

HOW FRAGILE ARE INFORMATION CASCADES?

BY YUVAL PERES¹, MIKLÓS Z. RÁCZ^{2,*}, ALLAN SLY^{2,†} AND IZABELLA STUHL^{3,‡}

¹yuval@yuvalperes.com

²Princeton University, *mracz@princeton.edu; †asly@math.princeton.edu

³Penn State University, ‡ius68@psu.edu

It is well known that sequential decision making may lead to information cascades. That is, when agents make decisions based on their private information, as well as observing the actions of those before them, then it might be rational to ignore their private signal and imitate the action of previous individuals. If the individuals are choosing between a right and a wrong state, and the initial actions are wrong, then the whole cascade will be wrong. This issue is due to the fact that cascades can be based on very little information.

We show that if agents occasionally disregard the actions of others and base their action only on their private information, then wrong cascades can be avoided. Moreover, we study the optimal asymptotic rate at which the error probability at time t can go to zero. The optimal policy is for the player at time t to follow their private information with probability $p_t = c/t$, leading to a learning rate of c'/t , where the constants c and c' are explicit.

1. Introduction. Many everyday situations involve sequential decision making, where one makes a decision based on some private information and the previous actions of others. Consider, for example, the following classroom experiment (see [12], Chapter 16). An experimenter puts an urn containing three balls at the front of the room. This urn is either *majority blue*, containing two blue balls and one yellow ball, or *majority yellow*, containing two yellow balls and one blue ball; both urns are equally likely to be chosen. The students then come to the front of the room one by one. Each student draws a ball at random from the urn and puts it back without showing it to the rest of the class. The student then has to guess the majority color of the urn, announcing her guess publicly. Each student thus makes their decision based on their draw and the announcements of those gone before them.

Let us consider how such an experiment proceeds. The first student only has her own draw to go by, so she will announce the drawn color as her best guess of the majority color. The second student knows this, so together with her own draw she has two independent draws as information. If the colors of the two agree, then the second student announces this color. If the colors of the two draws differ, then she has to use a tie-breaking rule—let us assume that she breaks ties by following her own draw. With this choice we see that the second student also announces the color of the ball she drew. Hence the third student has three independent draws as information and her best guess for the majority color of the urn will be the majority color among the three draws.

Notice that if the first two announced colors were blue, then the third student announces blue *regardless of the color of her draw*. The fourth student knows this and hence the announcement of the third student has no information value. The fourth student is thus in the same situation as the third one and will also just announce blue following the first two students. Following the same logic, all subsequent students will announce blue, regardless of the color of their draw. This phenomenon is known as *herding* or as an *information cascade*.

Received April 2018; revised July 2019.

MSC2020 subject classifications. 91A26, 60C05.

Key words and phrases. Sequential decision making, information cascades, fragility, asymptotic learning, optimal learning rate.

Its study was originated by Banerjee [5] and by Bikhchandani, Hirshleifer, and Welch [7], independently and concurrently; we refer to Easley and Kleinberg [12], Chapter 16, for an exposition.

The main issue with information cascades is that they can be *wrong*. For instance, it might be that the urn is majority yellow, but the first two students draw a blue ball and announce blue, and hence *everyone* announces blue as their guess. The main cause of this is that information cascades can be based on very little information: the actions of a few initial actors can determine all subsequent actions. This also explains why information cascades are *fragile*: if additional information is revealed (e.g., someone reveals not only their action but also their private signal) or if some people deviate from rational behavior, then wrong information cascades can be broken. The focus of this paper is to analyze the fragility of information cascades quantitatively.

To this end, we study a variant of the simple model of sequential decision making studied previously, where not all agents are Bayesian: some are “revealers”, who disregard the actions of others and act solely based on their private signal. This is motivated by both empirical results from laboratory experiments on human behavior in such a setting [4], as well as theoretical considerations [6]; see Section 1.3 for further discussion of related work. We assume that the player at time t is a revealer with probability p_t , independently of everything else, and is a Bayesian otherwise. While agents do not know whether those before them were Bayesians or revealers, this process still introduces additional information that can be useful for making inferences. Are wrong information cascades broken in such a model? That is, do people eventually learn the “correct” action?

We show that the answer is yes: there exist revealing probabilities $\{p_t\}_{t=1}^\infty$ such that learning occurs. Moreover, we study the optimal asymptotic rate at which the error probability at time t can go to zero. We show that the optimal policy is for the player at time t to follow their private information with probability $p_t = c/t$, leading to a learning rate of c'/t , where the constants c and c' are explicit.

1.1. Model and main result. We describe the simplest case of the model first, in order to focus on the conceptual points; we discuss generalizations at the end of the paper. The state of the world is $\theta \in \{1, 2\}$, chosen uniformly at random, that is, $\mathbb{P}(\theta = 1) = \mathbb{P}(\theta = 2) = 1/2$. At times $t = 1, 2, 3, \dots$ players try to guess the state of the world, based on their private information, as well as observing the actions (guesses) of those before them.

The private signals are drawn in the following way. There is an urn that contains two types of balls: type 1 balls are blue and type 2 balls are yellow. Given θ , there are a balls of type θ in the urn and b balls of the other type, where we assume that $a > b > 0$. Each player draws a single ball (with replacement) from the urn, its color is their private signal. In other words, the private signals X_1, X_2, \dots are i.i.d. with the following distribution:

$$\begin{aligned} \mathbb{P}(X_1 = 1 \mid \theta = 1) &= \frac{a}{a+b}, & \mathbb{P}(X_1 = 2 \mid \theta = 1) &= \frac{b}{a+b}, \\ \mathbb{P}(X_1 = 1 \mid \theta = 2) &= \frac{b}{a+b}, & \mathbb{P}(X_1 = 2 \mid \theta = 2) &= \frac{a}{a+b}. \end{aligned}$$

The goal of the players is to guess the majority color (type) of the balls in the urn. We denote the actions (guesses) of the players by Z_1, Z_2, \dots . We assume that each player is one of two kinds:

- a *Bayesian*, whose guess is the maximum a posteriori (MAP) estimate;¹ or

¹We assume that if the posteriors are equal, then a Bayesian follows their private signal. See Appendix B for a discussion on why the choice of tie-breaking rule is immaterial: our results hold regardless of how ties are broken.

- a *revealer*, whose guess is their private signal.

We assume that player t is a revealer with probability p_t , independently of everything else. Formally, let I_1, I_2, \dots be independent Bernoulli random variables (and also independent of everything else) such that $\mathbb{E}[I_t] = p_t$. If $I_t = 0$, then player t is a Bayesian and hence $Z_t = \text{MAP}(Z_1, \dots, Z_{t-1}, X_t)$, while if $I_t = 1$, then player t is a revealer and hence $Z_t = X_t$. Note that players do not know whether the players before them are Bayesians or revealers. We do assume, however, that the players know the probabilities $\{p_t\}_{t=1}^\infty$.

The players aim to learn the majority color/type of the urn, that is, to learn θ , and also to minimize the probability of an incorrect guess. Denote by

$$E_t := \mathbb{P}(Z_t \neq \theta)$$

the probability that the guess of player t is incorrect. We aim to understand the optimal asymptotic rate at which the error probability E_t can go to zero; the following theorem is our main result.

THEOREM 1.1. *Consider the model described above and let*

$$(1.1) \quad \kappa_\star \equiv \kappa_\star(a, b) := \frac{1}{1 + \frac{a/b-1}{\log a/b} (\log(\frac{a/b-1}{\log a/b}) - 1)}.$$

We have that

$$(1.2) \quad \inf_{\{p_t\}_{t=1}^\infty} \limsup_{t \rightarrow \infty} t E_t = \kappa_\star(a, b).$$

That is, the optimal rate of learning is $1/t$, and we obtain the specific constant as well in (1.1). As we shall see, one can get arbitrarily close to the optimum by taking

$$(1.3) \quad p_t = (1 + \varepsilon) \frac{a + b}{b} \cdot \frac{\kappa_\star(a, b)}{t} \wedge 1$$

for $t \geq 1$, where $\varepsilon > 0$ is arbitrary (and where we use the standard notation $x \wedge y := \min\{x, y\}$).

1.2. Heuristic explanation of the optimal rate of learning. We now provide intuition for why $1/t$ is the optimal order for the rate of learning, as well as the reasons behind the constant in (1.1). First, note that the probability that player t is a revealer and draws a ball of the minority color is $\frac{b}{a+b} p_t$. When this occurs, player t guesses incorrectly, implying that $E_t \geq \frac{b}{a+b} p_t$. So in order for the error probability E_t to go to zero, p_t must go to zero as $t \rightarrow \infty$.

On the other hand, p_t cannot go to zero too quickly. If $\sum_{t=1}^\infty p_t < \infty$, then by the Borel–Cantelli lemma there will be only finitely many revealers almost surely. This leads to a situation similar to when there are no revealers: if a correct cascade has not started before the last revealer, then there is a constant probability of ending up in a wrong cascade.

In fact, p_t should decay as $1/t$ to achieve the optimal rate of learning. To see a lower bound of this order, let $p_t = \delta/t$ for δ small. By a Chernoff bound, with high probability there will be at most $2\delta \log t$ revealers among the first t players. If the first two players and all revealers until time t draw balls of the minority color, then so will every player until time t . This event has probability at least $c(b/(a + b))^{2\delta \log t}$ for some constant c , which is greater than $t^{-\varepsilon}$ if $\delta > 0$ is small enough.

