On the efficiency of social learning

Dinah Rosenberg*and Nicolas Vieille†

November 3, 2017

Abstract

We revisit well-known models of learning in which a sequence of agents make a binary decision on the basis of a private signal and additional information. We introduce efficiency measures, aimed at capturing the speed of learning in such contexts. Whatever the distribution of private signals, we show that the learning efficiency is the same, whether each agent observes the entire sequence of earlier decisions, or only the previous decision. We provide a simple necessary and sufficient condition on the signal distributions under which learning is efficient. This condition fails to hold in many prominent cases of interest. Extensions are discussed.

When making partially informed decisions, agents routinely rely on both their private information, and on the observed behavior of other agents. Since the past decisions of these other agents partially reflect their information, privately held information may spread within the economy or society, and agents acting later make better informed decisions.

The stylized models developed in the literature on observational and social learning offer influential setups to formalize and analyze this insight. In the paradigmatic setup introduced in Bikhandani, Hirshleifer and Welch (1992) and Banerjee (1992), each agent in an infinite sequence tries to guess an unknown, binary state of nature. Agents are endowed with private signals, which are conditionally *i.i.d.* given the underlying state,

*HEC, 1, rue de la Libération, 78351 Jouy en Josas, France, rosenberg@hec.fr

†HEC, 1, rue de la Libération, 78351 Jouy en Josas, France, vieille@hec.fr.

and in addition observe the entire sequence of earlier guesses. If private signals offer a binary and noisy version of the state, a cascade occurs: eventually, all agents find it optimal to ignore their private information, and to repeat the previous guess. With positive probability, the limit guess is incorrect and learning then fails.

Following these seminal contributions, much attention has been paid to the situation where the precision of private signals is unbounded, as formalized in Smith and Sørensen (2000): no matter how strong the existing evidence in favor of one state is, there is a positive probability that the signal of the current agent will overturn this evidence. In such a case, the long-term outcomes of learning are as good as in the ideal case where all signals are public: eventually, all agents make the correct guess., w.p. 1. The driving force behind this *asymptotic learning* result is that, whenever a sequence of incorrect guesses is developing, and even though the belief assigned to the incorrect state is reinforced with time, some agent will eventually put an end to it.

The unbounded precision assumption facilitates asymptotic learning even when agents do not observe the entire sequence of earlier guesses, but only a sample. Assuming that each agent observes only the action of the previous agent, Celen and Kariv (2004) have provided one example where the sequence of guesses does not settle on the correct guess, yet the probability that the n-th agent makes a wrong guess converges to zero, as $n \to +\infty$.[1] This positive result was proven by Acemoglu, Dahleh, Lobel and Ozdaglar (2011) to hold much more generally, under the assumption that samples are independent across agents, and that each agent is aware of the "identity" of the agents he is sampling.[2] When one agent is sampled, whose identity is *not* disclosed, Smith and Sørensen (1996, 2013) have shown that successive guesses converge to the correct guess.[3] This is true even when agents are not even aware of their own rank in the sequence, as shown in Monzon and Rapp (2014).

This set of elegant theoretical results leaves open the issue of how quickly learning takes place. Our goal here is to refine such results by assessing the *efficiency* of learning, as a function of the distribution of private signals and of the nature of the informational feedback that is available. This impact is *a priori* unclear. For given signal distributions,

[1] Thus, the sequence of guesses converges to the correct one in probability, but not almost surely.

[2] Under some weak restriction on the sampling distributions. When instead samples are correlated across agents, alternative measures of information aggregation, coined information *diffusion*, have been developed in Lobel and Stadler (2015).

[3] Again in probability, but also almost surely in Cesaro mean, yet not almost surely.

and if scarce feedback is available, sequences of identical guesses carry less information, and are more easily overturned: intuitively, strings of consecutive identical guesses tend to be shorter than when more feedback is provided. In terms of efficiency, this is good news if these guesses are incorrect, but not otherwise. On the other hand, for a given feedback, and if more precise signals are made more likely, then the chances that the *first* agent will make a correct guess are of course increased. But this first guess also carries more information. Should it be incorrect, it will take a very precise signal for agent 2 to act against this evidence. Whether the second agent (and later ones as well) benefits from the increased precision is ambiguous.

We consider two measures of efficiency. One criterion looks at how long one needs to wait until *one* agent makes a correct guess. We say that learning is *efficient* (in state θ) if the expectation (given θ) of the rank of the *first* agent making a correct guess is finite. We view this measure as a simple and natural proxy for the likelihood of long, incorrect herds. The other, closely related criterion, measures whether the expected number of incorrect guesses is finite or not. The latter criterion is sometimes more amenable to a formal analysis.

We mostly focus on two polar setups, and assume either that each agent observes the entire sequence of earlier guesses, or that each agent observes only the previous one. Our main result is two-fold. First, and possibly surprisingly, whether learning is efficient is *independent* of the setup: for every signal distribution, learning is efficient in one setup if and only if it is efficient in the other one. Second, whether learning is efficient can be readily ascertained. Denote by H and L the two states. Denote by F the cdf of the belief an agent assigns to the state being H, on the basis of his private signal. The function F is derived from the prior belief over the state, and from the state-contingent distribution of signals. With these notations, learning is efficient in state L if and only if the indefinite integral $\int_0^1 \frac{1}{F(p)} dp$ is finite.[4] Thus, not only should very informative signals have a positive probability ($F(p) > 0$ for all $p > 0$), but such signals should be relatively likely. Just how likely is given by the integral criterion.

Such a simple, necessary and sufficient condition for efficiency has obvious benefits. On the one hand, it allows to derive some comparative statics insights. For example, as is easily checked, improving the precision of signals – in the sense of a second-order

[4]Plainly, a dual version holds when exchanging the two states.

stochastic dominance of the distribution of posterior beliefs –, benefits efficiency.[5] It also enables to assess efficiency for specific distributions of signals. Learning is inefficient in the textbook case where beliefs are uniformly distributed ($F(p) = p$ for all p). Learning is also inefficient when signals are normally distributed, with a state-dependent distribution, or when signals follow Levy distributions. When signals follow Pareto distributions, or normal distributions with a state-dependent variance, efficiency only holds if the parameters in the two states are sufficiently far apart. Thus, our rather conservative efficiency criterion is rarely met. This stands in sharp contrast with the benchmark case where all signals would be publicly observed, in which case learning is always efficient, no matter F.

These conclusions raise further questions. We address three of them. To start with, one may still hope to rank the two setups in the inefficient case, based on how quickly the expected number of incorrect guesses among the first n agents increases to infinity. A general detailed study is beyond our scope. We limit ourselves to one example, where the performance of the two scenarios under this refined approach is quite similar.

Next, assume a designer who chooses the setup, with the goal of promoting efficiency. Does there exist alternative setups with better efficiency properties? Our (partial) results are negative. First, we prove that no general irrelevance result holds, that is, there are setups with unambiguously *poorer* efficiency properties. When each agent randomly samples one of the earlier guesses,[6] learning is inefficient, no matter F. Yet, whether there are settings with *better* efficiency properties is unclear. A natural extension is to introduce observational delay, that is, to assume that all decisions become eventually public, but possibly not immediately. For instance, let agents be grouped in successive generations, of possibly varying size. Agents in a given generation observe all previous generations, but not agents within the same generation. When posterior beliefs are uniform, we show that learning is inefficient, whenever the sizes of the successive generations do not increase at a super-exponential rate.[7]

Assume finally a social planner with the ability to dictate individual decision rules (subject to the informational constraints of the game), and with the goal of promoting

[5] As was stressed, it is obvious that improving the precision of *one* individual signal benefits this agent. That improving the precision of *all* signals also has positive effects is *a priori* unclear.

[6] Without being told who made it.

[7] As we argue, the case of uniform posteriors is a limiting one, in which we should expect efficiency to obtain more easily, and is therefore a natural choice for such an exercise.

efficiency. We are not able to solve for the first-best optimum of such a social planner. Yet, we prove that there exist individual strategies such that learning is efficient, no matter F. This leads to a better appreciation of the sources of inefficiency. Indeed, Bayesian learning differs in two obvious respects from the benchmark case where all signals are public (in which learning takes place exponentially quickly). First, rich private signals are encoded into a binary guess, so that agents get only a coarse version of the existing information. Next, this encoding is endogenous, and the result of equilibrium play. Arieli and Mueller-Frank (2016) have shown that, when there is a continuum of available actions, there are cases where the encoding of private signal into an action is one-to-one, so that private signals can be inferred from observed actions. The sequence of guesses is then the same as if all signals were publicly observed, restoring efficiency. Together with our observation, this tentatively suggests that inefficient learning requires the combination of a coarse action set, and of a strategic behavior.

Much of the literature has focused on the long-term outcomes of learning. In comparison, existing work on the efficiency of learning is limited. Chamley (2004) reports some numerical evidence on the speed of learning, obtained by Monte-Carlo simulation. For specific belief distributions, some formal bounds on the probability e_n that agent n makes an incorrect guess have been established. None of these bounds implies, nor is implied by, our results. Under the assumption that beliefs are uniformly distributed, Lobel, Acemoglu, Dahleh and Ozdaglar (2009) and Monzon and Rapp (2014) provide upper bounds on e_n, respectively for the case where one earlier agent is randomly sampled and observed, and for the case where the previous agent is observed. For an example close to the one in Monzon and Rapp (2014), Celen and Kariv (2004) show that learning is inefficient.[8] Under the assumption that signals are normally distributed, and all previous guesses observed, Hann-Caruthers, Martynov, and Tamuz (2017) provide a lower bound on e_n.

Our paper assumes, as all the literature reviewed above, that agents are perfect Bayesian, and that this is common knowledge. That is, when forming posterior beliefs, agents interpret the available evidence using Bayes rule and taking for granted that other agents are acting optimally. There is also a large literature where agents use instead some heuristics to update beliefs, see *e.g.* De Marzo, Vayanos and Zwiebel (2003), Golub and Jackson (2010) or Acemoglu and Ozdaglar (2010) for a survey.

[8]This is a reinterpretation in our own terms of their result, which does not fit our framework.

Our paper also relates to the literature on dynamic games with pure informational externalities (such as strategic experimentation games). For such games, Rosenberg, Solan and Vieille (2009) prove that at any equilibrium, a "consensus" eventually emerges about what is the optimal action. Mossel, Sly and Tamuz (2014, 2015) show that this consensus is "correct" in many important cases. Harel, Mossel, Strack and Tamuz (2017) provide results on how quickly such a consensus emerges.

The paper is organized as follows. Section 1 contains the model and some introductory material. Section 2 contains our main results, together with the benchmark case. It also includes an informal discussion of the proofs, and provides consequences of our results. Section 3 contains the additional results discussed above.

1 The Model

1.1 Framework and first definitions

We describe the general framework in which our results will be cast. Nature chooses a binary state $\theta \in \{L, H\}$. Both states are assumed equally likely (this is only for notational simplicity). A sequence of short-lived agents make decisions in turn. In round n, agent n chooses a binary guess $a_n \in \{l, h\}$, with the objective of matching the underlying unknown state θ. To be specific, agent n forms a posterior belief based on the available information and picks the guess a_n which matches the most likely state. The information available to agent n comes from two sources. First, agent n gets a private signal s_n. In addition, some information about guesses made by earlier agents is provided.

Successive private signals are assumed to be conditionally *iid* given θ. Obviously, the distribution of signals matters only through the distribution of beliefs it induces. We denote by $q_n := \mathbf{P}(\theta = H \mid s_n)$ the posterior belief that agent n assigns to the state being H, based on his signal only. We refer to q_n as the *private belief* of agent n. We denote by F the cdf of the distribution of q_n, and by F_θ the conditional cdf of q_n given θ.

We assume throughout that q_n has a continuous distribution with continuous density f. This continuity assumption ensures that signals do convey information about θ. But its main role is to ensure that no two signals have the same informational content. Under

this assumption, there is a zero probability that some agent is indifferent between the two guesses. Consequently, the distribution of the stochastic sequence (a_n) of guesses is uniquely defined.

We focus on the case where the information on the state conveyed by a private signal is potentially arbitrarily precise. Formally, following Smith and Sørensen (2000), we assume that $0 < F(p) < 1$ for all $p \in (0,1)$.[9] Thus, no matter how strong the available evidence is, there is a strictly positive probability that the signal s_n will overturn it.

In addition, we maintain the assumption that there exists $C > 0$ and $\alpha > 0$ such that the inequality $F(p) \leq Cp^\alpha$ holds for all p small enough, and maintain a symmetric one on $\tilde{F}$. Since α is unrestricted, we view this condition as a purely technical restriction.[10]

The *interim belief* π_n is the probability attached to $\theta = H$, given all the information available to agent n, excluding s_n. For instance, when all actions are public, $\pi_n = \mathbf{P}(\theta = H \mid a_1, \ldots, a_{n-1})$. The belief π_n is computed under the (equilibrium) assumption that earlier agents acted optimally.

