Sharp Thresholds of the Information Cascade Fragility Under a Mismatched Model

Wasim Hueihel
Tel-Aviv University

Ofer Shayevitz
Tel-Aviv University

Abstract

We analyze a sequential decision making model in which decision makers (or, players) take their decisions based on their own private information as well as the actions of previous decision makers. Such decision making processes often lead to what is known as the *information cascade* or *herding* phenomenon. Specifically, a cascade develops when it seems rational for some players to abandon their own private information and imitate the actions of earlier players. The risk, however, is that if the initial decisions were wrong, then the whole cascade will be wrong. Nonetheless, information cascade are known to be fragile: there exists a sequence of *revealing* probabilities $\{p_\ell\}_{\ell\geq 1}$, such that if with probability p_ℓ player ℓ ignores the decisions of previous players, and rely on his private information only, then wrong cascades can be avoided. Previous related papers which study the fragility of information cascades always assume that the revealing probabilities are known to all players perfectly, which might be unrealistic in practice. Accordingly, in this paper we study a mismatch model where players believe that the revealing probabilities are $\{q_\ell\}_{\ell\in\mathbb{N}}$ when they truly are $\{p_\ell\}_{\ell\in\mathbb{N}}$, and study the effect of this mismatch on information cascades. We consider both adversarial and probabilistic sequential decision making models, and derive closed-form expressions for the optimal learning rates at which the error probability associated with a certain decision maker goes to zero. We prove several novel phase transitions in the behaviour of the asymptotic learning rate.

Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS) 2020, Palermo, Italy. PMLR: Volume 108.

1 INTRODUCTION

There are myriad economic and social scenarios where our decisions are influenced by the actions of others. For example, voters are inclined to vote in favor of what opinion polls predicts will win. Academic researchers choose to work on topics that are of broad and current interest. Fertility decisions, e.g., how many children to have, are known to be influenced by what other people in the same geographical location are doing. Opinions we hold, products we buy, and technologies we use, etc., are all potentially affected by our surroundings. The above is a non-exhaustive list of scenarios where our rational behaviour guides us to follow the actions of others despite the fact that these may contradict our own information. This is exactly the situation where information cascades [Banerjee, 1992, Bikhchandani et al., 1992] develop.

To illustrate the way information cascades evolve we consider the following simple and classical herding experiment, proposed and studied in [Anderson and Holt, 1996, Anderson and Holt, 1997] (see also [David and Jon, 2010, Ch. 16]). In this experiment, we place an urn that contains three marbles in front of a bunch of players. The urn contains either one red marble and two blue marbles (*majority blue*), or, two red marbles and one blue marble (*majority red*). Players do not know whether the urn is majority blue or red, while both urns are equally likely to be chosen. In a successive manner, each player randomly draws a single marble from the urn, memorizes its color, and returns it to the urn, while not showing it to the other players. Then, each player in his turn publicly announce his guess for the urn majority color. The players guesses are based on both their own private draws as well as the actions/announcements of previous players.

Next, we explain how the above experiment evolves. It is clear that the first two players will announce their private signals as their guesses. Indeed, the first player gets to see his own draw only, and thus his best guess for the urn majority color is the color he draw. The second player is aware of that, and thus, together with his own draw he gets to see two independent draws from the urn. Accordingly, if the colors agree, then the player announces this color; otherwise, there is a tie, and in such a case let us assume that the player follows his own draw. Therefore, it is clear again that the second player announces the color of the marble he draw. Continuing, the third player now see three independent draws from the urn and consequently his best decision is the majority color among these draws.

The most important observation here is that the rational guess of any subsequent player may not reflect its own private information. For example, if the first two announced colors were blue, then the third player guess is blue *irrespectively* of the color of the marble he picked. Evidently, due to the fact that his guess will not reveal any information about the urn to any subsequent player, every subsequent player will guess the urn to be majority blue. This is where an *information cascade* developed: while no one is under the impression that every player draw a blue marble, since the first two guesses were blue, future rational guesses must be blue as well. To wit, an information cascade is a sequence of decisions where it is optimal for players to ignore their own private information and imitate the decisions of players ahead of them. The problem with information cascade is that they can be *wrong*! Indeed, if for instance in the above example the urn is majority red, then everyone wrongly announced blue as their guesses. In fact, in the above experiment, it can be shown that with probability $1/5$, a wrong cascade develops, in which case, most players will guess the urn majority color wrongly.

The experiment above illustrates that information cascades can be wrong because they rely on very little actual information–the actions of the first few players can determine the actions of all future players. Nonetheless, this hints that information cascades can also be fundamentally very fragile. Indeed, suppose that in the above experiment two consecutive players ℓ and $\ell+1$ draw red marbles, and they cheat (or, act irrationally) by announcing their marbles despite the fact that a majority-blue cascade was already developed (say, the first two players announced blue). Then, it is clear that the wrong cascade can be broken: player $\ell+2$ sees four *informative* draws (two blues and two reds), and he will announce his own private signal as his guess. The conclusion is that infusion of new information can overturn/brake wrong information cascades even after they have persisted for a long time.