To see an upper bound, we now argue that if $p_t = C/t$ with C large enough, then the probability that a wrong cascade (among Bayesians) lasts until time t is $o(1/t)$. Indeed, in such a cascade some of the revealers will be visible—precisely those that deviate from the

cascade consensus. The player at time k has probability $\frac{a}{a+b} p_k$ to be a revealer who draws the majority color, and probability $\frac{b}{a+b} p_k$ to be a revealer who draws the minority color. Hence in a wrong cascade there will be, in expectation, $\frac{a}{a+b} C \log t$ deviations from the cascade consensus by time t , while in a right cascade the expected number of deviations is only $\frac{b}{a+b} C \log t$. The total number of deviations by time t will roughly be Poisson distributed. The probability that a $\text{Poi}(\lambda)$ random variable is larger by a constant factor than its mean is exponentially small in λ and here λ is on the order of $\log t$. Hence by taking C large this exponential in $C \log t$ will be $o(1/t)$.

In fact, the heuristics of the previous paragraph give the right constant as well. To distinguish between right and wrong cascades we need to distinguish between $\text{Poi}(\frac{a}{a+b} C \log t)$ and $\text{Poi}(\frac{b}{a+b} C \log t)$ random variables. The total variation distance between them satisfies

$$(1.4) \quad 1 - \text{TV}\left(\text{Poi}\left(\frac{a}{a+b} C \log t\right), \text{Poi}\left(\frac{b}{a+b} C \log t\right)\right) = t^{-(1+o(1))f(a,b,C)}$$

for some (explicit) function $f(a, b, C)$. The right-hand side of (1.4) is roughly the error probability if player t is a Bayesian. This term should be balanced with the term $\frac{b}{a+b} p_t = \frac{bC}{a+b} \cdot \frac{1}{t}$ coming from player t being a revealer and drawing a ball of the minority color. This balancing requires choosing $C = C(a, b)$ such that $f(a, b, C) = 1$, which occurs when $C = \frac{a+b}{b} \kappa_*(a, b)$, just as in (1.3).

1.3. *Related work.* The special case of the model described in the previous subsection where every agent is Bayesian (i.e., with $p_t = 0$ for every $t \geq 1$) is identical to the model described in the exposition of Easley and Kleinberg [12], Chapter 16. The original model of Bikhchandani, Hirshleifer and Welch [7] differs only in its tie-breaking rule (breaking ties by flipping a fair coin), while that of Banerjee [5] differs in the signal distribution (false signals are drawn from a continuous distribution).² Despite these minor differences, these models all share the same phenomenological behavior as described in the introductory paragraphs.

In particular, Bikhchandani, Hirshleifer and Welch [7] emphasize the *fragility* of information cascades with respect to different types of shocks, as prior work on conforming behavior could not explain this phenomenon. They show examples from numerous fields (e.g., politics, zoology, medicine and finance) where cascades occur and are fragile. The current paper can be viewed as a more detailed quantitative exploration of the fragility of cascades. What amount of additional information is needed to break wrong cascades? What is the optimal rate of learning that can be achieved?

One possible source of additional information comes from people not acting in a rational, Bayesian manner. It is well documented that human behavior is often irrational (see, e.g., [18]). In the information cascades setting, laboratory experiments by Anderson and Holt [4] show that while most participants act rationally, many do not. When deviations from rational behavior occur, participants often act mainly or solely based on their private information, disregarding the information in the actions of those before them.³ Such individuals effectively reveal their private signal, which is valuable information for those coming after them. The model described in Section 1.1, which contains *Bayesians* and *revealers*, captures this empirically observed behavioral phenomenon.

A closely related model was introduced and studied by Bernardo and Welch [6]. This model also contains two types of individuals: (1) rational ones and (2) overconfident ones,

²These models trace back to the works of Cover [10] and Hellman and Cover [17] on sequential learning with finite memory. See [11, 26] for recent work in this direction.

³See also related experiments and results by Çelen and Kariv [8] for a setting with continuous signals.

termed “entrepreneurs”, who put more weight on their private signal than a rational individual would. As the authors mention in their paper, their motivation was not to show that information cascades can be broken by overconfident behavior, but rather to offer a simple explanation for the existence of overconfident individuals based on group selection principles. Nevertheless, since here our focus is on breaking wrong information cascades, we compare our work with theirs in this regard.

In [6] the overconfidence of entrepreneurs is termed “modest” if they still put positive weight on the information from individuals before them, and it is termed “extreme” if they act solely based on their private signal. If the overconfidence of entrepreneurs is modest, then still a wrong cascade occurs with positive probability, bounded away from zero, which is undesirable. Only if the overconfidence of entrepreneurs is extreme—as in the model in Section 1.1—can learning occur eventually with probability one. Bernardo and Welch study their model via simulations which, in the extreme overconfidence setting, suggest that a vanishing fraction of entrepreneurs is optimal. Our work is a rigorous and much more detailed study of this model;⁴ in particular, our results imply that the optimal number of entrepreneurs is *logarithmic* in the size of the group.

The recent work of Cheng, Hann-Caruthers and Tamuz [9] also considers sequential learning models with non-Bayesian agents and shows that wrong cascades can be avoided if there are some non-Bayesian agents. However, they assume (like in [6]) that each agent knows which of the previous agents were revealers, an assumption that we do not make. More importantly, the main contribution of the current paper is the explicit characterization of the optimal rate of learning.

Another possible source of additional information is using the first few agents as “guinea pigs”, that is, forcing them to follow their private signals; see the work of SgROI [24]. Le, Subramanian, and Berry [20] point out that this is related to the multi-armed bandit literature, with guinea pigs corresponding to agents used for exploring [3, 19]. They also mention that it follows from this literature that the optimal number of guinea pigs is logarithmic in the number of agents [3]; this is consistent with our results, albeit in a slightly different setting.

Finally, Wang and Djurić [27] show that a particular form of randomized decision making (instead of Bayesian updating) can also lead to asymptotic learning. Here the additional randomness is chosen in a way to indeed provide additional information. This should be contrasted with the work of Le, Subramanian and Berry [20], who show that randomness in the form of independent noise (coming from observation errors) does not lead to asymptotic learning.

The framework for sequential decision making described in this paper assumes finite discrete private signals. If the informativeness of private signals is unbounded (e.g., Gaussian signals), then wrong cascades do not form and asymptotic learning occurs [25]. In such settings the main question concerns the speed of asymptotic learning; see, for instance, the work of Hann-Caruthers, Martynov and Tamuz [15]. See also the concurrent work of Acemoglu, Makhdomi, Malekian and Ozdaglar [2], who study a model of sequential learning from ratings and reviews; they characterize when asymptotic learning occurs and determine the speed of learning when it does.

The framework of this paper also fits into the broader field of social learning. In particular, there is a large literature on learning in social networks. Acemoglu, Dahleh, Lobel and Ozdaglar [1] consider a model of sequential decision making where agents act only once, but each agent can only observe a *subset* of previous actions, based on a stochastic social

⁴We note that there are small differences in the model studied here and the model of Bernardo and Welch [6]; for instance, in [6] it is assumed that the identities of entrepreneurs are known whereas we do not make this assumption.

network. One of their results is that asymptotic learning occurs even if private signals have bounded informativeness if there are sufficiently many individuals whose neighborhoods are nonpersuasive and hence whose action will necessarily be influenced by their private signal. Similar to revealers in the model described in Section 1.1, these individuals provide a sufficient amount of information for those coming after them to lead to asymptotic learning. Specifying their results to our setting, they imply that asymptotic learning occurs when $\lim_{t \rightarrow \infty} p_t = 0$ and $\sum_{t=1}^{\infty} p_t = \infty$. In comparison, the main contribution of the current paper is the explicit characterization of the optimal rate of learning.

Another typical setting that is studied involves agents who take repeated actions based on their private signal, as well as observing the actions of their neighbors in the network. The main questions include whether or not all agents learn the correct action eventually, what is the speed of learning if it occurs, and how do these depend on the network topology. We highlight recent work of Harel, Mossel, Strack and Tamuz [16], which is similar in spirit to the current paper in that it provides a detailed study of the asymptotic rates of social learning in a mean field setting. A complete overview of the literature is beyond the scope of this article; we refer the reader to the two papers above, as well as to the works of Gale and Kariv [14], Mossel, Sly and Tamuz [21–23], and the references therein for more.

2. Proof of Theorem 1.1. The action of player t can be wrong in two ways: (i) if they act as a Bayesian and the MAP estimator is incorrect, or (ii) if they act on only their private signal and their draw from the urn is the minority type/color. Hence, conditioning on the coin flip deciding whether player t is a Bayesian or a revealer, we obtain that

$$E_t = \mathbb{P}(Z_t \neq \theta \mid I_t = 0)(1 - p_t) + \mathbb{P}(Z_t \neq \theta \mid I_t = 1)p_t.$$

Now recall that given $I_t = 1$, we have $Z_t = X_t$, and so

$$\mathbb{P}(Z_t \neq \theta \mid I_t = 1) = \mathbb{P}(X_t \neq \theta) = \frac{b}{a + b}.$$

Recall also that given $I_t = 0$, we have $Z_t = \text{MAP}(Z_1, \dots, Z_{t-1}, X_t)$. Thus

$$(2.1) \quad E_t = \mathbb{P}(\text{MAP}(Z_1, \dots, Z_{t-1}, X_t) \neq \theta)(1 - p_t) + \frac{b}{a + b}p_t.$$

So we need to understand the probability that the MAP estimator is incorrect at time t . We summarize the behavior of the MAP estimator in Lemma 2.1 and then prove Theorem 1.1 using this, before turning to the proof of the lemma. In the statement of the lemma and throughout the paper we use standard asymptotic notation; for instance $f(t) = o(g(t))$ as $t \rightarrow \infty$ if $\lim_{t \rightarrow \infty} f(t)/g(t) = 0$ and $f(t) = \omega(g(t))$ as $t \rightarrow \infty$ if $\lim_{t \rightarrow \infty} f(t)/g(t) = \infty$.