Agent n uses his interim belief π_n and his private belief q_n to form his *posterior belief* p_n that $\theta = H$. By Bayes rule, and since both states are *ex ante* initially likely, p_n is given by the likelihood ratio equation

$$\frac{p_n}{1-p_n} = \frac{\pi_n}{1-\pi_n} \times \frac{q_n}{1-q_n}.$$

Agent n guesses $a_n = h$ whenever $p_n \geq \frac{1}{2}$, that is $\dfrac{p_n}{1-p_n} \geq 1$, or equivalently,

$$\pi_n + q_n \geq 1. \tag{1}$$

Since the private belief is a continuous random variable, all agents have a strict preference for one action over the other, almost surely.

We conclude with a clarifying comment. While arbitrary distributions of signals can be conceived of, the induced distribution of private beliefs is constrained by Bayes

[9] When this assumption fails to hold, there is a positive probability that agents eventually settle on the incorrect guess. Thus, learning fails to be efficient.

[10] We know of no "standard" distributions of signals, for which the induced distribution of beliefs fails this condition. Distributions F that fail this condition but exhibit a regular behavior around zero can be accommodated. Hence, this condition is only needed to rule out pathological situations, for which our proofs would need serious amendments.

updating. Indeed, Bayes updating entails that the expected posterior belief is equal to the prior belief $\frac{1}{2}$, which writes here $\int_0^1 pf(p)dp = \frac{1}{2}$. Conversely, and given any such density function $f(\cdot)$, there exist state-contingent signal distributions such that the induced distribution of posterior beliefs has density f.[11] The conditional density f_θ of the posterior belief given θ, is uniquely determined by f, and given by

$$f_H(p) = 2pf(p) \text{ and } f_L(p) = 2(1-p)f(p). \tag{2}$$

Plainly, while f_θ is uniquely defined, there are many signal distributions which are consistent with f.

Equation (2) implies the following relations on cumulative distribution functions:

$$F_H(p) = 2\left(pF(p) - \int_0^p F(x)dx\right) \text{ and } F_L(p) = 2(1-p)F(p) + 2\int_0^p F(x)dx. \tag{3}$$

It is not difficult to deduce from (3) that $F_H(p) < F_L(p)$ for all $p < 1$, with $\lim_{p\to 0} \frac{F_H(p)}{F_L(p)} = 0$: posterior beliefs tend to be higher in state H than in state L, and posterior beliefs close to zero are relatively very unlikely in state H.

A leading example in the literature is that of *uniform posteriors*, in which $F(p) = p$ for all $p \in [0,1]$.[12]

1.2 Efficiency notions

We introduce the two notions of efficiency that we use. We denote by $\tau := \inf\{n \geq n : a_n = \theta\}$ the first agent whose guess is correct, and by $N := \#\{n \geq 1, a_n \neq \theta\}$ the total number of incorrect guesses.

According to our first notion, learning is *efficient in state* $\theta \in \{L, H\}$ if $\mathbf{E}_\theta[\tau] < +\infty$ and inefficient otherwise. According to our second notion, learning is efficient in state θ if $\mathbf{E}_\theta[N] < +\infty$, and inefficient otherwise. Obviously, the total number N of incorrect guesses is at least $\tau - 1$, hence the latter notion is at least as demanding as the former.

[11] This is by now a standard statement, first established in a finite setup by Aumann and Maschler (1995).

[12] Such a distribution of private beliefs arises *e.g.* when the conditional densities of the private signal are $f_L(s) = 2(1-s)$ and $f_H(s) = 2s$ for all $s \in [0,1]$, in which case the private belief q coincides with the private signal, and its unconditional density is $f(s) = \frac{1}{2}(f_H(q) + f_L(s)) = 1$.

The reason why we use one single terminology is as follows. When all guesses are public, the efficiency of learning is measured by how quickly an initial sequence of incorrect guesses is ended. Since $\tau < +\infty$ almost surely,[13] a natural proxy is whether the expectation of τ is finite or not. This is a weak notion of efficiency. Put it differently, an infinite expectation $E_\theta[\tau] = +\infty$ is a strong indication that, should the first agents coordinate on a wrong guess, putting an end to this sequence of incorrect guesses will require an unreasonable amount of time. When instead $\mathbf{E}_\theta[\tau] < +\infty$, we will show that the seemingly stronger result $\mathbf{E}_\theta[N] < +\infty$ holds as well. Thus, the two notions are equivalent.

When only the previous guess is observed, assessing whether $\mathbf{E}_\theta[\tau]$ is finite or not seems neither easy to answer, nor a reasonable measure of efficiency. Indeed, it is reasonable to expect that the sequence of guesses exhibits less persistency than when all guesses are public. That is, it is likely that $\mathbf{E}_\theta[\tau]$ will in most cases be finite, yet we believe it is a poor indication of efficiency. Accordingly, we will focus on the second notion.

We make no claim that these are the only reasonable measures of efficiency. Ultimately, an important rationale for ours is that they allow for general and clear-cut insights.

2 Main results

2.1 The benchmark case

As a benchmark, we use the ideal case where all signals are publicly observed. It is well-known that the rate of convergence of beliefs is exponential. For completeness, we prove an *ad hoc* statement.

Denote by $X := \ln \dfrac{q}{1-q}$ the log-likelihood ratio of the private belief, and introduce the Laplace transform $\phi_X(t) := \mathbf{E}_L\left[e^{tX}\right]$ of X. The map ϕ_X has non-negative but possibly infinite values, and its Cramer transform is $h_X(a) := \inf_{t\geq 0}\left(at - \ln \phi_X(t)\right)$. Using Chernoff bound, one has $\mathbf{P}_L(a_n = h) \leq e^{-nh_X(0)}$. As we show in the appendix,

[13] This follows from the stronger property that the total number N of wrong guesses is finite, almost surely.

$h_X(0) > 0$. Since $\mathbf{E}_L[N] = \sum_{n=1}^{+\infty} \mathbf{P}_L(a_n = h)$, this readily implies that learning is efficient.

Theorem 1 *Assume that signals are publicly observed. Then*

$$\mathbf{E}_L[N] < \frac{1}{e^{h_X(0)} - 1} < +\infty.$$

The conclusion in Theorem 1 holds much more generally as soon as private signals are informative, that is, whenever $0 < F(p) < 1$ for *some* $p \in (0, 1)$.

Theorem 1 allows to derive upper bounds on $\mathbf{E}_L[N]$ easily. For instance, with uniform private beliefs $F(q) = q$, the expected number of wrong guesses is strikingly low.

Corollary 2 *Assume $F(q) = q$ for all $q \in [0, 1]$. Then $\mathbf{E}_L[N] \leq \frac{\pi}{4 - \pi} \leq 3.66$.*

2.2 Observational learning

We now assume that all guesses are publicly observed, so that the information available to agent n consists of the ordered sequence $a_1, \ldots, a_{n-1}$ of earlier guesses, and of s_n. In that case, the interim belief π_n of agent n is given by $\pi_n = \mathbf{P}(\theta = H \mid a_1, \ldots, a_{n-1})$.

As Smith and Sørensen (2000) have shown, $\pi_n \to 0$, $\mathbf{P}_L$-almost surely. The proof of Smith and Sørensen relies on martingale theory: given $\theta = L$, the likelihood ratio $l_n := \frac{\pi_n}{1 - \pi_n}$ follows a positive martingale. It therefore converges, $\mathbf{P}_L$-a.s., to a variable $l_\infty \geq 0$ such that $\mathbf{E}_L[l_\infty] < +\infty$. Because the informativeness of private signals is unbounded, $l_\infty \in \{0, +\infty\}$, $\mathbf{P}_L$-a.s. Further, $\mathbf{E}_L[l_\infty] < +\infty$ implies $l_\infty < +\infty$, $\mathbf{P}_L$-a.s. Therefore, $l_\infty = 0$.

Our results below require an entirely different approach. Our first result identifies a condition on F under which the expected time of the first correct guess is finite.

Theorem 3 *Assume that actions are publicly observed. Then*

$$\mathbf{E}_L[\tau] < +\infty \Leftrightarrow \int_0^1 \frac{1}{F(p)} dp < +\infty.$$

This is a statement on efficiency in state L. There is of course a symmetric result on the efficiency in state H, obtained when exchanging the roles of the two states. Theorem 3 as stated only requires that $F(p) > 0$ for all $p > 0$. The dual version requires $F(p) < 1$ for all $p < 1$.

According to Theorem 3, efficient learning not only requires that extremely strong signals have positive probability, but also be relatively likely. A necessary (but not sufficient) condition for efficiency is that the density f of the private belief be unbounded. The exact form of the integral is somewhat difficult to interpret. But having a necessary and sufficient condition has obvious benefits.

Our second result is that, as announced, the two efficiency notions coincide under the following weak technical condition (**A1**): $\lim_{p \to 0} \frac{f(p)}{p^\beta} = c$, for some $c > 0$ and $\beta \in \mathbf{R}$, and a symmetric statement holds for $p \to 1$.[14]

Theorem 4 *Assume that actions are publicly observed. Then*

- $\mathbf{E}_L[\tau] < +\infty \Rightarrow \mathbf{E}_L[N] < +\infty$.
- *If (**A1**) holds, then* $\mathbf{E}_L[N] < +\infty \Rightarrow \mathbf{E}_L[\tau] < +\infty$.

We provide some insights into the proofs of Theorems 3 and 4 in Section 2.4. The full proofs are in the appendix and supplementary appendix respectively.

2.3 Social learning

We here assume that each agent n observes only a_{n-1} in addition to s_n. As Acemoglu et *al.* (2011) have shown, the probability $\mathbf{P}_L(a_n = h)$ that agent n makes a wrong guess converges to zero. Their proof makes uses of an improvement principle first used by Banerjee and Fudenberg (2004): since each agent has the option of copying the previous guess, the expected equilibrium payoff of agent n is increasing in n, that is, the probability of a wrong guess decreases with n. A delicate issue is to quantify this improvement, so as to ensure that the error probabilities decrease towards zero.

Efficiency in this setup will be shown to be equivalent to efficiency in the setup of the previous section. However, we have no direct proof of equivalence. Instead, we again characterize efficiency by means of a necessary and sufficient condition on F, see Theorem 5 below, then check that this condition is equivalent to the condition in Theorem 3, see Proposition 3.

[14] Consider two different priors $\pi_1 < \pi_1'$ on the state, and the associated public beliefs π_2 and π_2' held by the second agent, in the event where $a_1 = h$. The role of **(A1)** is to ensure that $\pi_2 < \pi_2'$, at least if π_1 is close to one. At the cost of notational complexity, **(A1)** could be significantly relaxed. We suspect that it can altogether be dispensed with.

For convenience, we will assume that the distribution of the private belief follows a (common) symmetry condition (**A2**), which can be stated as $f_H(p) = f_L(1-p)$. This condition is equivalent to $F(p) + F(1-p) = 1$ for all $p \in [0,1]$.

Theorem 5 *Assume that each agent observes the previous action and that (**A2**) holds. Then*

$$\mathbf{E}_L[N] < +\infty \Leftrightarrow \int_0^1 \frac{p}{\int_0^p F(x)dx} dp < +\infty.$$

The proof follows the same logic as that of Theorem 3, though the computations differ much.

Proposition 1 *One has*

$$\int_0^1 \frac{1}{F(p)} dp < +\infty \Leftrightarrow \int_0^1 \frac{p}{\int_0^p F(x)dx} dp < +\infty.$$

Proof. Note that $\int_0^p F(x)dx \leq pF(p)$. This implies $\int_0^1 \frac{p}{\int_0^p F(x)dx} dp \geq \int_0^1 \frac{1}{F(p)} dp$, hence the reverse implication holds. For the direct implication, we note that $\int_0^p F(x)dx \geq \frac{p}{2} F\left(\frac{p}{2}\right)$, hence

$$\int_0^1 \frac{p}{\int_0^p F(x)dx} dp \leq 2 \int_0^1 \frac{1}{F\left(\frac{p}{2}\right)} dp = 4 \int_0^{\frac{1}{2}} \frac{1}{F(p)} dp.$$

■

Having the two integral formulations is useful. The former version is often easier to use with specific distributions of signals, and will be relied upon in Section 2.5. The latter version allows us to establish that increasing the precision of signals has a positive impact on efficiency.

Corollary 6 *Consider two distributions of beliefs, with cdf F and G respectively. Assume that F stochastically dominates G in the second-order sense. If learning is efficient when the distribution of beliefs is given by F, then it is also efficient when the distribution of beliefs is given by G.*

Because F stochastically dominates G in the second-order sense, G can be obtained from F through a sequence of mean-preserving spreads. A mean-preserving spread in the distribution of beliefs corresponds to an increased precision of signals.

Proof. Since learning is efficient with F, one has $\displaystyle\int_0^1 \frac{p}{\int_0^p F(x)dx} dp < +\infty$. Since F stochastically dominates G, $\displaystyle\int_0^p F(x)dx \leq \int_0^p G(x)dx$ for all $p \in [0,1]$. In turn, this implies $\displaystyle\int_0^1 \frac{p}{\int_0^p G(x)dx} dp < +\infty$, hence learning is efficient with G. ■

2.4 Heuristics

We start with Theorem 3. The proof requires neither conceptual insights nor sophisticated tools. It instead relies on a careful analysis of Bayesian updating. We here provide some insights. Mathematical details are in the appendix.