Motivated by the above observation, we consider a simple sequential decision making model in which not all players are rational, but rather certain players act irrationally, by revealing their own private signal and discarding the actions of previous players. More specifically, we focus on a recent model proposed by [Peres et al., 2018], where it is assumed that the ℓ^{th} player is irrational/revealer with some probability p_ℓ, and is rational/Bayesian with complementary probability. While players *do not* know whether other players before them were rational or not, it is assumed in [Peres et al., 2018] nonetheless that players *know* the revealing probabilities $\{p_\ell\}_{\ell\in\mathbb{N}}$. As mentioned in [Peres et al., 2018], this model is prompted by both empirical laboratory experiments [Anderson and Holt, 1997, Huck and Oechssler, 2000, Weizscker, 2010], as well as several theoretical reasons [Bernardo and Welch, 2001]. One of the intriguing questions here, is whether wrong information cascades are broken in the above model? Or, stated differently, do people eventually learn the correct action? As was shown in [Peres et al., 2018], the answer to this question is positive. Namely, it can be shown that there exists a sequence of revealing probabilities $\{p_\ell\}_{\ell\in\mathbb{N}}$ such that learning occurs. In particular, the optimal policy minimizing the probability of error is for the ℓ^{th} player to reveal its private signal with probability $p_\ell = c/\ell$, which in turn implies a learning rate of c'/ℓ, where c and c' are explicit constants. While these results are neat, as mentioned in [Peres et al., 2018] they rely heavily on the assumption that players are fully coordinated, i.e., they know the revealing probabilities exactly, which might be unrealistic in practice. This sets precisely the main goal of our paper: we aim to understand how this coordination affects information cascades. To this end, we introduce a mismatch model where players believe that the revealing probabilities are $\{q_\ell\}_{\ell\in\mathbb{N}}$ when they truly are $\{p_\ell\}_{\ell\in\mathbb{N}}$. We are interested in understanding whether asymptotic learning occur in this case? and if so, under what conditions and at what learning rate? In particular, what is the cost of this mismatch?

The above mismatch model might be relevant in many real-world applications, such as, rumor spreading over social networks, online movie rating, etc., where it is well-documented that human behavior is sometimes irrational (e.g., [Kahneman and Tversky, 1973]). Furthermore, it is well-known in the social learning literature that a fully rational model often places unreasonable computational demands on Bayesian players (e.g., [Mossel and Tamuz, 2017]), hence understanding the impact of simpler more efficient strategies is desirable. As we explain in the paper, this situation can be partially captured by our model, since our mismatch framework allow for a family of sub-optimal strategies parameterized by the mismatch sequence $\{q_\ell\}_{\ell\in\mathbb{N}}$.

Main Contributions. The main contributions of this paper are as follows. We start by formulating a general adversarial/worst-case model where the placement of irrational players is *arbitrary*, and not being governed by any probabilistic/statistical rule. We show that this kind of model is in fact too stringent leading to trivial results. Combined with [Peres et al., 2018], this fact motivates us to study the more flexible probabilistic model described above. For this model we characterize the asymptotic learning rate exactly, which turns out to exhibit several novel interesting phase transitions. Specifically, we first show that for either too "optimistic" or "pessimistic" assumptions, i.e., $q_t = o(p_t)$ or $q_t = \omega(p_t)$, asymptotic learning *does not* occur, namely, the error probability is high, and the total number of wrong decisions is significant. We then consider the case where $q_t = \Theta(p_t)$, and show that asymptotic learning occurs, but at a reduced rate which loosely speaking depends on the ratio between the ℓ_1-norms of the matched and mismatched revealing probabilities. This is true, as long as the magnitude of this ratio is moderate laying between two thresholds, otherwise, asymptotic learning does not occur!

Related work. Sequential decision making has been studied in various areas including politics, economics and computer science. In particular, the case where all players are Bayesian (i.e., $p_\ell = 0, \forall \ell$) was considered in [David and Jon, 2010, Banerjee, 1992, Bikhchandani et al., 1992]. Notably, [Bikhchandani et al., 1992] gives many interesting real-world examples from diverse fields, where information cascades develop and are fragile. Similarly to [Peres et al., 2018], our paper provides a more detailed theoretical study of the fragility phenomenon. In practice, however, it is well documented, that human behavior deviates from rationality, and rather irrational decisions are made often, see, e.g., [Kahneman and Tversky, 1973, Huck and Oechssler, 2000] and [Weizscker, 2010]. Indeed, several laboratory experiments in [Anderson and Holt, 1996] illustrate that many individuals act irrationally, by ignoring the actions of other individuals and relying mainly on their own private information. Our model captures this empirically observed behavioral phenomenon.