LEMMA 2.1. *Consider the setting of Theorem 1.1 and fix $\varepsilon > 0$.*

(a) *Suppose that*

$$(2.2) \quad p_t = (1 + \varepsilon) \frac{a + b}{b} \cdot \frac{\kappa_*(a, b)}{t} \wedge 1$$

for every $t \geq 1$. Then

$$(2.3) \quad \mathbb{P}(\text{MAP}(Z_1, \dots, Z_{t-1}, X_t) \neq \theta) = o\left(\frac{1}{t}\right) \text{ as } t \rightarrow \infty.$$

(b) *Suppose that*

$$(2.4) \quad \limsup_{t \rightarrow \infty} tp_t \leq (1 - \varepsilon) \frac{a + b}{b} \kappa_*(a, b).$$

Then

$$(2.5) \quad \mathbb{P}(\text{MAP}(Z_1, \dots, Z_{t-1}, X_t) \neq \theta) = \omega\left(\frac{1}{t}\right) \quad \text{as } t \rightarrow \infty.$$

PROOF OF THEOREM 1.1. Choose $\{p_t\}_{t=1}^\infty$ as in (2.2). Lemma 2.1 says that then (2.3) holds, and hence by (2.1) we have that $\mathbb{E}_t = (1 + o(1))(1 + \varepsilon)\kappa_\star(a, b)/t$ as $t \rightarrow \infty$. Since $\varepsilon > 0$ is arbitrary, we have that

$$(2.6) \quad \inf_{\{p_t\}_{t=1}^\infty} \limsup_{t \rightarrow \infty} t\mathbb{E}_t \leq \kappa_\star(a, b).$$

To show that this is optimal, first note that $\mathbb{E}_t \geq \frac{b}{a+b} p_t$. Hence, if

$$(2.7) \quad \limsup_{t \rightarrow \infty} t\mathbb{E}_t < \kappa_\star(a, b),$$

then the corresponding sequence of probabilities $\{p_t\}_{t=1}^\infty$ must satisfy (2.4) for some $\varepsilon > 0$. But then Lemma 2.1 says that $\mathbb{P}(\text{MAP}(Z_1, \dots, Z_{t-1}, X_t) \neq \theta) = \omega(1/t)$. So by (2.1) we have that $\mathbb{E}_t = \omega(1/t)$, which contradicts (2.7). Thus the inequality in (2.6) is, in fact, an equality. \square

The rest of this section consists of the proof of Lemma 2.1. We start in Section 2.1 by introducing notation and making basic observations about the MAP estimator that are useful for both bounds in Lemma 2.1. Then we turn to the proof of Lemma 2.1(a) in Section 2.2 and we conclude with the proof of Lemma 2.1(b) in Section 2.3.

2.1. *The MAP estimator.* In this subsection we introduce some notation and make basic observations about the MAP estimator that are useful for the bounds in Lemma 2.1, which is proven subsequently.

For $i \in \{1, 2\}$, let \mathbb{P}_i denote the probability measure conditioned on $\theta = i$, that is, $\mathbb{P}_i(\cdot) := \mathbb{P}(\cdot \mid \theta = i)$. Similarly, \mathbb{E}_i denotes expectation conditioned on $\theta = i$. For $i \in \{1, 2\}$ and $t \geq 1$, denote by P_i^t the distribution of (Z_1, \dots, Z_t) . That is, for $(z_1, \dots, z_t) \in \{1, 2\}^t$, let

$$P_i^t(z_1, \dots, z_t) := \mathbb{P}_i(Z_1 = z_1, \dots, Z_t = z_t).$$

Similarly, for $i \in \{1, 2\}$ and $t \geq 1$, denote by Q_i^t the distribution of $(Z_1, \dots, Z_{t-1}, X_t)$. Define also the corresponding likelihoods:

$$L_i^t := P_i^t(Z_1, \dots, Z_t),$$

$$D_i^t := Q_i^t(Z_1, \dots, Z_{t-1}, X_t),$$

with $D_i^0 = L_i^0 = 1$ for $i \in \{1, 2\}$. An outside observer who records the actions of the first t players can compute the likelihoods L_1^t and L_2^t , while player t can compute the likelihoods D_1^t and D_2^t . If player t is a Bayesian, then their guess is based on the likelihoods D_1^t and D_2^t . Specifically, since the prior on θ is uniform, we have that

$$(2.8) \quad \text{MAP}(Z_1, \dots, Z_{t-1}, X_t) = \begin{cases} 1, & \text{if } D_1^t > D_2^t, \\ 2, & \text{if } D_1^t < D_2^t, \\ X_t, & \text{if } D_1^t = D_2^t, \end{cases}$$

where the last line is due to the tie-breaking rule; recall that if the posteriors are equal then a Bayesian follows their private signal.

For $i \in \{1, 2\}$ and $x \in \{1, 2\}$, define

$$\varphi_i(x) := \frac{a}{a+b} \mathbf{1}_{\{x=i\}} + \frac{b}{a+b} \mathbf{1}_{\{x \neq i\}},$$

and note that, since X_t is independent of everything else, we have that $D_i^t = L_i^{t-1}\varphi_i(X_t)$ for $i \in \{1, 2\}$ and $t \geq 1$. Thus in order to understand the likelihoods D_1^t and D_2^t , we need to analyze L_1^t and L_2^t . Define the L -likelihood and the D -likelihood ratios as

$$R_t := \frac{L_1^t}{L_2^t} \quad \text{and} \quad R'_t := \frac{D_1^t}{D_2^t},$$

respectively. We can write

$$(2.9) \quad R'_t = R_{t-1} \frac{\varphi_1(X_t)}{\varphi_2(X_t)}$$

and hence we can determine the action of player t for given R_{t-1} and X_t . Note that the random variable $\varphi_1(X_t)/\varphi_2(X_t)$ takes values in $\{b/a, a/b\}$, and hence we have the following three cases.

- If $R_{t-1} < b/a$, then $R'_t < 1$ and hence $D_1^t < D_2^t$, regardless of the value of X_t . Hence $Z_t = 1$ if player t is a revealer and $X_t = 1$, and $Z_t = 2$ otherwise.
- If $R_{t-1} \in [b/a, a/b]$, then $Z_t = X_t$. This can be checked by considering both cases. If $X_t = 1$ then $\varphi_1(X_t)/\varphi_2(X_t) = a/b$. Therefore by (2.9) we have that $R'_t \geq 1$ and hence $D_1^t \geq D_2^t$. If $D_1^t > D_2^t$ then $Z_t = 1 = X_t$ by the definition of the MAP estimator, while if $D_1^t = D_2^t$ then $Z_t = X_t$ by the tie-breaking rule. The case of $X_t = 2$ is analogous.
- If $R_{t-1} > a/b$, then $R'_t > 1$ and hence $D_1^t > D_2^t$, regardless of the value of X_t . Hence $Z_t = 2$ if player t is a revealer and $X_t = 2$, and $Z_t = 1$ otherwise.

The three cases above describe how the action of a player depends on their private signal and on the actions of those who acted before them. This allows us to analyze how the L -likelihood ratio R_t evolves.

The probability that the MAP estimator makes an error at time t can be expressed using the D -likelihood ratio as follows. To abbreviate the notation for vectors, we write $z_1^t \equiv (z_1, \dots, z_t)$. First, conditioning on the value of θ we obtain that

$$(2.10) \quad \mathbb{P}(\text{MAP}(Z_1^{t-1}, X_t) \neq \theta) = \frac{1}{2}\mathbb{P}_1(\text{MAP}(Z_1^{t-1}, X_t) = 2) + \frac{1}{2}\mathbb{P}_2(\text{MAP}(Z_1^{t-1}, X_t) = 1).$$

The two terms on the right-hand side of (2.10) are equal due to symmetry, so

$$(2.11) \quad \mathbb{P}(\text{MAP}(Z_1^{t-1}, X_t) \neq \theta) = \mathbb{P}_1(\text{MAP}(Z_1^{t-1}, X_t) = 2).$$

Using (2.8) we obtain the following upper and lower bounds:

$$\mathbb{P}_1(R'_t < 1) \leq \mathbb{P}(\text{MAP}(Z_1^{t-1}, X_t) \neq \theta) \leq \mathbb{P}_1(R'_t \leq 1).$$

It is more convenient to work with the L -likelihood ratio, so using the fact that $\frac{b}{a}R_{t-1} \leq R'_t \leq \frac{a}{b}R_{t-1}$ (which again follows from (2.9)) we obtain that

$$(2.12) \quad \mathbb{P}_1\left(R_{t-1} < \frac{b}{a}\right) \leq \mathbb{P}(\text{MAP}(Z_1^{t-1}, X_t) \neq \theta) \leq \mathbb{P}_1\left(R_{t-1} \leq \frac{a}{b}\right).$$

To obtain parts (a) and (b) of Lemma 2.1 we bound from above and below the probabilities appearing in (2.12).

2.2. An upper bound.

PROOF OF LEMMA 2.1(a). By (2.12) our goal is to show that

$$(2.13) \quad \mathbb{P}_1\left(R_t \leq \frac{a}{b}\right) = o\left(\frac{1}{t}\right)$$

as $t \rightarrow \infty$, and recall that we assume that the revealing probabilities $\{p_t\}_{t=1}^\infty$ are as in (2.2).