We have seen (see (1)) that agent n guesses $a_n = h$ iff his private belief $q_n = \mathbf{P}(\theta = H \mid s_n)$ and the public belief $\pi_n = \mathbf{P}(\theta = H \mid a_1, \ldots, a_{n-1})$ are such that $q_n + \pi_n \geq 1$. Conditional on the sequence $a_1, \ldots, a_{n-1}$ of earlier guesses, this occurs with probability $1 - F_\theta(1 - \pi_n)$ in state θ. In that event, the public belief π_{n+1} in the next round is given by Bayes rule:

$$\frac{\pi_{n+1}}{1 - \pi_{n+1}} = \frac{\pi_n}{1 - \pi_n} \times \frac{1 - F_H(1 - \pi_n)}{1 - F_L(1 - \pi_n)}. \tag{4}$$

We denote by (π_n^*) the sequence which is defined recursively by (4), with $\pi_1^* = \frac{1}{2}$. Thus, the sequence (π_n^*) describes the evolution of the public belief as long all agents keep guessing h.

With these notations, the probability $u_n := \mathbf{P}_L(\tau > n)$ that the first correct guess occurs no sooner than in round $n+1$, is given by

$$u_n = \mathbf{P}_L(\tau > n) = \mathbf{P}_L(a_1 = \cdots = a_n = h) = \prod_{k=1}^{n} (1 - F_L(1 - \pi_k^*)). \tag{5}$$

Since $\mathbf{E}_L(\tau) = \displaystyle\sum_{n=0}^{+\infty} u_n$, the question is to determine whether the series $\sum u_n$ converges, when (u_n) is given by (5) and (π_n^*) by (4).

Since $F_H(q) < F_L(q)$ for all $q < 1$, the sequence (π_n^*) converges to one. This implies that the ratio $\dfrac{1 - F_H(1 - \pi_n)}{1 - F_L(1 - \pi_n)}$ converges to one as well, hence $\displaystyle\lim_{n \to +\infty} \frac{\pi_{n+1}}{\pi_n} = 1$. The intuition is straightforward. Prior to τ, the belief assigned to H is progressively reinforced.

Consequently, the probability that agent n switches to l, that is, $\mathbf{P}(\tau = n \mid \tau \geq n)$, is decreasing with time. Because agents are more widely expected to conform to earlier guesses, the reinforcement in the belief assigned to H becomes more limited.

The proof is divided in two parts. The first part contains the bulk of the analysis. We introduce the log-likelihood ratio $r_n := \ln \dfrac{\pi_n}{1-\pi_n}$ of the public belief, and prove that the two series $\sum u_n$ and $\sum e^{-r_n}$ have the same nature. This follows along the following lines. Using (4), one has

$$u_n = e^{-r_{n+1}} \prod_{k=1}^{n} \left(1 - F_H(1-\pi_k)\right). \tag{6}$$

Thus $u_n \leq e^{-r_{n+1}}$ and the convergence of $\sum e^{-r_n}$ implies the convergence of the series $\sum u_n$. The proof of the reverse implication is more subtle. If $\sum_n u_n < +\infty$, then $1-F_L(1-\pi_n^*)$ converges to zero rather slowly. Since $\lim_{\pi\to 1} \frac{F_H(1-\pi)}{F_L(1-\pi)} = 0$, this nevertheless implies that $F_H(1-\pi_n)$ converges to zero sufficiently quickly, so that the infinite product $\prod_{k=1}^{+\infty} (1 - F_H(1-\pi_k))$ is non-zero. Using (6), this establishes the reverse implication.

On the other hand, a Taylor expansion in the updating formula (4) shows that $r_{n+1} - r_n$ is approximately equal to $2F(e^{-r_n})$ as $n \to +\infty$. Building on these observations, learning efficiency in state L is equivalent to $\sum_u e^{-r_n} < +\infty$, where (r_n) is recursively defined by $r_{n+1} - r_n = 2F(e^{-r_n})$.

We conclude the analysis in a second step. The specific form of the latter problem, a summation criterion subject to a difference equation, lends itself to a reformulation in continuous time. We indeed prove that learning efficiency in state L is equivalent to $\int_0^{+\infty} e^{-r(t)} dt < +\infty$, where $r(\cdot)$ is a solution to the differential equation $r'(t) = 2F(e^{-r(t)})$.[15] A simple change of variable then yields the result.

This is reminiscent of stochastic approximation techniques, as surveyed in Benaim (1999) or Pemantle (2007). However, these techniques aim at understanding the asymptotics of (r_n), and not at providing estimates on the entire sequence (r_n). Here instead, the limiting behavior of (r_n) is clear. The existing literature does not seem to apply. This approach also reminds of Euler schemes in analysis. But we here use a continuous-time formulation to approximate a discrete-time one, rather than the converse.

[15] Plainly, this is trickier than simply comparing a series with an integral, as the sequence (r_n) is defined recursively.

We next provide some intuition behind Theorem 4. The first implication is obvious since $N \geq \tau - 1$. Assume now that $\mathbf{E}_L[\tau] < +\infty$. In round τ, the current agent makes the guess l, and starts a herd of *correct* guesses. Unlike herds of incorrect guesses, a correct herd lasts forever with positive probability. However, with positive probability, this first correct herd will stop in finite time and some agent will switch back to h, thus starting a second incorrect herd. Since $\mathbf{E}_L[\tau] < +\infty$, the expected duration of this second herd is finite as well. Once the second incorrect herd ends, there is a probability *bounded away from zero* that the second correct herd will last forever, *etc.* Viewing the sequence of guesses as a sequence of herds of random duration, alternating between correct l-herds and incorrect h-herds, the expected number of h-herds is thus finite, and the expected duration of each h-herd is finite as well. This however does not quite imply that the expectation of N is finite, as the (random) number of h-herds is unbounded. To cut through the difficulty, we add the weak technical requirement (**A1**), which serves to ensure that the conditional expected duration of each h-herd is uniformly bounded.

2.5 Applications

We here report some results on the efficiency of learning for specific distributions of signals. Observe first that for uniform priors, $F(p) = p$, one has $\int_0^1 \frac{1}{F(p)} dp = +\infty$, hence learning is inefficient.[16] More generally, if $F(p)/p^\beta$ has a positive limit at 0 for some $\beta > 0$, then learning is efficient iff $\beta < 1$. Thus, the case of uniform priors appears as a limiting one, separating the efficient case from the inefficient one.

To apply Theorems 3 and 5 with given distribution of signals, one needs to get estimates on the cdf F of beliefs. The relation between a signal and the induced belief is not always straightforward. In particular, extreme signals do not necessarily turn into extreme beliefs, nor do fat tails in the signal distributions necessarily imply that learning is efficient. What matters is how likely informative signals are, and this is the result of an interplay between the two state-contingent distributions.

We start with the prominent case where private signals are normally distributed: $s_n \sim \mathcal{N}(\mu_\theta, \sigma_\theta^2)$, conditional on θ. We assume first different means, but identical variances.

Proposition 2 *Assume $\sigma_H^2 = \sigma_L^2 = \sigma^2$, and $\mu_H \neq \mu_L$. Then learning is inefficient.*

[16] A result equivalent to this one appears in Celen and Kariv (2004) in the case where only the previous action is observed.

Proof. We assume w.l.o.g. that $\Delta\mu := \mu_H - \mu_L > 0$. We will prove that for each $k \in \mathbf{N}$, $F(p) \leq p^k$ for all p close enough to zero. Thus, F is extremely flat at zero, in strong contrast to the case of uniform beliefs. Denote by g_θ the conditional density of the signal given θ and let $k \in \mathbf{N}$ be given. For a signal realization s and by Bayes rule, the private belief q is given by

$$\frac{q}{1-q} = \frac{g_H(s)}{g_L(s)} = \exp\left\{\frac{\Delta\mu}{\sigma^2}\left(s - \frac{\mu_H + \mu_L}{2}\right)\right\}$$

and is therefore increasing in s. Thus, for given $\bar{q} \in (0,1)$, one has

$$\begin{aligned} F_L(\bar{q}) = \mathbf{P}_L(q \leq \bar{q}) &= \mathbf{P}_L\left(\ln\frac{q}{1-q} \leq \ln\frac{\bar{q}}{1-\bar{q}}\right) \\ &= \mathbf{P}_L\left(s \leq \frac{\mu_H + \mu_L}{2} + \frac{\sigma^2}{\Delta\mu}\ln\frac{\bar{q}}{1-\bar{q}}\right) \end{aligned}$$

Since, conditional on $\theta = L$, the random variable $\dfrac{s - \mu_L}{\sigma}$ follows a $\mathcal{N}(0,1)$-distribution, one has for each $q \in [0,1]$,

$$F_L(q) = G(x_q) := G\left(\frac{\Delta\mu}{2\sigma} + \frac{\sigma}{\Delta\mu}\ln\frac{q}{1-q}\right),$$

where G is the cumulative distribution function of the $\mathcal{N}(0,1)$ distribution.

By (3), $F_L(q) \sim_0 2F(q)$,[17] hence it is sufficient to prove that $G(x_q) \leq q^k$ for q close to zero. We rely on the inequalities

$$\frac{|x|}{x^2+1}e^{-x^2/2} \leq G(x) \leq \frac{1}{|x|}e^{-x^2/2},$$

which hold for all $x < 0$, see Revuz and Yor (1999), p. 30. This implies that $G(x) \sim_{-\infty} \dfrac{1}{|x|}e^{-\frac{x^2}{2}}$.

Set $a := \dfrac{\Delta\mu}{\sigma}$, so that $x_q := \dfrac{a}{2} + \dfrac{1}{a}\ln\dfrac{q}{1-q}$. Standard computations yield

$$G(x_q) \sim_0 a\frac{\sqrt{q}}{|\ln q|}\exp\left\{-\frac{1}{2a^2}(\ln q)^2\right\}$$

and therefore, $G(x_q) \leq \exp\left\{-\dfrac{1}{2a^2}(\ln q)^2\right\}$ for q close enough to zero. This easily implies that $G(x_q) \leq q^k$ for q close to zero, as desired. ■

Conclusions are different when the variance is state-dependent.

[17] Given two functions f and g, and $a \in \mathbf{R} \cup \{-\infty, +\infty\}$, we write $f \sim_a g$ as a shortcut to $\lim_{x\to a}\frac{f(x)}{g(x)} = 1$.

Proposition 3 *Assume $\mu_H = \mu_L = 0$, and $\sigma_L > \sigma_H$. In state L, learning is efficient if and only if $\sigma_L^2 > 2\sigma_H^2$.*

Note that the ratio g_H/g_L of the conditional densities of the signal does not exceed σ_L/σ_H. Consequently, the support of the private belief is equal the interval $\left[0, \frac{\sigma_L}{\sigma_L+\sigma_H}\right]$. Thus, while $F(p) > 0$ for all $p > 0$, one has $F(\frac{\sigma_L}{\sigma_L+\sigma_H}) = 1$. Therefore, learning is inefficient in state H, but Theorem 3 applies as stated.[18]

Similar computations can be made for classical distributions. We only report results.[19] If signals follow a Levy distribution, with density $g_\theta(s) = \sqrt{\frac{c_\theta}{2\pi}} e^{-c_\theta/2s} \frac{1}{s^{3/2}}$ over $\mathbf{R}_+$, and $c_H > c_L > 0$, learning is always inefficient in state L. If signals follow an exponential distribution with density $g_\theta(s) = a_\theta e^{-a_\theta s}$ over $\mathbf{R}_+$, and $a_H > a_L > 0$, learning is efficient in state L iff $a_H > 2a_L$. If signals follow a Pareto distribution with density $g_\theta(s) = \alpha_\theta / s^{\alpha_\theta+1}$ over $\mathbf{R}_+$, and $\alpha_H > \alpha_L > 0$, learning is efficient in state L iff $\alpha_H > 2\alpha_L$.[20]

3 Further results

3.1 Rates of convergence

Assume that the distribution of signals is such that learning is inefficient. One may still wonder whether the two setups can be ranked on the basis of how quickly the expected number of incorrect guesses among the first n agents increases to $+\infty$. In a sense, this is a second-order issue, compared with the first-order inefficiency result, and getting clear insights under general distributional assumptions is technically challenging. We only provide some preliminary evidence, which suggests that rates of convergence to $+\infty$ are similar for the two setups. We focus on the case of uniform beliefs, and we denote by $N_n := |\{k \leq n, a_k \neq \theta\}|$ the number of incorrect guesses among the first n agents.

Proposition 4 *Assume $F(p) = p$ for all $p \in [0,1]$.*

- *If all actions are public,* $\mathbf{E}_L[\min(\tau, n)] \sim_{+\infty} \frac{1}{\pi} \ln n$.

[18] Because the symmetry assumption (**A2**) does not hold here, Theorem 5 does not apply.

[19] Proofs are available upon request.

[20] In all these cases, learning is inefficient in state H, because the ratio g_H/g_L is bounded, hence $F(p) = 1$ for some $p < 1$.

- *If only the previous action is observed,* $\mathbf{E}_L[N_n] \sim_{+\infty} \ln n$.