In [Bernardo and Welch, 2001], a model which combines both rational and partially irrational (who put more weight on their private signal) types of individuals, as in our model, was studied. It is assumed, however, that players know which of the previous players were revealers. Using simulations it was suggested that learning is achievable only when completely irrational individuals exist, and that their fraction should vanish. These observations were rigorously proved in [Peres et al., 2018], showing that the optimal number of revealers is logarithmic in the size of the group. Our results show that in many cases learning is possible even if there is a mismatch. Recently, [Cheng et al., 2018] also considered a model with irrational players, but assume that each agent knows which agents were revealers. While they show that asymptotic learning occurs, the optimal learning rate was not characterized, which is another contribution of our paper when there is a mismatch. Sequential decision models with unbounded private signals (e.g., Gaussian) were studied in [Smith and Srensen, 1996, Hann-Caruthers et al., 2017]. We also mention the study of sequential decision making models over social random graphs, e.g., [Acemoglu et al., 2011, Anunrojwong and Sothanaphan, 2018]. Finally, we mention another somewhat related literature which studies the situation where agents take repeated decisions (rather than just a single decision) based on their own private information and the actions of others, e.g., [Harel et al., 2018], and one is interested in understanding whether all decision makers learn the correct state eventually, and if so at what speed.

2 PROBLEM SETUP

In this section, we present our model and formulate the problem of interest. To convey neatly the main ideas of this paper, we focus on a simple setting of the information cascade model. Nonetheless, several generalizations listed at the end of this paper, can be derived using the same techniques used in this paper. Let $\theta \in \{1,2\}$ denote the state of the world, chosen uniformly at random. At times $t = 1, 2, 3, \ldots$, players one by one try to guess θ, relying on their own *private* signals, as well as the *global actions* (guesses) of players who played before them.

We next describe the way private signals are formed. There is an urn that contains two types of marbles: Type–I marbles are blue, and Type–II are red. There are two hypotheses depending on the value of θ. Specifically, given θ, there are α marbles of type θ in the urn, and β marbles of the other type, where we assume that $\alpha > \beta > 0$. We conduct the following experiment: each player draws a single marble from the urn and replace it. The color he draw is his private signal. Thus, the private signals denoted by $X_1, X_2, \ldots$ are i.i.d., and:

$$\mathbb{P}_1(X_t = 1) = \frac{\alpha}{\alpha+\beta}; \qquad \mathbb{P}_1(X_t = 2) = \frac{\beta}{\alpha+\beta},$$
$$\mathbb{P}_2(X_t = 1) = \frac{\beta}{\alpha+\beta}; \qquad \mathbb{P}_2(X_t = 2) = \frac{\alpha}{\alpha+\beta},$$

where $\mathbb{P}_i(\cdot) \triangleq \mathbb{P}(\cdot|\theta = i)$, and $t \in \mathbb{N}$. It is clear that the ratio α/β, rather than the actual values of α and β, is important. We denote this ratio by $\gamma \triangleq \alpha/\beta$. Note that an alternative equivalent description of the above setting is that at each time t, the t^{th} player private signal is $X_t = \theta$ with some probability, and $X_t = 3 - \theta$, with the complementary probability. With θ being latent, the players goal is to guess the type of the marbles in the urn correctly. We denote by Z_t the guess of the t^{th} player. As mentioned in the introduction, if all players act rationally by announcing their majority decision, then with positive probability (wrong) information cascade will occur. Accordingly, to break this wrong information cascade, we assume that each player, can operate in either one of the following two modes:

- A revealer/irrational player, whose guess is simply its private signal, i.e., $Z_t = X_t$.
- A Bayesian/rational player, whose guess is its best estimate given his private signal, previous guesses by other players, and any additional auxiliary information.

Given the two modes above, it is left to specify the *way* players are "chosen" to be in either one of the above modes, which is the last aspect of our model. We start with a general *adversarial/worst-case* machinery, which turns out to be too stringent, and as a consequence leads to trivial results. Nonetheless, this model will motivate our second *probabilistic* setting, which is the focus of this paper.