Let $\{\mathcal{F}_t\}_{t \geq 0}$ denote the filtration defined by the random variables $\{Z_t\}_{t \geq 1}$. Observe that, given $\theta = 1$, the inverse of the L -likelihood ratio, $\{1/R_t\}_{t \geq 0}$, is a martingale with respect to $\{\mathcal{F}_t\}_{t \geq 0}$ (see (2.14) below with $\lambda = 1$; when $\lambda = 1$, the sum over $i \in \{1, 2\}$ is 1). In particular, this implies that $\mathbb{E}_1[R_t^{-1}] = 1$. Since $x \mapsto x^\lambda$ is a concave function for $x \in (0, \infty)$ when $\lambda \in [0, 1]$, we have that, given $\theta = 1$, the sequence $\{R_t^{-\lambda}\}_{t \geq 0}$ is a supermartingale with respect to $\{\mathcal{F}_t\}_{t \geq 0}$. Thus $\mathbb{E}_1[R_t^{-\lambda} | \mathcal{F}_{t-1}] \leq R_{t-1}^{-\lambda}$. We now compute the conditional expectation explicitly:

$$\begin{aligned}
 \mathbb{E}_1[R_t^{-\lambda} | \mathcal{F}_{t-1}] &= R_{t-1}^{-\lambda} \sum_{i \in \{1,2\}} \mathbb{P}_1(Z_t = i | \mathcal{F}_{t-1}) \left(\frac{\mathbb{P}_1(Z_t = i | \mathcal{F}_{t-1})}{\mathbb{P}_2(Z_t = i | \mathcal{F}_{t-1})} \right)^{-\lambda} \\
 (2.14) \qquad &= R_{t-1}^{-\lambda} \sum_{i \in \{1,2\}} \mathbb{P}_1(Z_t = i | \mathcal{F}_{t-1})^{1-\lambda} \mathbb{P}_2(Z_t = i | \mathcal{F}_{t-1})^\lambda.
 \end{aligned}$$

The values of the conditional probabilities in (2.14) depend on the value of R_{t-1} . As described previously, we have three cases:

$$(2.15) \qquad \mathbb{P}_1(Z_t = 1 | \mathcal{F}_{t-1}) = \begin{cases} \frac{a}{a+b} p_t & \text{if } R_{t-1} < \frac{b}{a}, \\ \frac{a}{a+b} & \text{if } R_{t-1} \in \left[\frac{b}{a}, \frac{a}{b} \right], \\ 1 - \frac{b}{a+b} p_t & \text{if } R_{t-1} > \frac{a}{b}, \end{cases}$$

and also

$$(2.16) \qquad \mathbb{P}_2(Z_t = 1 | \mathcal{F}_{t-1}) = \begin{cases} \frac{b}{a+b} p_t & \text{if } R_{t-1} < \frac{b}{a}, \\ \frac{b}{a+b} & \text{if } R_{t-1} \in \left[\frac{b}{a}, \frac{a}{b} \right], \\ 1 - \frac{a}{a+b} p_t & \text{if } R_{t-1} > \frac{a}{b}. \end{cases}$$

Plugging these back into (2.14) we obtain the conditional expectation in the three cases:

$$\begin{aligned}
 (2.17) \qquad \mathbb{E}_1 \left[\left(\frac{R_t}{R_{t-1}} \right)^{-\lambda} \mid R_{t-1} < \frac{b}{a} \right] \\
 = \frac{a^{1-\lambda} b^\lambda}{a+b} p_t + \left(1 - \frac{a}{a+b} p_t \right)^{1-\lambda} \left(1 - \frac{b}{a+b} p_t \right)^\lambda,
 \end{aligned}$$

$$\begin{aligned}
 (2.18) \qquad \mathbb{E}_1 \left[\left(\frac{R_t}{R_{t-1}} \right)^{-\lambda} \mid R_{t-1} \in \left[\frac{b}{a}, \frac{a}{b} \right] \right] \\
 = \frac{a^{1-\lambda} b^\lambda + a^\lambda b^{1-\lambda}}{a+b},
 \end{aligned}$$

$$\begin{aligned}
 (2.19) \qquad \mathbb{E}_1 \left[\left(\frac{R_t}{R_{t-1}} \right)^{-\lambda} \mid R_{t-1} > \frac{a}{b} \right] \\
 = \frac{a^\lambda b^{1-\lambda}}{a+b} p_t + \left(1 - \frac{b}{a+b} p_t \right)^{1-\lambda} \left(1 - \frac{a}{a+b} p_t \right)^\lambda.
 \end{aligned}$$

The right-hand side of (2.18) is strictly less than 1, while the right-hand sides of (2.17) and (2.19) converge to 1 as $t \rightarrow \infty$. To estimate the quantities in (2.17) and (2.19), note that

$(1 - \delta)^\lambda = 1 - \lambda\delta + \Theta(\delta^2)$ as $\delta \rightarrow 0$. Defining

$$f_\lambda \equiv f_\lambda(a, b) := \frac{\lambda a + (1 - \lambda)b - a^\lambda b^{1-\lambda}}{a + b},$$

we have from (2.17)—(2.19) that

$$(2.20) \quad \mathbb{E}_1 \left[\left(\frac{R_t}{R_{t-1}} \right)^{-\lambda} \mid R_{t-1} < \frac{b}{a} \right] = 1 - f_{1-\lambda}(a, b)p_t + O(p_t^2),$$

$$(2.21) \quad \mathbb{E}_1 \left[\left(\frac{R_t}{R_{t-1}} \right)^{-\lambda} \mid R_{t-1} \in \left[\frac{b}{a}, \frac{a}{b} \right] \right] = 1 - (f_\lambda(a, b) + f_{1-\lambda}(a, b)),$$

$$(2.22) \quad \mathbb{E}_1 \left[\left(\frac{R_t}{R_{t-1}} \right)^{-\lambda} \mid R_{t-1} > \frac{a}{b} \right] = 1 - f_\lambda(a, b)p_t + O(p_t^2),$$

as $t \rightarrow \infty$. On the interval $\lambda \in [0, 1]$ the function $\lambda \mapsto f_\lambda$ is concave and nonnegative with $f_0 = f_1 = 0$, it attains its maximum at

$$\lambda_\star \equiv \lambda_\star(a, b) := \frac{\log(\frac{a/b-1}{\log a/b})}{\log a/b},$$

and its maximum value is

$$(2.23) \quad f_{\lambda_\star}(a, b) = \frac{b}{a + b} \cdot \frac{1}{\kappa_\star(a, b)},$$

where recall the definition of κ_\star from (1.1). Note also that $\lambda_\star \in (1/2, 1)$, due to the fact that $a > b$.

We group the cases of (2.20) and (2.21) together, but treat them separately from the case of (2.22), which leads to defining the following random sets:

$$A_t := \left\{ i \in [t] : R_{i-1} > \frac{a}{b} \right\}, \quad B_t := \left\{ i \in [t] : R_{i-1} \leq \frac{a}{b} \right\}.$$

In words, the set A_t is the set of time indices when the MAP estimator is equal to 1 regardless of the private signal at this time. Define also

$$R_t^{(1)} := \prod_{i \in A_t} \frac{R_i}{R_{i-1}}, \quad R_t^{(2)} := \prod_{i \in B_t} \frac{R_i}{R_{i-1}},$$

and note that $R_t = R_t^{(1)} R_t^{(2)}$, since $\{A_t, B_t\}$ is a partition of $[t]$. By (2.22), we have that there exists $C = C(a, b)$ such that for any $\lambda \in [0, 1]$ we have that

$$(2.24) \quad \mathbb{E}_1 \left[\left(\frac{R_t}{R_{t-1}} \right)^{-\lambda} e^{f_\lambda p_t} \mid R_{t-1}, t \in A_t \right] \leq e^{C p_t^2}.$$

Similarly, by (2.20), together with the fact that the right-hand side of (2.20) is greater than the right-hand side of (2.21) for all t large enough, we have that there exists $C' = C'(a, b)$ such that for any $\lambda \in [0, 1]$ we have that

$$(2.25) \quad \mathbb{E}_1 \left[\left(\frac{R_t}{R_{t-1}} \right)^{-\lambda} e^{f_{1-\lambda} p_t} \mid R_{t-1}, t \in B_t \right] \leq e^{C' p_t^2}.$$

Let $M_t := \sum_{i=1}^t p_i$, and note that by the choice of $\{p_t\}_{t=1}^\infty$, together with (2.23), we have that

$$(2.26) \quad M_t = (1 + o(1))(1 + \varepsilon) \frac{a + b}{b} \kappa_\star(a, b) \log t = (1 + o(1)) \frac{(1 + \varepsilon)}{f_{\lambda_\star}(a, b)} \log t$$

as $t \rightarrow \infty$. Define the random variable $\Gamma_t := \sum_{i \in A_t} p_i$. Putting together (2.24) and (2.25), it follows by induction that there exists $C = C(a, b)$ such that for any $\lambda_1, \lambda_2 \in [0, 1]$ we have that

$$(2.27) \quad \mathbb{E}_1[(R_t^{(1)})^{-\lambda_1} e^{f_{\lambda_1} \Gamma_t} (R_t^{(2)})^{-\lambda_2} e^{f_{1-\lambda_2} (M_t - \Gamma_t)}] \leq \exp\left(C \sum_{i=1}^t p_i^2\right).$$

Since $\sum_{i=1}^\infty p_i^2 < \infty$, the expectation in (2.27) is bounded above by a constant independent of t . That is, there exists $C_0 = C_0(a, b) < \infty$ such that

$$(2.28) \quad \mathbb{E}_1[(R_t^{(1)})^{-\lambda_1} e^{f_{\lambda_1} \Gamma_t} (R_t^{(2)})^{-\lambda_2} e^{f_{1-\lambda_2} (M_t - \Gamma_t)}] \leq C_0.$$

Next, we claim that $R_t \leq \frac{a}{b}$ implies that

$$(2.29) \quad R_t^{(1)} \leq 1,$$

and furthermore that there exists a constant $C_1 < \infty$ such that

$$(2.30) \quad \mathbb{P}_1(R_t^{(1)} \leq t^{-C_1}) \leq t^{-2}.$$

We defer the proofs of both of these claims to Appendix A.