When actions are public, one has the additional bound $\mathbf{E}_L[\min(\tau, n)] \leq \mathbf{E}_L[N_n] \leq a\mathbf{E}_L[\min(\tau, n)]$ for some a and all $n \in \mathbf{N}$. Hence, Proposition 4 provides a sharp estimate of the rate of convergence to infinity.

In general, the exact rate of convergence obviously depends on the distribution F. For normally distributed signals (with state-independent variance), since F is extremely flat at zero, we would expect the expected number of incorrect guesses to be significantly higher.

3.2 Alternative feedbacks

An intriguing issue is whether the irrelevance result between the two settings generalizes to other ones. Alternatively, would a designer with the ability to set the rules of the game be able to foster efficiency?

It turns out that there is no general irrelevance result: the efficiency properties of some alternative settings are even poorer. Assume indeed, as in Smith and Sørensen (1996, 2013) that each agent observes the action of one earlier agent, randomly sampled, but is not told the rank of the sampled agent.[21]

Proposition 5 *Under these assumptions and (**A2**), learning is inefficient, irrespective of the distribution of signals.*

Yet, the setup of Proposition 5 does not qualify as an intermediate setup between our two main ones. Indeed, the feedback is coarser than when the previous action is observed, and in addition made noisier by the feature that identities are not disclosed. The question remains whether the amount of feedback provided to each agent can be fine-tuned so as to improve efficiency. We provide a partial analysis of this issue, for uniform beliefs, $F(p) = p$. One rationale for the choice of this case is that, as explained above, it represents a limiting one, between efficiency and inefficiency. Fostering efficiency here might therefore prove easier than for other signal distributions.

[21] We assume that sampling is independent across agents, and that all previous agents are equally likely to be sampled. If samples are positively correlated across agents, it is natural to expect that the efficiency properties will be even bleaker. On the other hand, we expect the negative result of Proposition 5 to remain valid if more recent agents are more likely to be sampled than earlier ones.

There are two natural and mutually exclusive approaches to parametrize alternative setups. One is to assume that memory fades away with time, so that an agent "remembers" only the most recent actions, but not necessarily only the last one. In the setting we instead focus on, all actions are eventually observed, but possibly with delay. To be more specific, agents are partitioned into successive generations. Each agent observes all actions from the previous generations of agents, but not the decisions of his peers. In effect, agents within a generation use the same "public" information in addition to their private signal. As a result, actions within a generation are conditionally *iid*, given θ.

We denote by d_k the size of the k-th generation, and by $\Delta_k := d_1 + \cdots + d_k$ the cumulative size of the first k generations. As soon as successive generations do not grow too quickly, learning remains inefficient.

Theorem 7 *Assume that $F(p) = p$ for all $p \in [0,1]$. Assume that the sequence $(d_{k+1}/d_k)_k$ of ratios of the sizes of consecutive generations is bounded. Then $\mathbf{E}_L[\tau] = +\infty$.*

The requirement that the sequence $(d_{k+1}/d_k)_k$ be bounded allows for generations whose size increases exponentially fast – but not faster.

In the absence of any constraint on the sequence (d_k), it is easy to construct recursively a sequence (d_k) such that $\mathbf{E}_L[\tau] < +\infty$. Consider indeed an arbitrary generation k and denote by $\pi(k) := \mathbf{P}(\theta = H \mid \tau > \Delta_{k-1})$ the (public) belief in generation k, in the event where all agents in all previous generations made the wrong guess h. Because the guesses in generation k are conditionally *iid* (and since the informativeness of private signals is not bounded), it is not difficult to check that for fixed Δ_{k-1}

$$\lim_{d_k \to +\infty} \Delta_k \mathbf{P}_L(\tau > \Delta_k \mid \tau > \Delta_{k-1}) = 0.$$

Setting inductively d_k large enough so that, say $\Delta_k \mathbf{P}_L(\tau > \Delta_k \mid \tau > \Delta_{k-1}) \leq \frac{1}{2^k}$ for all k, one then has $\mathbf{E}_L[\tau] < +\infty$. However, it is quite doubtful whether this is a reasonable measure of efficiency here. Indeed, if only a small fraction of generation k makes a correct guess, this is not clear evidence for generation $k+1$: the so-called overturning principle (see Smith and Sørensen (2000)) does not hold here, and it is not clear whether the expected number of incorrect guesses $\mathbf{E}_L[N]$ is finite or not.

3.3 Towards the first-best

In the benchmark case where signals are observed, information keeps flowing in an exogenous manner, allowing for exponentially fast learning. Our setups differ in two important and complementary dimensions. First, and because guesses are binary, information is garbled: even when all guesses are observed, agents only get a coarsened version of the available information. Second, the encoding process of a signal into a guess is the result of optimizing behavior by the agents, and is therefore endogenous. Which of these two dimensions, coarse information transmission and strategic behavior, is accountable for the lack of efficiency?

Assume that there is a continuum of possible actions. Under many distributional assumptions (and many specifications of the agents' objectives), one may expect that individual actions will be finely tuned to one's private signal, in a one-to-one manner. As a result, actions can be uniquely decoded, and the private signals recovered from the sequence of actions: the sequence of guesses is the same as if all private signals were observed. Such a situation has been studied, and coined *perfect learning* in Arieli and Mueller-Frank (2016), or *efficient learning* in Chamley (2004).

Assume now that guesses are binary. Would an uninformed social planner, with the ability to dictate strategies subject to the informational constraints of the game, be able to restore efficiency? As we show in Theorem 8, the answer is positive, no matter what the distribution of signals is.[22] Thus, this suggests that inefficiency is the result of the combination of a coarse action set, and of a strategic behavior.

We do not attempt a characterization of the social optimum, and only aim at finding individual strategies such that $\mathbf{E}_L[N] < +\infty$. This allows us to focus on simple decision rules with reasonably good properties (as opposed to optimal, complex ones). Not surprisingly, we focus on cut-off rules – choose $a_n = h$ if and only if the private belief of agent n is above some history-dependent cut-off. This cut-off is restricted to depend only on (i) the previous guess and on (ii) when the last guess switch took place. Thus, the cut-off is a function of the length of the current string of identical guesses, and of the previous guess. In addition, we focus on stationary profiles of such decision rules. Such profiles are described by two sequences $(\bar{q}_k)_k$ and $(\tilde{q}_k)_k$ in $(0,1)$ with the interpretation that $\bar{q}_k$ (resp., $\tilde{q}_k$) is the cut-off used by agent n if $a_{n-1} = l$ (resp., $a_{n-1} = h$) and the last action switch occurred exactly k rounds earlier.

[22] As long as $0 < F(p) < 1$ holds for $p \in (0,1)$.

We recall that $N := |\{n : a_n \neq \theta\}|$ is the total number of wrong guesses and we denote by $\tilde{\tau} := \inf\{n : a_n \neq \theta\}$ the first agent making a wrong guess. We prove Theorem 8 below under the additional requirement (**A3**): for some $c > 0$ and $\gamma > 0$, one has $F(p) \geq cp^\gamma$ in a neighborhood of zero, thereby ruling out distributions of private beliefs with extremely thin tails.[23]

Theorem 8 *Assume $0 < F(p) < 1$ for all $p \in (0,1)$ and (**A3**). There are sequences $(\bar{q}_k)$ and $(\tilde{q}_k)$ such that*

$$\mathbf{E}_\theta[N] < +\infty \text{ and } \mathbf{E}_\theta[\tilde{\tau}] = +\infty, \text{ for each } \theta \in \{L, H\}.$$

This is the strongest statement one may hope for. In expectation, the first incorrect guess is made infinitely far in the future, and the total number of incorrect guesses is finite.

At equilibrium, an agent's behavior only depends on his posterior belief p_n being above/below $\frac{1}{2}$, and learning inefficiency occurs because switches occur too rarely. The sequences $(\bar{q}_k)$ and $(\tilde{q}_k)$ in Theorem 8 have the feature that agents are more reluctant to follow a herd than equilibrium behavior would dictate. Put differently, if the previous guess was $a_{n-1} = h$, agent n is required to switch to $a_n = l$ even his posterior belief is somewhat above $\frac{1}{2}$.

To get insights into the magnitude of the departure from equilibrium play, we again investigate what happens with uniform beliefs, and focus on the very first herd. In the event where $a_1 = \cdots = a_{n-1} = h$, agent n holds an *interim* belief that is determined by the cut-off values $\bar{q}_1, \ldots, \bar{q}_{n-1}$. We denote by $\bar{p}_n$ the value of the *posterior* belief of agent n when his *private* belief is equal to the cut-off value $\bar{q}_n$.

Proposition 6 *Assume $F(p) = p$ for all $p \in [0,1]$. For each $a \in (\frac{1}{2}, 1)$, the conclusions of Theorem 8 hold with $\bar{q}_k = \frac{1}{(k+1)^a} = 1 - \tilde{q}_k$. For such choices, one has $\bar{p}_n \geq 1 - e^{-n^a}$ for all n large enough.*

At the posterior belief $\bar{p}_n$, the statistical evidence in favor of H is overwhelming. Thus, if many agents have guessed h so far, agent n is instructed to switch to $a_n = l$ as soon as he is less than certain that $\theta = H$. This shows that the behavior of agent n at the cut-off $\bar{q}_n$ markedly differs from what is individually optimal.

[23] Assumption (**A3**) is not innocuous, unlike our standing assumption on F. For instance, normally distributed signals fail (**A3**).

References

[1] Acemoglu, D., Dahleh, M., Lobel, I. and A. Ozdaglar (2011). "Bayesian learning in social networks", *Review of Economic Studies*, **78**, 1201–1236.

[2] Acemoglu, D. and A. Ozdaglar (2010). "Opinion Dynamics and Learning in Social Networks", *Dynamic Games and Applications*, **1**, 3–49.

[3] Arieli, I. and M. Mueller-Frank (2016). "Inferring beliefs from actions", *Games and Economic Behavior*, forthcoming.

[4] Aumann, R.J. and M. Maschler (1995). *Repeated Games with Incomplete Information*, MIT Press.

[5] Benaim, M. (1999). *Dynamics of stochastic approximation algorithms*, Séminaire de Probabilités (Strasbourg), tome 3, pages 1–68, Springer-Verlag.

[6] Banerjee, A.V. (1992). "A simple model of herd behavior", *Quaterly Journal of Economics*, **1992**, 797–817.

[7] Banerjee, A..V. and D. Fudenberg (2004). "Word-of-mouth learning", *Games and Economic Behavior*, **46**, 1–22.

[8] Bikhchandani, D. Hirshleifer and I. Welch (1992). "A theory of fads, fashion, custom, and cultural change as informational cascades", *Journal of Political Economy*, **100**, 992–1026.

[9] Celen, B. and S. Kariv (2004). "Observational Learning Under Imperfect Information", *Games and Economic Behavior*, **47**, 72–86.

[10] Chamley, C. (2004). *Rational herds: economic models of social learning*, Cambridge University Press.

[11] DeMarzo, P., D. Vayanos and J. Zwiebel (2003). "Persuasion Bias, Social Influence, and Unidimensional Opinions", *Quarterly Journal of Economics*, **118**, 909–968.

[12] Hann-Caruthers, W., Martynov, V. and O. Tamuz (2017). "The speed of sequential asymptotic learning", mimeo.

[13] Golub, B. and M. Jackson (2010). “Naive Learning in Social Networks and the Wisdom of Crowds", **2**, 112–149.

[14] Harel, M., Mossel, E., Strack, P. and O. Tamuz (2017). “The speed of social learning", mimeo.

[15] Lobel, I. and E. Sadler (2015). “Information Diffusion in Networks through Social Learning", *Theoretical Economics*, **10**, 807–851.

[16] Lobel, I., D. Acemoglu, Dahleh, M., and A. Ozdaglar (2009). “Rate of Convergence of Learning in Social Networks", *Proceedings of the American Control Conference (ACC)*,

[17] Mossel, E., A. Sly and O. Tamuz (2014). “Asymptotic learning on bayesian social networks", *Probability Theory and Related Fields*, **158**, 127–157.

[18] Mossel, E., A. Sly and O. Tamuz (2015). “Strategic learning and the topology of social networks", *Econometrica*, **83**, 1755–1794.

[19] Monzon, I. and M. Rapp (2014). “Observational Learning with Position Uncertainty", *Journal of Economic Theory*, 375–402.

[20] Pemantle, R (2007). “A survey of random processes with reinforcement", *Probability Surveys*, **4**, 1–79.

[21] Rosenberg, D., E. Solan and N. Vieille (2009). “Informational externalities and emergence of consensus", *Games and Economic Behavior*,**66**, 979–994.

[22] Revuz, D. and M. Yor (1999). *Continous Martingales and Brownian Motion*, Springer-Verlag, 3rd edition.

[23] Smith, L. and P. Sørensen (1996, 2013). “Rational Social Learning with Random Sampling", mimeo.

[24] Smith, L. and P. Sørensen (2000). “Pathological Outcomes of Observational Learning", *Econometrica*, **68**, 371–398.

