23] Lehrer E, and Smorodinsky R, Repeated large games with incomplete information, *Games Econ. Behav.* 18 (1997), 116-134.
- [24] Nyarko R, Bayesian learning in repeated games leads to correlated equilibria in normal form games, *Econ. Theory* 4 (1994), 821-841.
- [25] Rabin M, Psychology and economics, *J. Econ. Lit.* 36 (1998), 11-46.
- [26] Rabin M, Inference by believers in the law of small numbers, *Quart. J. Econ.* 117 (2002), 775-816.
- [27] Sandroni A, Necessary and sufficient conditions for convergence to Nash equilibrium: the almost absolute continuity hypothesis, *Games Econ. Behav.* 22 (1998), 121-147.
- [28] Vives X, How fast do rational agents learn?, *Rev. Econ. Stud.* 60 (1993), 329-347.