From (2.29) we get the following bound on the probability of interest:

$$(2.31) \quad \mathbb{P}_1\left(R_t \leq \frac{a}{b}\right) \leq \mathbb{P}_1\left(R_t \leq \frac{a}{b}, t^{-C_1} \leq R_t^{(1)} \leq 1\right) + \mathbb{P}_1(R_t^{(1)} \leq t^{-C_1}).$$

By (2.30) the second term is at most t^{-2} , so in order to show (2.13) it suffices to bound from above the first term in the display above. We can break the event $\{R_t \leq \frac{a}{b}, t^{-C_1} \leq R_t^{(1)} \leq 1\}$ into subevents based on the value of $R_t^{(1)}$. Recall that $R_t = R_t^{(1)} R_t^{(2)}$, so if $R_t^{(1)} \in [e^{-(x+1)}, e^{-x}]$ and $R_t \leq \frac{a}{b}$, then $R_t^{(2)} \leq \frac{a}{b} e^{x+1}$. Letting $C_3 := 1 + \log \frac{a}{b}$ we obtain the bound

$$(2.32) \quad \begin{aligned} \mathbb{P}_1\left(R_t \leq \frac{a}{b}, t^{-C_1} \leq R_t^{(1)} \leq 1\right) &\leq \sum_{x=0}^{C_1 \log t} \mathbb{P}_1(R_t^{(1)} \in [e^{-(x+1)}, e^{-x}], R_t^{(2)} \leq e^{x+C_3}) \\ &\leq \sum_{x=0}^{C_1 \log t} \mathbb{P}_1(R_t^{(1)} \leq e^{-x}, R_t^{(2)} \leq e^{x+C_3}). \end{aligned}$$

We estimate each term in this sum. First, we can rewrite this probability as follows:

$$\begin{aligned} &\mathbb{P}_1(R_t^{(1)} \leq e^{-x}, R_t^{(2)} \leq e^{x+C_3}) \\ &= \mathbb{P}_1((R_t^{(1)})^{-\lambda_\star} \geq e^{\lambda_\star x}, (R_t^{(2)})^{-(1-\lambda_\star)} \geq e^{-(1-\lambda_\star)(x+C_3)}) \\ &= \mathbb{P}_1((R_t^{(1)})^{-\lambda_\star} e^{f_{\lambda_\star} \Gamma_t} \geq e^{\lambda_\star x + f_{\lambda_\star} \Gamma_t}, \\ &\quad (R_t^{(2)})^{-(1-\lambda_\star)} e^{f_{1-\lambda_\star} (M_t - \Gamma_t)} \geq e^{-(1-\lambda_\star)(x+C_3) + f_{1-\lambda_\star} (M_t - \Gamma_t)}). \end{aligned}$$

If both inequalities hold in the display above, then also the product of the expressions on the left-hand side are greater than or equal to the product of the expressions on the right-hand side. We thus obtain the following bound:

$$\begin{aligned} &\mathbb{P}_1(R_t^{(1)} \leq e^{-x}, R_t^{(2)} \leq e^{x+C_3}) \\ &\leq \mathbb{P}_1((R_t^{(1)})^{-\lambda_\star} e^{f_{\lambda_\star} \Gamma_t} (R_t^{(2)})^{-(1-\lambda_\star)} e^{f_{1-\lambda_\star} (M_t - \Gamma_t)} \geq e^{(2\lambda_\star - 1)x + f_{\lambda_\star} M_t - (1-\lambda_\star)C_3}). \end{aligned}$$

Using Markov’s inequality, together with (2.28) with $\lambda_1 = \lambda_\star$ and $\lambda_2 = 1 - \lambda_\star$, we obtain that

$$(2.33) \quad \begin{aligned} \mathbb{P}_1(R_t^{(1)} \leq e^{-x}, R_t^{(2)} \leq e^{x+C_3}) &\leq C_0 e^{(1-\lambda_\star)C_3} \exp\{-(2\lambda_\star - 1)x - f_{\lambda_\star} M_t\} \\ &\leq C_0 e^{(1-\lambda_\star)C_3} \exp\{-f_{\lambda_\star} M_t\}, \end{aligned}$$

where the second inequality follows from the facts that $x \geq 0$ and $\lambda_\star > 1/2$. Recalling from (2.26) that $f_{\lambda_\star} M_t = (1 + o(1))(1 + \varepsilon) \log t$, and using (2.32) and (2.33), we arrive at the following bound:

$$\mathbb{P}_1\left(R_t \leq \frac{a}{b}, t^{-C_1} \leq R_t^{(1)} \leq 1\right) \leq \frac{C_4 \log t}{t^{(1+o(1))(1+\varepsilon)}}$$

for some constant $C_4 < \infty$. Putting this together with (2.31) and (2.30) we obtain (2.13). \square

2.3. A lower bound.

PROOF OF LEMMA 2.1(b). By (2.12) our goal is to show that

$$(2.34) \quad \mathbb{P}_1\left(R_t < \frac{b}{a}\right) = \omega\left(\frac{1}{t}\right)$$

as $t \rightarrow \infty$, and recall that we assume that the revealing probabilities $\{p_t\}_{t=1}^\infty$ satisfy (2.4). As in the proof of the upper bound, let $M_t := \sum_{i=1}^t p_i$ and note that (2.4) implies that

$$(2.35) \quad M_t \leq (1 + o(1))(1 - \varepsilon) \frac{a + b}{b} \kappa_\star(a, b) \log t = (1 + o(1)) \frac{1 - \varepsilon}{f_{\lambda_\star}(a, b)} \log t$$

as $t \rightarrow \infty$, where recall the definition of f_λ and λ_\star from Section 2.2. We may also assume that

$$(2.36) \quad M_t \geq \delta \log t$$

for all t large enough, where $\delta \equiv \delta(a, b) := \frac{1}{2 \log(1+a/b)}$; if $M_t < \delta \log t$ then a simple argument shows that (2.5) holds, which we defer to Appendix A.3. Define also $\tau(s) := \min\{t \geq 1 : M_t \geq s\}$, and note that (2.35) and (2.36) imply that $\exp((1 + o(1)) \frac{f_{\lambda_\star}}{1-\varepsilon} s) \leq \tau(s) \leq \exp(s/\delta)$ as $s \rightarrow \infty$. In particular $\tau(s) < \infty$ for every $s < \infty$.

In order to show (2.34), we define three events that together imply that $R_t < b/a$ and show that the probability that they all occur, given $\theta = 1$, is $\omega(1/t)$ as $t \rightarrow \infty$. First, let $t_0 := \tau(2 \frac{a+b}{a-b} \log(a/b) + 2)$ and define

$$(2.37) \quad \mathcal{E}_0 := \left\{R_{t_0} \leq \left(\frac{b}{a}\right)^4\right\}.$$

This initial event takes the L -likelihood ratio below b/a , and the events we now define ensure that it stays below b/a . Let $J_t := \log R_t$ denote the L -log-likelihood ratio and define the stopping time $T := \min\{s \geq t_0 : J_s \notin [-\log t, 2 \log \frac{b}{a}]\}$. Define now the events

$$\begin{aligned} \mathcal{E}_1 &:= \{J_T \leq -\log t\}, \\ \mathcal{E}_2 &:= \left\{\min_{s \in [t]} J_s \geq -\log^{3/4} t\right\}. \end{aligned}$$

Observe that $\mathcal{E}_0, \mathcal{E}_1$ and \mathcal{E}_2 together imply that $J_s \in [-\log^{3/4} t, 2 \log \frac{b}{a}]$ for all $s \in [t_0, t]$. In particular, they imply that $R_t < b/a$ and so

$$(2.38) \quad \mathbb{P}_1\left(R_t < \frac{b}{a}\right) \geq \mathbb{P}_1(\mathcal{E}_0 \cap \mathcal{E}_1 \cap \mathcal{E}_2).$$

In what follows we show that the right-hand side of the display above is $\omega(1/t)$ as $t \rightarrow \infty$.

We start with the initial event \mathcal{E}_0 . Note that the first two individuals follow their private signal, that is, $Z_1 = X_1$ and $Z_2 = X_2$, and hence if $X_1 = X_2 = 2$, then $R_2 = (b/a)^2$. If $X_i = 2$ for all $i \in \{3, 4, \dots, t_0\}$ then also $Z_i = 2$ for all $i \in \{3, 4, \dots, t_0\}$, regardless of whether the corresponding players are Bayesians or revealers. Consequently, by (2.15) and (2.16), in this event the L -likelihood ratio at time t_0 is equal to

$$(2.39) \quad R_{t_0} = \left(\frac{b}{a}\right)^2 \prod_{i=3}^{t_0} \frac{1 - \frac{a}{a+b} p_i}{1 - \frac{b}{a+b} p_i}.$$

Now using

$$\frac{1 - \frac{a}{a+b} p_i}{1 - \frac{b}{a+b} p_i} = 1 - \frac{\frac{a-b}{a+b} p_i}{1 - \frac{b}{a+b} p_i} \leq 1 - \frac{a-b}{a+b} p_i \leq e^{-\frac{a-b}{a+b} p_i},$$

and also, by the definition of t_0 , the fact that

$$\sum_{i=3}^{t_0} p_i \geq M_{t_0} - 2 \geq 2 \frac{a+b}{a-b} \log \frac{a}{b},$$

we obtain from (2.39) that $R_{t_0} \leq (b/a)^4$. Hence

$$\mathbb{P}_1(\mathcal{E}_0) \geq \mathbb{P}_1(X_i = 2 \ \forall i \in [t_0]) = \left(\frac{b}{a+b}\right)^{t_0}.$$

Since t_0 is a constant, we have that $\mathbb{P}_1(\mathcal{E}_0)$ is strictly bounded away from zero. Thus by (2.38) it suffices to show that

$$(2.40) \quad \mathbb{P}_1(\mathcal{E}_1 \cap \mathcal{E}_2 \mid \mathcal{E}_0) = \omega\left(\frac{1}{t}\right)$$

as $t \rightarrow \infty$.