A The benchmark case: Proof of Theorem 1

We denote by g_θ the conditional density of a private signal, given θ. Remember that agent n picks action h if and only if the posterior probability p_n assigned to the state being H is at least $\frac{1}{2}$, or equivalently, if the log-likelihood ratio of this belief is non-negative. When all signals $s_1, \ldots, s_n$ are observed, and by Bayes rule, one has $\dfrac{p_n}{1-p_n} = \displaystyle\sum_{i=1}^{n} \ln \frac{g_H(s_i)}{g_L(s_i)}$ or equivalently, $\dfrac{p_n}{1-p_n} = \displaystyle\sum_{i=1}^{n} \ln \frac{q_i}{1-q_i}$.

We set $X_i := \ln \dfrac{q_i}{1-q_i}$, so that $\mathbf{P}_L(a_n = h) = \mathbf{P}_L\left(\displaystyle\sum_{i=1}^{n} X_i \geq 0\right)$. Introduce the Laplace transform $\phi_X(t) := \mathbf{E}_L\left[e^{tX_1}\right]$ for all $t \geq 0$, and set

$$h_X(a) := \sup_{t \geq 0} \left(at - \ln \phi_X(t)\right) \text{ for } a \in \mathbf{R}.$$

The function ϕ_X is always non-negative, but may assume infinite values, in which case we set $\ln \phi_X(t) = +\infty$. By Chernoff bound, one has

$$\mathbf{P}_L\left(\sum_{i=1}^{n} X_i \geq na\right) \leq e^{-nh_X(a)} \text{ for all } n \in \mathbf{N} \text{ and } a \in \mathbf{R}.$$

For $a = 0$, we obtain $\mathbf{P}_L(a_n = h) \leq e^{-nh_X(0)}$.

Lemma 1 *One has $h_X(0) > 0$.*

Proof. By definition,

$$\phi_X(t) = \mathbf{E}_L\left[e^{t \ln \frac{q}{1-q}}\right] = \int_0^1 f_L(q) \frac{q^t}{(1-q)^t} dq.$$

Obviously, $\phi_X(0) = 1$. On the other hand, the private belief q_n is a sufficient statistic for θ given s_n: by the iterated law of conditional expectations, $\mathbf{E}\left[1_{\theta=H} \mid q_n\right] = \mathbf{E}\left[q_n \mid q_n\right] = q_n$. Consequently, $\dfrac{q}{1-q} = \dfrac{f_H(q)}{f_L(q)}$ hence

$$\phi_X(1) = \int_0^1 f_L(p) \frac{p}{1-p} dp = \int_0^1 f_H(p) dp = 1$$

as well.

Define $T := \{t \in \mathbf{R}, \mathbf{E}_L\left[e^{tX_1}\right] < +\infty\}$. Since $t \mapsto e^{tx}$ is strictly convex for almost every x, the set T is an interval and ϕ_X is strictly convex on T. Since $\phi_X(0) = \phi_X(1) = 1$, it follows that $\phi_X(t) < 1$ for all $t \in (0,1)$. Therefore, $h_X(0) \geq -\inf_{t\in[0,1]} \ln \phi_X(t) > 0$. ■

Since the total number of wrong guesses is $N = \sum_{n\geq 1} 1_{a_n \neq \theta}$, one obtains

$$\mathbf{E}_L[N] = \sum_{n\geq 1} \mathbf{P}_L(a_n = h) \leq \sum_{n\geq 1} e^{-nh_X(0)} = \frac{1}{e^{h_X(0)} - 1}.$$

This proves Theorem 1.

When private beliefs are uniformly distributed ($F(p) = p$), one has $f_L(p) = 2(1-p)$, and $\phi_X(t) = 2\int_0^1 p^t(1-p)^{1-t}dp$. The latter integral is finite if and only if $t \in (-1,2)$, hence ϕ_X is well-defined and convex on $(-1,2)$. In addition, $\phi_X(t) = \phi_X(1-t)$ for each $t \in (-1,2)$. Hence $\inf_{t\geq 0} \phi_X(t) = \inf_{t\in(-1,2)} \phi_X(t) = \phi_X\left(\frac{1}{2}\right)$. Since

$$\phi_X\left(\frac{1}{2}\right) = 2\int_0^1 \sqrt{p(1-p)}dp = \frac{\pi}{4},$$

this yields $h_X(0) = -\ln\frac{\pi}{4}$, hence $\mathbf{E}_L[N] \leq \frac{\pi}{4-\pi}$, as claimed.

B All guesses are public: Proofs of Theorem 3 and 4

B.1 Proof of Theorem 3

Recall that $q_n := \mathbf{P}(\theta = H \mid s_n)$ is the private belief, $\pi_n = \mathbf{P}(\theta = H \mid a_1, \ldots, a_{n-1})$ the interim belief, and $p_n := \mathbf{P}(\theta = H \mid s_n, a_1, \ldots, a_{n-1})$ the posterior belief of agent n.

Lemma 2 *One has*

$$\mathbf{E}_L[\tau] = 1 + \sum_{n=1}^{+\infty} \prod_{k=1}^{n} \left(1 - F_L(1-\pi_k^*)\right),$$

where the sequence (π_k^*) *is given by* $\pi_1^* = 1/2$ *and*

$$\frac{\pi_{n+1}^*}{1-\pi_{n+1}^*} = \frac{\pi_n^*}{1-\pi_n^*} \times \frac{1-F_H(1-\pi_n^*)}{1-F_L(1-\pi_n^*)}.$$

Proof. This identity follows from Bayesian updating. Let $n \geq 1$ be given. Recall from (1) that $a_n = h$ if and only if $q_n \geq 1 - \pi_n$. Hence $\mathbf{P}_\theta(a_n = h \mid a_1, \ldots, a_{n-1}) = 1 - F_\theta(1 - \pi_n)$.

On the event $a_n = h$, the updating of the interim/public beliefs is given by

$$\frac{\pi_{n+1}}{1 - \pi_{n+1}} = \frac{\pi_n}{1 - \pi_n} \times \frac{1 - F_H(1 - \pi_n)}{1 - F_L(1 - \pi_n)}. \tag{7}$$

The event $\{\tau > n\}$ coincides with the event $\{a_1 = \cdots = a_n = h\}$ and therefore with $\bigcap_{k=1}^{n} \{q_k \geq 1 - \pi_k^*\}$, where $(\pi_k^*)_k$ is the sequence defined recursively in (7), or as in the statement of the lemma. Because the private beliefs (q_k) are conditionally *iid* given θ, it follows that

$$\mathbf{P}_L(\tau > n) = \prod_{k=1}^{n} \left(1 - F_L(1 - \pi_k^*)\right).$$

The result now follows from the identity

$$\mathbf{E}_L[\tau] = \sum_{n=1}^{+\infty} \mathbf{P}_L(\tau \geq n).$$

■

In the rest of the proof, we will write π_n instead of π_n^*, keeping in mind that it is the (deterministic) sequence of public beliefs held along the specific history where all agents play h. We set $r_n := \ln \frac{\pi_n}{1 - \pi_n}$ and $u_n := \prod_{k=1}^{n} (1 - F_L(1 - \pi_k)) = \mathbf{P}_L(\tau \geq n + 1)$, so that $\mathbf{E}_L[\tau] = 1 + \sum_{n=1}^{+\infty} u_n$. The updating equation (7) rewrites

$$r_{n+1} - r_n = \ln \frac{1 - F_H(\frac{1}{1+e^{r_n}})}{1 - F_L(\frac{1}{1+e^{r_n}})}. \tag{8}$$

We note that $\lim_{n \to +\infty} r_n = +\infty$, or equivalently, $\lim_{n \to +\infty} \pi_n = 1$. To see why, observe that (r_n) is increasing by (8). If the sequence (r_n) were bounded from above, and using (8), $r_{n+1} - r_n$ would be bounded below by a positive number – indeed, F_H/F_L is bounded away from one on any interval $[0, \bar{p}]$ such that $\bar{p} < 1$. This would in turn imply $r_n \to +\infty$ – a contradiction. That $r_n \to \infty$ implies that $r_{n+1} - r_n \to 0$, using again (8).

We prove the result through a series of equivalences.

Lemma 3 *One has*

$$\sum_{n=1}^{+\infty} u_n < +\infty \;\Leftrightarrow\; \sum_{n=1}^{+\infty} e^{-r_n} < +\infty.$$

Proof. Observe that

$$1 - F_L(1-\pi_n) = (1 - F_H(1-\pi_n))\, e^{r_n - r_{n+1}} \tag{9}$$

which, since $r_1 = 0$, implies

$$u_n = e^{-r_{n+1}} \prod_{k=1}^{n} (1 - F_H(1-\pi_k)).$$

Thus, $u_n \leq e^{r_{n+1}}$ for all n, so that the reverse implication of the lemma trivially holds.

Assume now that the series $\sum u_n$ is convergent. It suffices to prove that the infinite product $\beta := \prod_{k=1}^{+\infty} (1 - F_H(1-\pi_k))$ is strictly positive. Indeed, one then has $e^{-r_{n+1}} \leq \frac{1}{\beta} u_n$, hence the convergence of $\sum e^{-r_n}$ follows from that of $\sum u_n$.

Assume to the contrary that the infinite product β is zero. This implies that

$$\sum_{k=1}^{+\infty} \ln(1 - F_H(1-\pi_k)) = -\infty \text{ or equivalenty, } \sum_{k=1}^{+\infty} F_H(1-\pi_k) = +\infty.$$

Note now that

$$F_H(1-\pi_n) \leq F_H(e^{-r_n}) \leq 2e^{-r_n} F(e^{-r_n}) \leq 2Ce^{-(1+\alpha)r_n} :$$

the first inequality follows from the inequality $1 - \pi_n \leq e^{-r_n}$, the second one from (3), and the last one from our standing assumption on F. Hence, $\displaystyle\sum_{n=1}^{+\infty} e^{-(1+\alpha)r_n} = +\infty$.

Pick $\varepsilon > 0$ such that $(1-\varepsilon)(1+\alpha) > 1$. Since $\lim_{n\to+\infty}(1-\pi_n) = 0$ and $\lim_{p\to 0} \frac{F_H(p)}{F_L(p)} = 0$, one has $\displaystyle\lim_{n\to+\infty} \frac{\ln(1 - F_H(1-\pi_n))}{\ln(1 - F_L(1-\pi_n))} = 0$. Hence, there exists N_0 such that

$$\ln(1 - F_H(1-\pi_n)) \geq \frac{\varepsilon}{2} \ln(1 - F_L(1-\pi_n)) \text{ for all } n \geq N_0. \tag{10}$$

Since $F_H \leq F_L$, one has $\displaystyle\sum_{n=1}^{+\infty} \ln(1 - F_L(1-\pi_n)) = -\infty$ as well, and there is $N_1 \in \mathbf{N}$ such that

$$\sum_{k=1}^{N_0-1} \ln(1 - F_L(1-\pi_k)) \geq \frac{\varepsilon}{2} \sum_{k=1}^{n} \ln(1 - F_L(1-\pi_k)) \text{ for all } n \geq N_1$$

so that, by (10), one has for all $n \geq N_1$,

$$\begin{aligned}\sum_{k=1}^{n} \ln(1 - F_H(1-\pi_k)) &\geq \sum_{k=1}^{N_0-1} \ln(1 - F_L(1-\pi_k)) + \sum_{k=N_0}^{n} \ln(1 - F_H(1-\pi_k)) \\ &\geq \varepsilon \sum_{k=1}^{n} \ln(1 - F_L(1-\pi_k)). \qquad (11)\end{aligned}$$

Taking logarithms in (9), summing over $k \in \{1, \ldots, n\}$, and using (11), we get

$$(1-\varepsilon) \sum_{k=1}^{n} \ln(1 - F_L(1-\pi_k)) \geq -r_{n+1} \text{ for all } n \geq N_1$$

and therefore, by the choice of ε,

$$\sum_{k=1}^{n} \ln(1 - F_L(1-\pi_k)) \geq -(1+\alpha) r_{n+1} \text{ for all } n \geq N_1.$$

That is, $u_n \geq e^{-(1+\alpha)r_{n+1}}$ for all $n \geq N_1$, hence $\sum_{n \in \mathbf{N}} u_n = +\infty$. This is the desired contradiction. ■

The series $\sum e^{-r_n}$ is somewhat easier to manipulate than the series $\sum u_n$. Yet, it still involves the recursively defined sequence (r_n). We next rely on a time-change technique to get a series $\sum v_n$, where (v_n) is an explicitly defined sequence. Define $\Phi_0(x) := \ln \dfrac{1 - F_H(\frac{1}{1+e^x})}{1 - F_L(\frac{1}{1+e^x})}$ for all x. Note that $\Phi_0(x) = \ln \dfrac{\tilde{F}_H\left(1 - \frac{1}{1+e^x}\right)}{\tilde{F}_L\left(1 - \frac{1}{1+e^x})\right)}$, where $\tilde{F}(q) := 1 - F(1-q)$ is the cdf of the (private) belief assigned to L. Since $q \mapsto \dfrac{\tilde{F}_H(q)}{\tilde{F}_L(q)}$ is non-decreasing, the map Φ_0 is non-increasing.

Note also that $r_{n+1} - r_n = \Phi_0(r_n)$ for all $n \geq 1$.

Lemma 4 *One has*

$$\sum_{n=1}^{+\infty} e^{-r_n} < +\infty \;\Leftrightarrow\; \sum_{k=0}^{+\infty} \frac{e^{-k}}{\Phi_0(k)} < +\infty.$$

Proof. For $k \geq 0$, we set $\omega_k := \inf\{n \geq 1 : r_n \geq k\}$. Since (r_n) increases to $+\infty$, one has $\omega_k < +\infty$ for each k. Since $(r_{n+1} - r_n)_n$ decreases to zero, the sequence (ω_k) is

eventually strictly increasing. Choose $k_0 \in \mathbf{N}$ such that $r_{n+1} - r_n < \frac{1}{2}$ for all $n \geq \omega_{k_0}$. and for each $K \geq k_0$, set

$$S_K := \sum_{n=\omega_{k_0}}^{\omega_K - 1} e^{-r_n} = \sum_{k=k_0}^{K-1} \sum_{n=\omega_k}^{\omega_{k+1}-1} e^{-r_n}.$$

The convergence of the sequence $(S_K)_K$ is equivalent to the convergence of the series $\sum e^{-r_n}$.

From the choice of k_0, one has $\omega_{k+1} > \omega_k$ for each $k \geq k_0$. Since $r_{\omega_k} \geq k > r_{\omega_k - 1}$ for each k, one gets in turn $r_{\omega_{k+1}} - r_{\omega_k} \geq (k+1) - (k + \frac{1}{2}) = \frac{1}{2}$ and $r_{\omega_{k+1}} - r_{\omega_k} \leq k + \frac{3}{2} - k = \frac{3}{2}$ for each $k \geq k_0$, hence

$$\frac{1}{2} \leq r_{\omega_{k+1}} - r_{\omega_k} \leq \frac{3}{2}. \tag{12}$$

On the other hand, note that

$$\sum_{k=k_0}^{K-1} (\omega_{k+1} - \omega_k) e^{-(k+1)} \leq S_K \leq \sum_{k=k_0}^{K-1} (\omega_{k+1} - \omega_k) e^{-k}. \tag{13}$$

Since $r_{n+1} - r_n = \Phi_0(r_n)$ for all n,

$$r_{\omega_{k+1}} - r_{\omega_k} = \sum_{n=\omega_k}^{\omega_{k+1}-1} \Phi_0(r_n),$$

and since Φ_0 is non-increasing, we deduce from (12) an upper bound

$$r_{\omega_{k+1}} - r_{\omega_k} \geq (\omega_{k+1} - \omega_k) \Phi_0(r_{\omega_{k+1}-1}) \geq (\omega_{k+1} - \omega_k) \Phi_0(k+1)$$

and the lower bound

$$r_{\omega_{k+1}} - r_{\omega_k} \leq (\omega_{k+1} - \omega_k) \Phi_0(r_{\omega_k}) \leq (\omega_{k+1} - \omega_k) \Phi_0(k)$$

on $\omega_{k+1} - \omega_k$, hence

$$\frac{1}{2\Phi_0(k)} \leq \omega_{k+1} - \omega_k \leq \frac{3}{2\Phi_0(k+1)}.$$

Plugging into (13), one gets

$$\frac{1}{2e} \sum_{k=k_0}^{K-1} \frac{e^{-k}}{\Phi_0(k)} \leq S_K \leq \frac{3}{2} \sum_{k=k_0}^{K-1} \frac{e^{-k}}{\Phi_0(k+1)} = \frac{3e}{2} \sum_{k=k_0+1}^{K} \frac{e^{-k}}{\Phi_0(k)}.$$

Hence the convergence of the series $\sum e^{-r_n}$ is equivalent to the convergence of the series $\sum \frac{e^{-k}}{\Phi_0(k)}$, as desired. ■

Lemma 5 *One has*

$$\sum_{k=0}^{+\infty} \frac{e^{-k}}{\Phi_0(k)} < +\infty \;\Leftrightarrow\; \sum \frac{e^{-k}}{F\left(\frac{1}{1+e^k}\right)} < +\infty.$$

Proof. Since $\lim_{p\to 0} F_H(p)/F_L(p) = 0$, one has

$$\frac{1-F_H(p)}{1-F_L(p)} = 1 + F_L(p) + o\left(F_L(p)\right). \tag{14}$$

By (3) and since $\int_0^p F(x)dx \leq pF(p)$, one has

$$F_L(p) = 2(1-p)F(p) - 2\int_0^p F(x)dx \sim_0 2F(p).$$

Substituting into (14),

$$\frac{1-F_H(p)}{1-F_L(p)} = 1 + 2F(p) + o(F(p)).$$

Taking logarithms, this implies $\Phi_0(k) \sim_{+\infty} 2F\left(\frac{1}{1+e^k}\right)$. The result follows. ■

Lemma 6 *One has*

$$\sum \frac{e^{-k}}{F\left(\frac{1}{1+e^k}\right)} < +\infty \;\Leftrightarrow\; \int_0^{+\infty} \frac{e^{-x}}{F(\frac{1}{1+e^x})} dx < +\infty.$$

Proof. This follows from a standard comparison argument. Because $y \mapsto \frac{e^{-y}}{F\left(\frac{1}{1+e^y}\right)}$ is the quotient of two decreasing functions, one has

$$\frac{e^{-(n+1)}}{F\left(\frac{1}{1+e^n}\right)} \leq \int_n^{n+1} \frac{e^{-y}}{F\left(\frac{1}{1+e^y}\right)} dy \leq \frac{e^{-n}}{F\left(\frac{1}{1+e^{n+1}}\right)}.$$

By summation, this implies in turn that

$$\frac{1}{e}\sum_{k=0}^{n-1} \frac{e^{-k}}{F\left(\frac{1}{1+e^k}\right)} \leq \int_0^n \frac{e^{-y}}{F\left(\frac{1}{1+e^y}\right)} dy \leq e\sum_{k=1}^{n} \frac{e^{-k}}{F\left(\frac{1}{1+e^k}\right)}.$$

■

Lemma 7 *One has*

$$\int_0^{+\infty} \frac{e^{-x}}{F(\frac{1}{1+e^x})} dx < +\infty \;\Leftrightarrow\; \int_0^1 \frac{1}{F(p)} dp < +\infty.$$

Proof. A change of variable $p = \dfrac{1}{1+e^x}$ shows that

$$\int_0^{+\infty} \frac{e^{-x}}{F(\frac{1}{1+e^x})} dx = \int_0^{\frac{1}{2}} \frac{1}{(1-p)^2 F(p)} dp.$$

The latter integral is finite if and only if the integral $\displaystyle\int_0^1 \frac{1}{F(p)} dp$ is finite. ■

B.2 Proof of Theorem 4

We will prove that, under **(A1)**, $\mathbf{E}_L[\tau] < +\infty$ implies $\mathbf{E}_L[N] < +\infty$. We follow the sketch given in the main text, arguing that the total number of h-herds has a finite expectation, and that the (conditional) expected duration of each h-herd is uniformly bounded.

We start with some notation, next provide some technical preparations, before we present the proof.

Given θ, the sequence $(\pi_n)_n$ of interim/public beliefs is a Markov chain with values in $[0,1]$. Until now, the prior on $\theta = H$ has been kept equal to $\bar{\pi} = \frac{1}{2}$. As is often convenient with Markov chains, we will let this prior vary. In the rest of the proof, we write $\mathbf{P}_{\pi,\theta}$ the distribution on plays, conditional on θ, when the initial probability assigned to H is π. We note that $\pi_{n+1} < \frac{1}{2}$ if $a_n = l$ and $\pi_{n+1} > \frac{1}{2}$ if $a_n = h$.

Denote by (τ_k) and $(\tilde{\tau}_k)$ the successive dates at which agents switch from one action to the other. Formally, $\tau_1 = \inf\{n : a_n = l\}$ and, for $k \geq 1$, $\tilde{\tau}_k := \inf\{n > \tau_k : a_n = h\}$ if $\tau_k < +\infty$, and $\tau_{k+1} := \inf\{n > \tilde{\tau}_k : a_n = l\}$ if $\tilde{\tau}_k < +\infty$ (with $\inf \emptyset = +\infty$). With these notations, the first h-herd starts in round 1 and ends in round $\tau_1 - 1$ (it is therefore empty if $a_1 = h$). The $k+1$-th h-herd starts in round $\tilde{\tau}_k$, and ends in round $\tau_{k+1} - 1$.

We denote by $\mathcal{H}_n = \sigma(a_1, \ldots, a_{n-1})$ the (σ-algebra of) publicly available information in round n. Note that τ_k is not a stopping time, but that $\tau_k + 1$ is.

Let $\Phi_1 : [\bar{\pi}, 1) \to [\bar{\pi}, 1)$ be the continuous map which assigns to each interim belief π the interim belief $\Phi_1(\pi)$ in the next round, in the event where the current agent chooses h – that is, $\Phi_1(\pi)$ is defined by

$$\frac{\Phi_1(\pi)}{1-\Phi_1(\pi)} = \frac{\pi}{1-\pi} \times \frac{1-F_H(1-\pi)}{1-F_L(1-\pi)}.$$

It is clear that the map Φ_1 is continuous, and that the sequence $\left(\Phi_1^k(\pi)\right)_k$ of iterates is increasing for each π, with $\cup_{k\geq 0}\left[\Phi_1^k(\bar{\pi}), \Phi_1^{k+1}(\bar{\pi})\right) = [\bar{\pi}, 1)$. Note in addition that $\Phi_1(\pi) > \pi$ for all π.

As earlier, we denote by $\tilde{F}_\theta(p) = 1 - F_\theta(1-p)$ the cdf of the private belief assigned to state L. Note that $\tilde{F}'_\theta(p) = f(1-p)$.

Claim 9 *There exists $\pi_* < 1$, such that Φ_1 is non-decreasing on $(\pi_*, 1)$.*

Proof. By (**A1**), $f(p) \sim_0 cp^\beta$, hence $F(p) \sim_0 \frac{c}{\beta+1}p^{\beta+1}$. Since $\int_0^1 \frac{1}{F(p)}dp < +\infty$, we obtain $\beta \in (-1, 0)$, hence $\lim_{p\to 0} pf(p) = 0$.

Define

$$\Psi(x) := \frac{\Phi_1(x)}{1-\Phi_1(x)} = \frac{x}{1-x} \times \frac{\tilde{F}_H(x)}{\tilde{F}_L(x)}.$$

We will prove that Ψ is increasing on some interval $(\pi_*, 1)$. By (3), one has

$$\Psi(x) = \frac{x}{1-x} \times \frac{(1-x)\tilde{F}(x) + \int_0^x \tilde{F}(y)dy}{xF(x) - \int_0^x F(y)dy} = \frac{\phi(x)}{\phi(x) - \psi(x)},$$

where $\phi(x) := x(1-x)\tilde{F}(x) + x\int_0^x \tilde{F}(y)dy$ and $\psi(x) = \int_0^x \tilde{F}(y)$. Thus, the sign of $\Psi'(x)$ is the same as the sign of $-\phi'(x)\psi(x) + \phi(x)\psi'(x)$. Elementary algebra shows that $\Psi'(x) \geq 0$ if and only if

$$-\int_0^x \tilde{F}(y)dy\left((1-x)\tilde{F}(x) + \int_0^x \tilde{F}(y)dy + x(1-x)f(1-x)\right) + \tilde{F}(x)\left(x(1-x)\tilde{F}(x) + x\int_0^x \tilde{F}(y)dy\right) \tag{15}$$

Since $\displaystyle\int_0^1 F(p)dp = \frac{1}{2}$ and $\lim_{x\to 1}(1-x)f(1-x) = 0$, the left-hand side of (15) converges to $\frac{1}{4}$ when $x \to 1$. ■

We use the previous claim to prove in the next two claims that the probability that a correct herd lasts forever is bounded away from zero, as soon as the prior belief assigned to the correct state is bounded away from zero.

Claim 10 *One has $\lambda(\pi) := \mathbf{P}_{\pi,H}(\tau_1 = +\infty) > 0$ for all $\pi \in (0, 1)$.*

Proof. Assume to the contrary that $\lambda(\pi) = 0$ for some $\pi \in (0, 1)$, and observe that $\lambda(\pi) = \mathbf{P}_{\pi,H}(\tau_1 \geq 2) \times \lambda(\Phi_1(\pi))$. Since $\mathbf{P}_{\pi,\theta}(\tau_1 \geq 2) = \mathbf{P}_{\pi,\theta}(a_1 = h) > 0$, it follows that $\lambda(\Phi_1(\pi) = \mathbf{P}_{\Phi_1(\pi),H}(\tau_1 = +\infty) = 0$. Since $\mathbf{P}_{\Phi_1(\pi),L}(\tau_1 < +\infty) \leq \mathbf{P}_{\Phi_1(\pi),H}(\tau_1 < +\infty)$, this implies $\mathbf{P}_{\Phi_1(\pi)}(\tau_1 < +\infty) = 1$. Since the sequence $(\pi_n)_{n\geq 2}$ is a bounded martingale

under $\mathbf{P}_{\Phi_1(\pi)}$, it follows from the optional sampling theorem that $\mathbf{E}_{\Phi_1(\pi)}[\pi_{\tau_1+1}] = \Phi_1(\pi)$. But $\pi_{\tau_1+1} < \frac{1}{2}$, *a.s.*, while $\Phi_1(\pi) > \frac{1}{2}$. This is the desired contradiction. ■

Not only is $\lambda(\pi)$ positive, but it is uniformly bounded away from zero, when π is bounded away from zero. We recall that $\bar{\pi} = \frac{1}{2}$.