We now turn to estimating the probability of \mathcal{E}_1 and \mathcal{E}_2 , given \mathcal{E}_0 . Note that, given \mathcal{E}_0 , the L -log-likelihood ratio performs a random walk from time t_0 until the stopping time T . Specifically, for $s \geq t_0$ we can write

$$J_{s \wedge T} = J_{t_0} + \sum_{i=t_0+1}^{s \wedge T} \xi_i,$$

where the random variables $\{\xi_i\}_{i>t_0}$ are independent (of each other and everything else) with the following distribution under \mathbb{P}_1 :

$$\begin{aligned} \mathbb{P}_1\left(\xi_i = \log \frac{a}{b}\right) &= \frac{a}{a+b} p_i, \\ \mathbb{P}_1\left(\xi_i = \log \frac{1 - \frac{a}{a+b} p_i}{1 - \frac{b}{a+b} p_i}\right) &= 1 - \frac{a}{a+b} p_i. \end{aligned}$$

Note, in particular, that the random variables $\{\xi_i\}_{i>t_0}$ are uniformly bounded: $\log \frac{b}{a} \leq \xi_i \leq \log \frac{a}{b}$ for all $i > t_0$. Furthermore, we have that

$$\mathbb{E}_1[\xi_i] = \frac{1 + \frac{a}{b}(\log \frac{a}{b} - 1)}{1 + \frac{a}{b}} p_i + \Theta(p_i^2),$$

showing that the L -log-likelihood ratio has an upward drift in this regime. We perform a change of measure to remove this drift: define $\tilde{\mathbb{P}}_1$ such that under $\tilde{\mathbb{P}}_1$ the random variables

$\{\xi_i\}_{i>t_0}$ have expectation zero. That is, define $\tilde{\mathbb{P}}_1$ such that for all $i > t_0$ we have that

$$\begin{aligned} \tilde{\mathbb{P}}_1\left(\xi_i = \log \frac{a}{b}\right) &= q_i, \\ \tilde{\mathbb{P}}_1\left(\xi_i = \log \frac{1 - \frac{a}{a+b} p_i}{1 - \frac{b}{a+b} p_i}\right) &= 1 - q_i, \end{aligned}$$

where q_i is chosen such that $\tilde{\mathbb{E}}_1[\xi_i] = 0$. A short computation gives that

$$q_i = \frac{\log \frac{1 - \frac{b}{a+b} p_i}{1 - \frac{a}{a+b} p_i}}{\log\left(\frac{a}{b} \cdot \frac{1 - \frac{b}{a+b} p_i}{1 - \frac{a}{a+b} p_i}\right)}.$$

In the following we first estimate $\tilde{\mathbb{P}}_1(\mathcal{E}_1 \cap \mathcal{E}_2 \mid \mathcal{E}_0)$ and then show (2.40) by understanding the Radon–Nikodym derivative of $\mathbb{P}_1(\cdot \mid \mathcal{E}_0)$ with respect to $\tilde{\mathbb{P}}_1(\cdot \mid \mathcal{E}_0)$.

Under $\tilde{\mathbb{P}}_1$ we have that $\{J_{s \wedge T}\}_{s \geq t_0}$ is a martingale. By the optional stopping theorem we thus have that

$$(2.41) \quad \tilde{\mathbb{E}}_1[J_T \mid \mathcal{E}_0] = \tilde{\mathbb{E}}_1[J_{t_0}].$$

By the definition of T we have that either $J_T > 2 \log \frac{b}{a}$, in which case $J_T \in (2 \log \frac{b}{a}, \log \frac{b}{a}]$, or $J_T < -\log t$, in which case $J_T \in [-\log t - \log \frac{a}{b}, -\log t)$. Hence we have that

$$\begin{aligned} \tilde{\mathbb{E}}_1[J_T \mid \mathcal{E}_0] &= \tilde{\mathbb{E}}_1[J_T \mathbf{1}_{\{J_T < -\log t\}} \mid \mathcal{E}_0] + \tilde{\mathbb{E}}_1[J_T \mathbf{1}_{\{J_T > 2 \log(b/a)\}} \mid \mathcal{E}_0] \\ (2.42) \quad &\geq \tilde{\mathbb{P}}_1(\mathcal{E}_1 \mid \mathcal{E}_0) \times \left(-\log t - \log \frac{a}{b}\right) + (1 - \tilde{\mathbb{P}}_1(\mathcal{E}_1 \mid \mathcal{E}_0)) \times 2 \log \frac{b}{a} \\ &\geq 2 \log \frac{b}{a} + \tilde{\mathbb{P}}_1(\mathcal{E}_1 \mid \mathcal{E}_0) \times (-\log t). \end{aligned}$$

From the definition of t_0 we also have that

$$(2.43) \quad \tilde{\mathbb{E}}_1[J_{t_0}] \leq 4 \log \frac{b}{a}.$$

Thus putting together (2.41), (2.42), and (2.43) we obtain that

$$(2.44) \quad \tilde{\mathbb{P}}_1(\mathcal{E}_1 \mid \mathcal{E}_0) \geq \frac{2 \log \frac{a}{b}}{\log t}.$$

Applying Theorem 1.6 in [13] we have that

$$\tilde{\mathbb{P}}_1\left(\min_{s \in [t]} J_s < -\log^{3/4} t \mid \mathcal{E}_0\right) \leq \exp(-c \log^{1/2} t)$$

for some constant $c > 0$. Together with (2.44) this shows that

$$(2.45) \quad \tilde{\mathbb{P}}_1(\mathcal{E}_1 \cap \mathcal{E}_2 \mid \mathcal{E}_0) \geq \frac{\log \frac{a}{b}}{\log t}$$

for all t large enough.

Since we want to show (2.40), what remains is to estimate the Radon–Nikodym derivative of $\mathbb{P}_1(\cdot \mid \mathcal{E}_0)$ with respect to $\tilde{\mathbb{P}}_1(\cdot \mid \mathcal{E}_0)$ on the sigma-algebra \mathcal{F}_t and on the event $\mathcal{E}_1 \cap \mathcal{E}_2$. From the definition of $\tilde{\mathbb{P}}_1$ we can write this down explicitly: we have that

$$\begin{aligned} (2.46) \quad &\frac{d\mathbb{P}_1(\cdot \mid \mathcal{E}_0)}{d\tilde{\mathbb{P}}_1(\cdot \mid \mathcal{E}_0)} \Big|_{\mathcal{F}_t} \mathbf{1}_{\{\mathcal{E}_1 \cap \mathcal{E}_2\}} \\ &= \prod_{i=t_0+1}^t \left\{ \frac{a}{a+b} p_i \mathbf{1}_{\{\xi_i = \log \frac{a}{b}\}} + \frac{1 - \frac{a}{a+b} p_i}{1 - q_i} \mathbf{1}_{\{\xi_i = \log \frac{1 - \frac{a}{a+b} p_i}{1 - \frac{b}{a+b} p_i}\}} \right\} \mathbf{1}_{\{\mathcal{E}_1 \cap \mathcal{E}_2\}}. \end{aligned}$$

We estimate each factor in this product. Specifically, we claim that there exists a constant $C = C(a, b)$ such that for every $i \in \{t_0 + 1, \dots, t\}$ the following holds:

$$(2.47) \quad \frac{\frac{a}{a+b} p_i}{q_i} \mathbf{1}_{\{\xi_i = \log \frac{a}{b}\}} + \frac{1 - \frac{a}{a+b} p_i}{1 - q_i} \mathbf{1}_{\{\xi_i = \log \frac{1 - \frac{a}{a+b} p_i}{1 - \frac{a}{a+b} p_i}\}} \geq \exp((1 - \lambda_\star) \xi_i - f_{\lambda_\star} p_i - C p_i^2).$$

This inequality can be checked for both potential values of ξ_i by expanding the expressions in p_i . When $\xi_i = \log \frac{1 - \frac{a}{a+b} p_i}{1 - \frac{a}{a+b} p_i}$, both sides of (2.47) are equal to $1 + \frac{a-b-a \log \frac{a}{b}}{(a+b) \log \frac{a}{b}} p_i + \Theta(p_i^2)$, so by choosing C large enough, the quadratic term on the right-hand side will be smaller than that on the left-hand side, and hence (2.47) holds in this case if C is large enough. When $\xi_i = \log \frac{a}{b}$, the left-hand side of (2.47) is equal to $\frac{a \log \frac{a}{b}}{a-b} + \frac{a \log \frac{a}{b}}{a-b} (\frac{a-b}{(a+b) \log \frac{a}{b}} - \frac{1}{2}) p_i + \Theta(p_i^2)$, while the right-hand side of (2.47) is equal to $\frac{a \log \frac{a}{b}}{a-b} - \frac{a \log \frac{a}{b}}{a-b} f_{\lambda_\star} p_i + \Theta(p_i^2)$. Thus the constant terms are equal and it can be checked that the coefficient of the first order term is greater for the expression on the left-hand side than for the expression on the right-hand side. Hence in this case (2.47) holds for all i large enough regardless of the value of C , and if C is chosen large enough then it holds for all $i \in \{t_0 + 1, \dots, t\}$.