Claim 11 *λ is continuous on $(0,1)$ and $\inf_{\pi\in[\bar{\pi},1)} \mathbf{P}_{\pi,H}(\tau_1 = +\infty) > 0$.*

Proof. We first observe that λ is non-decreasing on $(\pi_*, 1)$, where π_* is the value obtained in Claim 9. Indeed, let $\pi \geq \pi' > \pi_*$. Then

$$\lambda(\pi) = \prod_{n=1}^{+\infty} \left(1 - F_H(1 - \Phi_1^{n-1}(\pi))\right) \geq \prod_{n=1}^{+\infty} \left(1 - F_H(1 - \Phi_1^{n-1}(\pi'))\right) = \lambda(\pi'),$$

where the inequality holds since $\Phi^{n-1}(\pi) > \pi_*$ for all n.

We now argue that λ is continuous on $(\pi_*, 1)$. Let indeed $(\pi(i))_i$ be a sequence in $(\pi_*, 1)$ such that $\pi(i) \to \pi > \pi_*$, and pick a lower bound $\pi^m > \pi_*$ for the sequence $(\pi(i))$. Since $\lambda(\pi^m) > 0$, one has $\sum_{n\geq 1} F_H(1 - \Phi^{n-1}(\pi^m)) < +\infty$. Since $F_H \circ (1 - \Phi^{n-1})$ is continuous and $F_H(1 - \Phi^{n-1}(\pi(i)) \leq F_H(1 - \Phi^{n-1}(\pi^m))$ for all i and n, it follows from the dominated convergence theorem that $\lim_{i\to+\infty} \lambda(\pi(i)) = \lambda(\pi)$.

We next prove that λ is continuous on $(0,1)$. Let $\pi \in (0,1)$ be given, and let $k \in \mathbf{N}$ be such that $\Phi_1^k(\pi) > \pi_*$. Let V be a neighborhood of π such that $\Phi_1^k(\pi') > \pi_*$ for each $\pi' \in V$. Since $\lambda(\pi') = \mathbf{P}_{\pi',H}(\tau_1 > k)\lambda(\Phi_1^k(\pi'))$, the map λ is continuous on V as a product of two continuous functions.

We turn to the second statement of the claim. Set $b := \min_{\pi\in[\bar{\pi},\Phi_1(\bar{\pi})]} \lambda(\pi) > 0$. Since $\lambda(\pi) = \mathbf{P}_{\pi,H}(\tau_1 > k)\lambda(\Phi_1^k(\pi))$, one has $\lambda(\Phi_1^k(\pi)) \geq \lambda(\pi)$ for all $\pi \in [\bar{\pi}, \Phi_1(\bar{\pi})]$ and $k \in \mathbf{N}$. Since $\left[\Phi_1^k(\bar{\pi}), \Phi_1^{k+1}(\bar{\pi})\right] \subseteq \Phi_1^k\left([\bar{\pi}, \Phi_1(\bar{\pi})]\right)$ by the intermediate value theorem, it follows that $\min\limits_{\left[\Phi_1^k(\bar{\pi}),\Phi_1^{k+1}(\bar{\pi})\right]} \lambda \geq b$. Thus, $\min_{[\bar{\pi},1)} \lambda \geq b$, as desired. ■

When exchanging the roles of two states, and setting $\tilde{\tau} := \inf\{n : a_n = h\}$, it follows that $\min_{\pi\leq\frac{1}{2}} \mathbf{P}_{\pi,L}(\tilde{\tau} = +\infty) > 0$ for all $\pi \in (0,1)$. We denote this minimum by $\tilde{b}$. We denote by $M := 1 + \max\{k, \tilde{\tau}_k < +\infty\}$ the total number of h-herds.

Corollary 12 *One has $\mathbf{E}_L[M] \leq \dfrac{1}{\tilde{b}}$.*

Proof. By the Markov property of the sequence (π_n), one has $\pi_n \leq \frac{1}{2}$ on the event $\tau_k + 1 = n$, and thus,

$$\mathbf{P}_{\bar{\pi},L}\left(\tilde{\tau}_k < +\infty \mid \mathcal{H}_n\right) = \mathbf{P}_{\pi_n,L}(\tilde{\tau} < +\infty) \leq 1 - \tilde{b}$$

for each k and n: conditional on the history leading to some l-herd, the probability that this l-herd ends in finite time does not exceed $1-\tilde{b}$. The result easily follows (*e.g.* via a coupling argument). ■

We now prove that the conditional expected durations of the successive h-herds are bounded.

Claim 13 *The map Φ_1 is bounded from above above on the interval $x \in (0, \frac{1}{2}]$.*

Proof. It suffices to prove that $\dfrac{\Phi_1(x)}{1-\Phi_1(x)} = \dfrac{x}{1-x} \times \dfrac{1-F_H(1-x)}{1-F_L(1-x)} = \dfrac{x}{1-x} \times \dfrac{\tilde{F}_H(x)}{\tilde{F}_L(x)}$ is bounded from above on $(0, \frac{1}{2}]$. As in the proof of Claim 9, one has

$$\frac{1-\Phi_1(x)}{\Phi_1(x)} = \frac{1-x}{x} \times \frac{x\tilde{F}(x) - \int_0^x \tilde{F}(y)dy}{(1-x)\tilde{F}(x) + \int_0^x \tilde{F}(y)dy}.$$

Since $1-x \geq \frac{1}{2}$ and $(1-x)\tilde{F}(x) + \displaystyle\int_0^x \tilde{F}(y)dy \leq \tilde{F}(x)$, one obtains

$$\frac{1-\Phi_1(x)}{\Phi_1(x)} \geq \frac{1}{2}\frac{x\tilde{F}(x) - \int_0^x \tilde{F}(y)dy}{x\tilde{F}(x)} = \frac{1}{2}\left(1 - \frac{\tilde{F}_M(x)}{\tilde{F}(x)}\right),$$

where $\tilde{F}_M(x) := \dfrac{1}{x}\displaystyle\int_0^x \tilde{F}(y)dy$ is the average value of $\tilde{F}$ on $[0, x]$.

From (**A1**), it follows that $\dfrac{\tilde{F}_M}{\tilde{F}}$ is bounded away from one on $(0, \frac{1}{2}]$, hence the result. ■

Set $d := \sup_{(0,\frac{1}{2}]} \Phi_1$. For all $k \in \mathbf{N}$, one has $\pi_{\tilde{\tau}_k} \leq \frac{1}{2}$, hence $\pi_{\tilde{\tau}_k+1} = \Phi_1(\pi_{\tilde{\tau}_k}) \leq d$ on the event $\tilde{\tau}_k < +\infty$.

Set $\Delta := \max_{\pi \in \left[\frac{1}{2}, d\right]} \mathbf{E}_{\pi,L}[\tau]$. For each k, and by the Markov property of the sequence (π_n) one has

$$\mathbf{E}_{\pi,L}[\tau_{k+1} - \tilde{\tau}_k \mid \mathcal{H}_{\tilde{\tau}_k+1}] = \mathbf{E}_{\pi_{\tilde{\tau}_k+1},L}[\tau] \leq \Delta,$$

on the event $\tilde{\tau}_k < +\infty$.

It follows from Corollary 12 that $\mathbf{E}_L[N] \leq \dfrac{\Delta}{\tilde{b}}$.

C Rates of convergence: Proof of Proposition 4

We start with the first claim. When all guesses are public, one has

$$u_n := \mathbf{P}_L(\tau > n) = \prod_{k=1}^{n} \left(1 - F_L(1 - \pi_k)\right).$$

With $F(p) = p$, one has $F_L(p) = p(2-p)$ and $F_H(p) = p^2$, hence $u_n = \prod_{k=1}^{n} \pi_k^2$. On the other hand, the belief updating equation

$$\frac{\pi_{n+1}}{1-\pi_{n+1}} = \frac{\pi_n}{1-\pi_n} \times \frac{1 - F_H(1-\pi_n)}{1 - F_L(1-\pi_n)}$$

reduces to $\dfrac{\pi_{n+1}}{1-\pi_{n+1}} = \dfrac{2-\pi_n}{1-\pi_n}$, from which it follows that $\left(\dfrac{1}{1-\pi_n}\right)_n$ is an arithmetic sequence, and $\pi_n = 1 - \frac{1}{2n}$ for each $n \geq 1$.

Consequently,

$$u_n = \left(\prod_{k=1}^{n} \left(1 - \frac{1}{2k}\right)\right)^2 = \left(\frac{(2n)!}{2^{2n}(n!)^2}\right)^2.$$

Using Stirling formula, it follows that $u_n \sim_{+\infty} \dfrac{1}{\pi n}$.

Since $\mathbf{E}_L\left[\min(\tau, n)\right] = \sum_{k=1}^{n} \mathbf{P}_L(\tau \geq k)$ and since $\sum \mathbf{P}_L(\tau \geq k)$ diverges, $\mathbf{E}_L\left[\min(\tau, n)\right] \sim_{+\infty} \sum_{k=1}^{n} \dfrac{1}{\pi k}$ and therefore $\mathbf{E}_L\left[\min(\tau, n)\right] \sim_{+\infty} \frac{1}{\pi} \ln n$, as desired.

We turn to the second claim. Recall that the probability $x_n := \mathbf{P}_L(a_n = h)$ of a wrong guess satisfies the recursive equation $x_{n+1} - x_n = -2 \int_0^{x_n} F(p)dp$. With uniform beliefs, this reduces to a special case of the discrete time logistic equation,

$$x_{n+1} = x_n(1 - x_n). \tag{16}$$

Since $x_1 \in (0,1)$, it is obvious from (16) that (x_n) is decreasing and must converge to zero. Since $x_n = x_1 \prod_{k=1}^{n-1}(1 - x_k)$ and since $\lim_n x_n = 0$, the infinite product $\prod_{k=1}^{+\infty}(1 - x_k)$ is zero, hence the series $\sum x_k$ fails to converge.

An easy induction shows that $x_n < \frac{1}{n+1}$ for all $n \geq 2$. Set now $y_n := n x_n$, and observe that

$$y_{n+1} - y_n = x_n\left(1 - (n+1)x_n\right) \geq 0.$$

The sequence (y_n) being non-decreasing with $y_n \leq 1$, it has a positive limit, which we denote by $l > 0$.

On the other hand,

$$y_{n+1} = (n+1)x_{n+1} = y_n \left(1 + \frac{1}{n}\right)\left(1 - \frac{y_n}{n}\right)$$

hence

$$y_{n+1} - y_n = \frac{y_n(1-y_n)}{n} - \frac{y_n^2}{n^2}.$$

Since the sequence (y_n) converges, the series $\sum(y_{n+1} - y_n)$ converges as well, hence $l = 1$.[24] We have thus shown that $x_n \sim_{+\infty} \frac{1}{n}$. This implies that

$$\mathbf{E}_L\left[N_n\right] = \sum_{k=1}^{n} x_k \sim_{+\infty} \sum_{k=1}^{n} \frac{1}{k} \sim_{+\infty} \ln n.$$

This concludes the proof of the second claim.

D Proof of Proposition 5

Recall that $q_n = \mathbf{P}(H \mid s_n)$ is the private belief of agent n. The interim belief of agent n is here the probability $\pi_n = \mathbf{P}(H \mid a)$ assigned by agent n to H, conditional on the random action $\tilde{a}_{n-1}$ observed by agent n being a. We use notations similar to those of Section **??**: $x_n := \mathbf{P}_L(a_n = h)$ is the probability that agent n picks the wrong action and $\lambda_n(a \mid \theta)$ is more generally the probability that agent n picks agent a conditional on the state being θ.

Finally, we denote by $\bar{x}_n := \dfrac{1}{n}\displaystyle\sum_{k=1}^{n} x_k$ the expected frequency of mistakes in the first n rounds.