After justifying (2.47), we can now turn back to estimating the quantity in (2.46) by multiplying (2.47) over all $i \in \{t_0 + 1, \dots, t\}$. Using the fact that $J_t = J_{t_0} + \sum_{i=t_0+1}^t \xi_i$ on the event $\mathcal{E}_1 \cap \mathcal{E}_2$, and recalling also that $M_t = \sum_{i=1}^t p_i$, we obtain that

$$(2.48) \quad \begin{aligned} & \frac{d\mathbb{P}_1(\cdot | \mathcal{E}_0)}{d\tilde{\mathbb{P}}_1(\cdot | \mathcal{E}_0)} \Big|_{\mathcal{F}_t} \mathbf{1}_{\{\mathcal{E}_1 \cap \mathcal{E}_2\}} \\ & \geq \exp\left((1 - \lambda_\star)(J_t - J_{t_0}) - f_{\lambda_\star}(M_t - M_{t_0}) - C \sum_{i=t_0+1}^t p_i^2 \right) \mathbf{1}_{\{\mathcal{E}_1 \cap \mathcal{E}_2\}} \\ & \geq \exp((1 - \lambda_\star) J_t - f_{\lambda_\star} M_t - C') \mathbf{1}_{\{\mathcal{E}_1 \cap \mathcal{E}_2\}}, \end{aligned}$$

where $C' := C \sum_{i=1}^\infty p_i^2$, and where in the second inequality we used that $J_{t_0} < 0$ (given \mathcal{E}_0) and that $\sum_{i=1}^\infty p_i^2 < \infty$ due to the assumption (2.4). Now recall from (2.35) that $f_{\lambda_\star} M_t \leq (1 + o(1))(1 - \varepsilon) \log t$ as $t \rightarrow \infty$. Recall also that on the event \mathcal{E}_2 we have that $J_t \geq -\log^{3/4} t$. Plugging these two estimates into the right-hand side of (2.48) we obtain that

$$\frac{d\mathbb{P}_1(\cdot | \mathcal{E}_0)}{d\tilde{\mathbb{P}}_1(\cdot | \mathcal{E}_0)} \Big|_{\mathcal{F}_t} \mathbf{1}_{\{\mathcal{E}_1 \cap \mathcal{E}_2\}} \geq \exp(-(1 - \lambda_\star) \log^{3/4} t - C') \times t^{-(1+o(1))(1-\varepsilon)} \mathbf{1}_{\{\mathcal{E}_1 \cap \mathcal{E}_2\}}.$$

Consequently we obtain that

$$\begin{aligned} \mathbb{P}_1(\mathcal{E}_1 \cap \mathcal{E}_2 | \mathcal{E}_0) &= \tilde{\mathbb{E}}_1 \left[\frac{d\mathbb{P}_1(\cdot | \mathcal{E}_0)}{d\tilde{\mathbb{P}}_1(\cdot | \mathcal{E}_0)} \Big|_{\mathcal{F}_t} \mathbf{1}_{\{\mathcal{E}_1 \cap \mathcal{E}_2\}} \Big| \mathcal{E}_0 \right] \\ &\geq \exp(-(1 - \lambda_\star) \log^{3/4} t - C') \times t^{-(1+o(1))(1-\varepsilon)} \times \tilde{\mathbb{P}}_1(\mathcal{E}_1 \cap \mathcal{E}_2 | \mathcal{E}_0) \\ &\geq \frac{\log \frac{a}{b}}{\log t} \times \exp(-(1 - \lambda_\star) \log^{3/4} t - C') \times t^{-(1+o(1))(1-\varepsilon)} = \omega\left(\frac{1}{t}\right), \end{aligned}$$

where in the second inequality we used (2.45). This concludes the proof of (2.40) as desired. \square

3. Conclusions and future work. In this paper we have shown that wrong information cascades can be *broken* by a small number of individuals who disregard the actions of others and only follow their private signal, leading to asymptotic learning of the correct action.

Moreover, we determined precisely the optimal asymptotic rate of decay of the error probability at a given time t .

This paper initiates a larger investigation into a broad family of problems that probe the fragility of information cascades. We collect here some natural directions for future work.

- *Alternative objectives.* While asymptotic learning is the primary objective, once this is achieved one can ask to minimize various notions of error. In this paper we focused on the optimal asymptotic rate at which the error probability E_t can go to zero, but other notions of error are natural to consider as well. For instance, what is the expected number of errors until time t and how can this be minimized? That is, what is

$$NE_t \equiv NE_t(a, b) := \inf_{\{p_i\}_{i=1}^t} \mathbb{E} \left[\sum_{i=1}^t \mathbf{1}_{\{Z_i \neq \theta\}} \right]$$

and what are the optimal revealing probabilities to achieve this? Theorem 1.1 implies that $NE_t \leq (1 + o(1))\kappa_\star \log t$. We conjecture that $NE_t = (1 + o(1))c \log t$ for some constant c such that $0 < c < \kappa_\star$.

- *Alternative behavioral models.* In this paper we considered a particular model: each player is either a Bayesian or follows their private signal blindly. More generally, one can consider any (causal) behavioral model that deviates from a pure Bayesian model. How do various notions of error depend on the model specifics?
- *More general setups.* In this paper we considered the simplest possible setup: two possible states of the world, a uniform prior over them, with each state corresponding to a distribution over two possible private signals, which are in a natural bijection with the possible states of the world. More generally one can ask the same questions with k possible states of the world, a general prior over them, and each possible state of the world corresponding to a distribution over ℓ possible private signals.
- *Unknown parameters.* We assumed that the players know the revealing probabilities $\{p_t\}_{t=1}^\infty$. What if these are unknown, or if the players believe that they are $\{p_t\}_{t=1}^\infty$ when they are actually $\{q_t\}_{t=1}^\infty$? Does asymptotic learning occur? If so, what is the optimal learning rate?

We suspect that, for many of these problems, the results and techniques of this paper will be useful in determining the correct order of magnitude for the relevant quantities. However, just like in this paper, determining the precise constants will require a deeper understanding of the specific problem of interest.

APPENDIX A: AUXILIARY PROOFS

A.1. Proof of (2.29). To prove (2.29), first note that the set A_t is either empty, in which case $R_t^{(1)} = 1$, or it can be decomposed into blocks of consecutive integers. That is, there exist integers $\ell_1, u_1, \ell_2, u_2, \dots, \ell_k, u_k$ (the lower and upper ends of the blocks) such that $\ell_i \leq u_i$ for every $i \in [k]$, $u_i + 1 < \ell_{i+1}$ for every $i \in [k - 1]$, and $A_t = \bigcup_{i=1}^k \{\ell_i, \ell_i + 1, \dots, u_i\}$. We can then write

$$(A.1) \quad R_t^{(1)} = \prod_{i=1}^k \prod_{j=\ell_i}^{u_i} \frac{R_j}{R_{j-1}} = \prod_{i=1}^k \frac{R_{u_i}}{R_{\ell_i-1}}.$$

Since $\ell_i \in A_t$ we have that $R_{\ell_i-1} > \frac{a}{b}$ for every $i \in [k]$, and since $u_i + 1 \notin A_t$ we have that $R_{u_i} \leq \frac{a}{b}$ for every $i \in [k]$; the condition $R_t \leq \frac{a}{b}$ is necessary to ensure that this latter fact holds for $i = k$ as well, that is, that $R_{u_k} \leq \frac{a}{b}$. Putting these together we obtain that each factor in (A.1) is at most 1, showing that (2.29) holds.

A.2. Proof of (2.30). We first rewrite the probability in question by taking logarithms:

$$\mathbb{P}_1(R_t^{(1)} \leq t^{-C_1}) = \mathbb{P}_1\left(\sum_{i \in A_t} \log \frac{R_i}{R_{i-1}} \leq -C_1 \log t\right).$$

Given $\theta = 1$ and $i \in A_t$, the random variable $\log \frac{R_i}{R_{i-1}}$ has the following distribution: it takes the value $\log \frac{b}{a}$ with probability $\frac{b}{a+b} p_i$ and it takes the value $\log \frac{1 - \frac{b}{a+b} p_i}{1 - \frac{a}{a+b} p_i}$ with probability $1 - \frac{b}{a+b} p_i$. Define $Y_i := (\log \frac{R_i}{R_{i-1}}) \wedge 0$. Thus, given $\theta = 1$ and $i \in A_t$, the random variable Y_i takes the value $\log \frac{b}{a}$ with probability $\frac{b}{a+b} p_i$ and it takes the value 0 otherwise. We then have that

$$\begin{aligned} \mathbb{P}_1\left(\sum_{i \in A_t} \log \frac{R_i}{R_{i-1}} \leq -C_1 \log t\right) &\leq \mathbb{P}_1\left(\sum_{i \in A_t} Y_i \leq -C_1 \log t\right) \\ &= \mathbb{P}_1\left(\sum_{i \in A_t} \frac{-Y_i}{\log \frac{a}{b}} \geq \frac{C_1}{\log \frac{a}{b}} \log t\right), \end{aligned}$$

and notice that $Y_i' := -Y_i / \log \frac{a}{b}$ is a Bernoulli random variable with expectation $\mathbb{E}Y_i' = \frac{b}{a+b} p_i$. Let $\{B_i\}_{i=1}^\infty$ be independent Bernoulli random variables with expectation $\mathbb{E}B_i = \frac{b}{a+b} p_i$. We then have the following bound:

$$\mathbb{P}_1\left(\sum_{i \in A_t} \frac{-Y_i}{\log \frac{a}{b}} \geq \frac{C_1}{\log \frac{a}{b}} \log t\right) \leq \mathbb{P}\left(\sum_{i=1}^t B_i \geq \frac{C_1}{\log \frac{a}{b}} \log t\right).$$

Note that $\mathbb{E} \sum_{i=1}^t B_i = \frac{a}{a+b} \sum_{i=1}^t p_i = (1 + o(1))C' \log t$ for some $C' = C'(a, b)$. Hence by the multiplicative Chernoff bound we obtain that

$$\mathbb{P}\left(\sum_{i=1}^t B_i \geq \frac{C_1}{\log \frac{a}{b}} \log t\right) \leq \exp(-2 \log t) = t^{-2}$$

if C_1 is large enough, concluding the proof.