Lemma 8 *One has*

$$\bar{x}_{n+1} = \bar{x}_n - \frac{2}{n+1}\int_0^{\bar{x}_n} F(t)dt. \tag{17}$$

Proof. Because agent $n+1$ randomly observes one of the previous guesses, the probability that agent $n+1$ observes action a in state θ is

$$\bar{\lambda}_n(a \mid \theta) = \frac{1}{n}\sum_{k=1}^{n} \lambda_k(a \mid \theta),$$

[24] Otherwise, $y_{n+1} - y_n$ would be equivalent to $l(1-l)/n$.

and his interim belief π_{n+1} obeys the updating equation $\dfrac{\pi_{n+1}}{1-\pi_{n+1}} = \dfrac{\bar{\lambda}_n(a \mid H)}{\bar{\lambda}_n(a \mid L)}$. On the other hand, agent $n+1$ picks action h if and only if $q_{n+1} + \pi_{n+1} \geq 1$. Hence, the probability $\lambda_{n+1}(h \mid \theta)$ that agent $n+1$ plays h is

$$\lambda_{n+1}(h \mid \theta) = \sum_a \bar{\lambda}_n(a \mid \theta)\mathbf{P}_\theta(q_{n+1} > 1 - \pi_{n+1}).$$

Elementary algebraic manipulations yield

$$x_{n+1} = (1-\bar{x}_n)F_H(\bar{x}_n) + \bar{x}_n F_H(1-\bar{x}_n),$$

using $x_n = \lambda_n(h \mid L)$.

Equivalently, and using (3),

$$x_{n+1} = \bar{x}_n - 2\int_0^{\bar{x}_n} F(t)dt. \tag{18}$$

Since $\bar{x}_{n+1} = \dfrac{n}{n+1}\bar{x}_n + \dfrac{1}{n+1}x_{n+1}$, the result follows from equation (18). ■

It follows from equation (17) that $\lim_{n\to+\infty} \bar{x}_n = 0$. Indeed, $(\bar{x}_n)$ is decreasing and non-negative, hence is convergent, with limit $l \geq 0$. If the limit l were strictly positive, one would have $\bar{x}_{n+1} - \bar{x}_n \sim_{+\infty} -\dfrac{C_2}{n+1}$, with $C_2 := \displaystyle\int_0^l F(t)dt$. This would imply $\lim_n \bar{x}_n = -\infty$ – this is the desired contradiction.

Using again equation (17), $|x_{n+1} - \bar{x}_n| \leq 2\bar{x}_n F(\bar{x}_n)$, hence $x_{n+1} \sim_{+\infty} \bar{x}_n$, so that the convergence of the series $\sum x_n$ is equivalent to that of $\sum \bar{x}_n$.

From our standing assumption on F, and when possibly lowering α, one has $F(t) \leq \frac{1}{2}(\alpha+1)t^\alpha$ in a neighborhood of zero. Plugging this into (17), we obtain

$$\bar{x}_{n+1} \geq \bar{x}_n - \frac{1}{n+1}\bar{x}_n^{1+\alpha}. \tag{19}$$

Fix some large integer N and define the auxiliary sequence (y_n) by $y_N = \bar{x}_N$ and $y_{n+1} - y_n = -\frac{1}{n+1}y_n^{1+\alpha}$. Since $\bar{x}_n \to 0$, and since the function $y \mapsto y - y^{1+\alpha}$ is increasing over a neighborhood of zero, it easily follows by induction that $\bar{x}_n \geq y_n$ for all $n \geq N$, provided N is chosen large enough.

Hence, it is sufficient to prove that the series $\sum y_n$ is not convergent. The sequence (y_n) is obviously decreasing, and converges to zero.[25] Hence

$$\frac{y_{n+1}}{y_n} = 1 - \frac{1}{n}y_n^\alpha = 1 + o\left(\frac{1}{n}\right),$$

[25] If (y_n) instead had a positive limit l, we would have $y_{n+1} - y_n \leq -\frac{l^\alpha}{n}$ which by summation would imply $\lim y_n = -\infty$.

and it follows from the Raabe-Duhamel criterion that $\sum_{n=N}^{+\infty} y_n = +\infty$.

E Proof of Theorem 7

For $k \geq 1$, we denote by $\Delta_k = d_1 + \cdots + d_k$ the cumulative duration of the first k generations, with $\Delta_0 = 1$, and recall that $\tau := \inf\{n : a_n = \theta\}$ is the identity of the first player who plays a correct guess. Since

$$\mathbf{E}_L[\tau] = \sum_{n=1}^{+\infty} \mathbf{P}_L(\tau \geq n) = \sum_{k=1}^{+\infty} \sum_{n=\Delta_{k-1}+1}^{\Delta_k} \mathbf{P}_L(\tau \geq n),$$

we have the following bounds on $\mathbf{E}_L[\tau]$:

$$\sum_{k=1}^{+\infty} d_k \mathbf{P}_L(\tau > \Delta_k) \leq \mathbf{E}_L[\tau] \leq \sum_{k=1}^{+\infty} d_k \mathbf{P}_L(\tau > \Delta_{k-1}).$$

Since the sequence $(d_k/d_{k-1})_k$ is bounded, the two series $\sum d_k \mathbf{P}_L(\tau > \Delta_k)$ and $\sum d_k \mathbf{P}_L(\tau > \Delta_{k-1})$ simultaneously converge or diverge. Hence $\mathbf{E}_L[\tau] < +\infty$ is equivalent to the series $\sum d_k \mathbf{P}_L(\tau > \Delta_k)$ being convergent.

Abusing notations, we denote by $\pi_i = \mathbf{P}(H \mid a_1 = \cdots = a_{\Delta_{i-1}} = h)$ the interim belief of agents from the i-th, in the specific event where all agents have played h so far. Recall that agent n in the k-th generation chooses $a_n = h$ if and only if the private and interim beliefs are such that $q_n + \pi_k \geq 1$, which occurs with probability $1 - F_L(1 - \pi_k)$ if the state is L.

With uniform private beliefs, one has $F_L(q) = q(2 - q)$ and $F_H(q) = q^2$. Thus, $1 - F_L(1 - \pi_k) = \pi_k^2$, and, since private signals are conditionally independent,

$$\mathbf{P}_L(\tau > \Delta_k) = \prod_{i=1}^{k-1} \pi_i^{2d_i}.$$

On the other hand, Bayes rule leads to the belief updating formula

$$\frac{\pi_{k+1}}{1 - \pi_{k+1}} = \frac{\pi_k}{1 - \pi_k} \times \left(\frac{1 - F_H(1 - \pi_k)}{1 - F_L(1 - \pi_k)} \right)^{d_k} = \frac{\pi_k}{1 - \pi_k} \times \left(\frac{2 - \pi_k}{\pi_k} \right)^{d_k}.$$

Setting $u_k := \frac{1}{2} \frac{\pi_k}{1 - \pi_k}$, we have $\pi_k = 1 - \frac{1}{1 + 2u_k}$, and the updating formula simplifies to

$$u_{n+1} = u_n \left(1 + \frac{1}{u_n} \right)^{d_n}. \tag{20}$$

We prove the result through a series of claims.

Claim 14 *One has $u_{n+1} \geq \Delta_n + 1$ for all n.*

Proof. We use the inequality $(1+x)^\alpha \geq 1 + \alpha x$ (for $\alpha > 1$, $x > 0$), which yields $u_{n+1} \geq u_n + d_n$. ■

Claim 15 *The series $\sum \frac{d_k}{(\Delta_k)^2}$ is convergent.*

Proof. Set $\omega_j := \inf\{k : d_k \geq j\}$ (with $\omega_1 = 1$) and $n_j := \omega_{j+1} - \omega_j$. Thus,

$$\sum_{j=1}^{J-1} \sum_{k=\omega_j}^{\omega_{j+1}-1} \frac{d_k}{(\Delta_k)^2} \leq \sum_{j=1}^{J-1} (j+1) \sum_{k=\omega_j}^{\omega_{j+1}-1} \frac{1}{(\Delta_k)^2}.$$

For $\omega_j \leq k < \omega_{j+1}$, one has

$$\Delta_k \geq \left(\sum_{i=1}^{j-1} i n_i \right) + j(k - \omega_j + 1).$$

When setting $b_j := \sum_{i=1}^{j-1} i n_i$, it follows that

$$\begin{aligned} \sum_{k=\omega_j}^{\omega_{j+1}-1} \frac{1}{(\Delta_k)^2} &\leq \sum_{p=1}^{n_j} \frac{1}{(b_j + jp)^2} = \frac{1}{n_j} \left(\frac{1}{n_j} \sum_{p=1}^{n_j} \frac{1}{(\frac{b_j}{n_j} + j\frac{p}{n_j})^2} \right) \\ &\leq \frac{1}{n_j} \int_0^1 \frac{1}{\left(\frac{b_j}{n_j} + jx\right)^2} dx = \frac{1}{j} \left(\frac{1}{b_j} - \frac{1}{b_{j+1}} \right). \end{aligned}$$

Therefore,

$$\begin{aligned} \sum_{j=1}^{J-1} \sum_{k=\omega_j}^{\omega_{j+1}-1} \frac{d_k}{(\Delta_k)^2} &\leq \sum_{j=1}^{J-1} \frac{1}{j} \left(\frac{1}{b_j} - \frac{1}{b_{j-1}} \right) \\ &= \sum_{j=1}^{J-1} \frac{1}{j b_j} - \sum_{j=2}^{J} \frac{1}{(j-1) b_j} \leq \frac{1}{b_1} \end{aligned}$$

and the result follows. ■

Claim 16 *The series $\sum \frac{d_k}{(u_k)^2}$ is convergent.*

Proof. Thanks to Claim 14 and since $u_1 = \frac{1}{2}$, one has

$$\sum_{k=1}^{\infty} \frac{d_k}{(u_k)^2} \leq 4d_1 + \sum_{k=2}^{\infty} \frac{d_k}{(\Delta_{k-1})^2} \leq 4d_1 + C_2 \sum_{k=1}^{+\infty} \frac{d_k}{(\Delta_k)^2},$$

where C_2 is an upper bound for the sequence $(d_{k+1}/d_k)_k$. ■

Claim 17 *The series* $\sum \frac{d_k}{u_k}$ *is divergent.*

Proof. For each n, the updating formula (20) yields

$$\frac{u_{n+1}}{u_n} = \left(1 + \frac{1}{u_n}\right)^{d_n} \leq e^{d_n/u_n},$$

which, taking products over n, implies

$$2u_{n+1} \leq \exp\left(\sum_{j=1}^{n} \frac{d_i}{u_i}\right).$$

Since $(u_n) \to +\infty$, the divergence of $\sum \frac{d_k}{u_k}$ follows. ■

Claim 18 *The convergence of* $\sum_k d_k \mathbf{P}_L(\tau > \Delta_k)$ *is equivalent to the convergence of* $\sum d_k e^{-\sum_{i=1}^{k-1} d_i/u_i}$.

Proof. We set $a_k := d_k \mathbf{P}_L(\tau > \Delta_k) = d_k e^{\sum_{i=1}^{k-1} 2d_i \ln \pi_i}$. We use the inequality $x - x^2 \leq \ln(1+x) \leq x$ which holds for all $x > -\frac{1}{2}$. Since $(\pi_i) \to 1$, this inequality implies

$$-\frac{1}{1+2u_i} - \frac{1}{(1+2u_i)^2} \leq \ln \pi_i \leq -\frac{1}{1+2u_i},$$

for all i large enough, which yields the slightly simpler bounds

$$-\frac{1}{2u_i} - \frac{1}{(2u_i)^2} \leq \ln \pi_i \leq -\frac{1}{2u_i} + \frac{1}{(2u_i)^2}.$$

Taking sums and exponentials, we obtain

$$d_k e^{-\sum_{k=1}^{i-1} \frac{d_i}{u_i}} e^{-\sum_{k=1}^{i-1} \frac{d_i}{2u_i^2}} \leq a_k \leq d_k e^{-\sum_{k=1}^{i-1} \frac{d_i}{u_i}} e^{+\sum_{k=1}^{i-1} \frac{d_i}{2u_i^2}},$$

and thus,

$$d_k e^{-\sum_{k=1}^{i-1} \frac{d_i}{u_i}} e^{-\frac{1}{2}\Sigma_{u^2}} \leq a_k \leq d_k e^{-\sum_{k=1}^{i-1} \frac{d_i}{u_i}} e^{+\frac{1}{2}\Sigma_{u^2}},$$

where Σ_{u^2} is the infinite sum of the convergent series $\sum \frac{d_k}{(u_k)^2}$. The result follows. ■

We conclude the proof by proving that the series $\sum d_k e^{-\sum_{i=1}^{k-1} d_i/u_i}$ fails to converge. Using the inequalities

$$\frac{1}{u_n} - \frac{1}{(u_n)^2} \leq \ln\left(1 + \frac{1}{u_n}\right) \leq \frac{1}{u_n}$$

and the identity $\frac{u_{n+1}}{u_n} = e^{d_n \ln\left(1+\frac{1}{u_n}\right)}$, one gets

$$e^{d_n/u_n} e^{-d_n/(u_n)^2} \leq \frac{u_{n+1}}{u_n} \leq e^{d_n/u_n},$$

or equivalently,

$$\frac{u_n}{u_{n+1}} e^{-d_n/(u_n)^2} \leq e^{-d_n/u_n} \leq \frac{u_n}{u_{n+1}}.$$

Taking products and multiplying by d_k, we get

$$u_1 \frac{d_k}{u_k} e^{-\Sigma_{u^2}} \leq d_k e^{-\sum_{i=1}^{k-1} d_i/u_i} \leq u_1 \frac{d_k}{u_k}.$$

Hence the divergence of $\sum d_k e^{-\sum_{i=1}^{k-1} d_i/u_i}$ follows from Claim 17.