A.3. Small revealing probabilities. Here we prove Lemma 2.1(b) in the case when $M_t := \sum_{i=1}^t p_i < \delta \log t$, where $\delta \equiv \delta(a, b) := \frac{1}{2 \log(1+a/b)}$. Note that the first two individuals follow their private signal, that is, $Z_1 = X_1$ and $Z_2 = X_2$. Hence if $X_1 = X_2 = 2$, then $R_2 = (b/a)^2$, and we are in the regime where the MAP estimator outputs 2. Thus if all the subsequent revealers have 2 as their private signal then the MAP estimator continues to output 2. Let $\text{Rev}_t := \{i \in [t] : I_i = 1\}$ denote the set of revealers until time t . Recalling (2.11) we thus obtain the following lower bound:

$$(A.2) \quad \mathbb{P}(\text{MAP}(Z_1, \dots, Z_{t-1}, X_t) \neq \theta) \geq \mathbb{P}_1(X_1 = X_2 = 2, X_i = 2 \forall i \in \text{Rev}_t).$$

Note that $|\text{Rev}_t|$ is independent of θ and that the Chernoff bound gives that

$$\mathbb{P}(|\text{Rev}_t| > 1.8\delta \log t) \leq 1/2.$$

Thus we have that

$$\begin{aligned} (A.3) \quad &\mathbb{P}_1(X_1 = X_2 = 2, X_i = 2 \forall i \in \text{Rev}_t) \\ &\geq \frac{1}{2} \mathbb{P}_1(X_1 = X_2 = 2, X_i = 2 \forall i \in \text{Rev}_t \mid |\text{Rev}_t| \leq 1.8\delta \log t) \\ &\geq \frac{1}{2} \left(\frac{b}{a+b}\right)^{2+1.8\delta \log t} = \frac{b^2}{2(a+b)^2} \times t^{-0.9}. \end{aligned}$$

Putting together (A.2) and (A.3) proves (2.5).

APPENDIX B: ON TIE-BREAKING RULES

A natural question is whether our results depend on the choice of the tie-breaking rule. Our main reason for choosing the tie-breaking rule that we chose is its simplicity (just like Le et al. [20] mention in their paper). It turns out that the choice of tie-breaking rule is immaterial: our result holds regardless of the tie-breaking rule—that is, the learning rate, together with the explicit constants, are unchanged—and essentially the same proof works to show this. We now briefly explain why this is the case.

First, note that the upper and lower bounds in (2.12), which bound the probability that the MAP estimator is incorrect, hold for any tie-breaking rule. These bounds are the starting point of the proof. They show that the analysis boils down to understanding the behavior of the L -likelihood ratio.

Turning now to the proof of the upper bound in Section 2.2, the following things are changed and unchanged for different tie-breaking rules. In (2.15) and (2.16), the middle cases depend on the tie-breaking rule, but the other cases are unchanged. If the values in these two middle cases are both strictly bounded away from 0 and 1, then the proof goes through in the same way, as we now explain. This means that (2.17) and (2.19) are unchanged, while (2.18) will change depending on the tie-breaking rule, but regardless will be a constant that is strictly less than 1. The same goes for the next three numbered displays: (2.20) and (2.22) are unchanged, while (2.21) will change depending on the tie-breaking rule, but regardless will be a constant that is strictly less than 1. Since what matters in the rest of the proof is (2.20) and (2.22) (with (2.21) bounded by (2.20)), everything else goes through unchanged.

So the only issue is if the two middle cases in (2.15) and (2.16) are not bounded away from 0 and 1. This means that (2.18) and (2.21) could potentially be 1. In this case we have to additionally argue that the L -likelihood ratio quickly gets out of the interval $[b/a, a/b]$ and hence (2.20) and (2.22) govern the behavior in this case as well.

In the proof of the lower bound in Section 2.3, the only thing that changes is the lower bound estimate for the probability $\mathbb{P}_1(\mathcal{E}_0)$. With a different tie-breaking rule we would not have that $Z_2 = X_2$. Nonetheless, for any tie-breaking rule we can show (possibly with changing t_0 to a different constant) that $\mathbb{P}_1(\mathcal{E}_0)$ is bounded from below by a constant. The rest of the proof goes through unchanged.

Acknowledgements. We thank Louigi Addario–Berry for helpful comments. We also thank two anonymous referees for their feedback.

REFERENCES

- [1] ACEMOGLU, D., DAHLEH, M. A., LOBEL, I. and OZDAGLAR, A. (2011). Bayesian learning in social networks. *Rev. Econ. Stud.* **78** 1201–1236. MR2896020 <https://doi.org/10.1093/restud/rdr004>
- [2] ACEMOGLU, D., MAKHDOUNI, A., MALEKIAN, A. and OZDAGLAR, A. (2019). Fast and slow learning from reviews. Working paper.
- [3] AGRAWAL, R. (1995). Sample mean based index policies with $O(\log n)$ regret for the multi-armed bandit problem. *Adv. in Appl. Probab.* **27** 1054–1078. MR1358906 <https://doi.org/10.2307/1427934>
- [4] ANDERSON, L. R. and HOLT, C. A. (1997). Information cascades in the laboratory. *Am. Econ. Rev.* **87** 847–862.
- [5] BANERJEE, A. V. (1992). A simple model of herd behavior. *Q. J. Econ.* **107** 797–817.
- [6] BERNARDO, A. E. and WELCH, I. (2001). On the evolution of overconfidence and entrepreneurs. *J. Econ. Manag. Strategy* **10** 301–330.
- [7] BIKHCHANDANI, S., HIRSHLEIFER, D. and WELCH, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Polit. Econ.* **100** 992–1026.
- [8] ÇELEN, B. and KARIV, S. (2004). Distinguishing informational cascades from herd behavior in the laboratory. *Am. Econ. Rev.* **94** 484–498.
- [9] CHENG, Y., HANN-CARUTHERS, W. and TAMUZ, O. (2017). A deterministic protocol for sequential asymptotic learning. Available at [arXiv:1802.06871](https://arxiv.org/abs/1802.06871).

- [10] COVER, T. M. (1969). Hypothesis testing with finite statistics. *Ann. Math. Stat.* **40** 828–835. MR0240906 <https://doi.org/10.1214/aoms/1177697590>
- [11] DRAKOPOULOS, K., OZDAGLAR, A. and TSITSIKLIS, J. N. (2013). On learning with finite memory. *IEEE Trans. Inf. Theory* **59** 6859–6872. MR3106869 <https://doi.org/10.1109/TIT.2013.2262037>
- [12] EASLEY, D. and KLEINBERG, J. (2010). *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge Univ. Press, Cambridge. MR2677125 <https://doi.org/10.1017/CBO9780511761942>
- [13] FREEDMAN, D. A. (1975). On tail probabilities for martingales. *Ann. Probab.* **3** 100–118. MR0380971 <https://doi.org/10.1214/aop/1176996452>
- [14] GALE, D. and KARIV, S. (2003). Bayesian learning in social networks. *Games Econom. Behav.* **45** 329–346. MR2023667 [https://doi.org/10.1016/S0899-8256\(03\)00144-1](https://doi.org/10.1016/S0899-8256(03)00144-1)
- [15] HANN-CARUTHERS, W., MARTYNOV, V. V. and TAMUZ, O. (2018). The speed of sequential asymptotic learning. *J. Econom. Theory* **173** 383–409. MR3737265 <https://doi.org/10.1016/j.jet.2017.11.009>
- [16] HAREL, M., MOSSEL, E., STRACK, P. and TAMUZ, O. (2017). Groupthink and the failure of information aggregation in large groups. Preprint available at <https://arxiv.org/abs/1412.7172>.
- [17] HELLMAN, M. E. and COVER, T. M. (1970). Learning with finite memory. *Ann. Math. Stat.* **41** 765–782. MR0272100 <https://doi.org/10.1214/aoms/1177696958>
- [18] KAHNEMAN, D. and TVERSKY, A. (1973). On the psychology of prediction. *Psychol. Rev.* **80** 237–251.
- [19] LAI, T. L. and ROBBINS, H. (1985). Asymptotically efficient adaptive allocation rules. *Adv. in Appl. Math.* **6** 4–22. MR0776826 [https://doi.org/10.1016/0196-8858\(85\)90002-8](https://doi.org/10.1016/0196-8858(85)90002-8)
- [20] LE, T. N., SUBRAMANIAN, V. G. and BERRY, R. A. (2017). Information cascades with noise. *IEEE Trans. Signal Inf. Process. Netw.* **3** 239–251. MR3661698 <https://doi.org/10.1109/TSIPN.2017.2682798>
- [21] MOSSEL, E., SLY, A. and TAMUZ, O. (2012). On agreement and learning. Preprint available at <https://arxiv.org/abs/1207.5895>.
- [22] MOSSEL, E., SLY, A. and TAMUZ, O. (2014). Asymptotic learning on Bayesian social networks. *Probab. Theory Related Fields* **158** 127–157. MR3152782 <https://doi.org/10.1007/s00440-013-0479-y>
- [23] MOSSEL, E., SLY, A. and TAMUZ, O. (2015). Strategic learning and the topology of social networks. *Econometrica* **83** 1755–1794. MR3414192 <https://doi.org/10.3982/ECTA12058>
- [24] SGROI, D. (2002). Optimizing information in the herd: Guinea pigs, profits, and welfare. *Games Econom. Behav.* **39** 137–166. MR1894952 <https://doi.org/10.1006/game.2001.0881>
- [25] SMITH, L. and SØRENSEN, P. (2000). Pathological outcomes of observational learning. *Econometrica* **68** 371–398. MR1748010 <https://doi.org/10.1111/1468-0262.00113>
- [26] TAY, W. P., TSITSIKLIS, J. N. and WIN, M. Z. (2008). On the subexponential decay of detection error probabilities in long tandems. *IEEE Trans. Inf. Theory* **54** 4767–4771. MR2591139 <https://doi.org/10.1109/TIT.2008.929032>
- [27] WANG, Y. and DJURIĆ, P. M. (2015). Social learning with Bayesian agents and random decision making. *IEEE Trans. Signal Process.* **63** 3241–3250. MR3350019 <https://doi.org/10.1109/TSP.2015.2421486>