In the adversarial setting, we assume that out of a total of $\mathsf{N} \in \mathbb{N}$ players $\mathsf{V}_\mathsf{N} \in \mathbb{N}$ are irrational, and are chosen in an *arbitrary* manner. We let Π_N be the set of these irrational players. Players *do not* know whether previous players were rational or not, but they do have the value of V_N in advance. We define the probability of incorrect decision of the t^{th} player, assuming that he is rational, as follows,

$$\mathsf{P}_{\mathsf{adv},t}(\mathsf{V}_\mathsf{N}) \triangleq \inf_{\hat{\theta}_t \in \hat{\Theta}} \sup_{\Pi_\mathsf{N} \subset [\mathsf{N}]:\ |\Pi_\mathsf{N}| = \mathsf{V}_\mathsf{N}} \mathsf{P}_{e,t}(\hat{\theta}_t, \Pi_\mathsf{N}), \quad (1)$$

and

$$\mathsf{P}_{e,t}(\hat{\theta}_t, \Pi_\mathsf{N}) \triangleq \mathbb{P}\left[\hat{\theta}_t(Z_1^{t-1}, X_t) \neq \theta \middle| \mathsf{Rev}_\mathsf{N} = \Pi_\mathsf{N}\right], \quad (2)$$

where Z_1^{t-1} is a shorthand notation for the sequence $(Z_1, Z_2, \ldots, Z_{t-1})$, Rev_N designates the set of revealers, the maximum is taken over all possible sets of revealers of size V_N, and the minimum is over the set of all possible estimators $\hat{\Theta}$, which are the Boolean maps $\{1,2\}^t \to \{1,2\}$. To wit, we look at the worst-case error probability over all possible choices of V_N irrational players out of N players. An alternative objective is to minimize the expected total number of errors, that is,

$$\mathsf{TE}(\mathsf{V}_\mathsf{N}) \triangleq \inf_{\hat{\theta} \in \hat{\Theta}} \sup_{\Pi_\mathsf{N} \subset [\mathsf{N}]:\ |\Pi_\mathsf{N}| = \mathsf{V}_\mathsf{N}} \sum_{t=1}^{\mathsf{N}} \mathsf{P}_{e,t}(\hat{\theta}_t, \Pi_\mathsf{N}). \quad (3)$$

For both objectives, it is a-priori unclear what is the optimal strategy $\hat{\theta}_t$. One option, which is simple and widely used in practice, is to assume that each rational player guesses the value of θ using the majority decision, denoted by $\mathsf{Maj}(Z_1^{t-1}, X_t)$. A more complicated approach is to minimize over the Boolean functions used by the rational players, i.e., to solve a minimax problem. We would like to find the asymptotic behaviour of $\mathsf{P}_{\mathsf{adv},t}(\mathsf{V}_\mathsf{N})$ and $\mathsf{TE}(\mathsf{V}_\mathsf{N})$, as a function of V_N and t. In particular, it is interesting to understand the structure of the worst-case choice of the set of the irrational players Π_N. It turns out that the above model/objective, however, is too stringent. Specifically, we show in Appendix C.4 that the error probability in (1) associated with any estimator is lower bounded by a constant, and accordingly, the total number of errors in (3) is proportional to the number of players N. This implies that there are no guessing strategies that are *robust* against an arbitrary adversarial revealers assignment. Therefore, a more flexible model is needed.

To this end, we consider the probabilistic setting introduced in [Peres et al., 2018]. Here, we assume that the t^{th} player is irrational with probability p_t, independently of the other players. Accordingly, this means that if player t is irrational, then $Z_t = X_t$, while if he is rational, then $Z_t = \hat{\theta}_t(Z_1, \ldots, Z_{t-1}, X_t)$, where $\hat{\theta}$ is a certain estimator for θ. The main important ingredient of our model is that we assume that players are completely oblivious to whether previous players were revealers or not. To wit, contrary to previous related works (e.g., [Bernardo and Welch, 2001, Cheng et al., 2018, Peres et al., 2018]), we assume that revealers neither know the exact positioning of revealers, nor the underlying probabilistic law of their placements. Instead, players assume that other players can be revealers with probabilities $\mathcal{Q} \equiv \{q_t\}_{t=1}^\infty$, which might be different than the actual underlying probabilities $\mathcal{P} \equiv \{p_t\}_{t=1}^\infty$. In that case, we say that there is a mismatch. Thus, estimators $\hat{\theta}_t$ might be in fact a function of $\mathcal{Q}$ as well. The matched case case where $\mathcal{P} = \mathcal{Q}$ was considered in [Peres et al., 2018].

Whenever a player is rational we assume that he tries to do his best in guessing θ under the knowledge of $\mathcal{Q}$. Namely, a rational player employs the (mismatched) maximum a posteriori probability (MAP) estimator, which is simply the MAP estimator, but with $\mathcal{P}$ replaced by $\mathcal{Q}$. Specifically, for $i \in \{1,2\}$ and $t \geq 1$, define the distributions:

$$W_i^t(z_1^t) \triangleq \mathbb{P}_i\left(Z_1^t = z_1^t\right), \tag{4}$$
$$H_i^t(z_1^{t-1}, x_t) \triangleq \mathbb{P}_i\left(Z_1^{t-1} = z_1^{t-1}, X_t = x_t\right), \tag{5}$$

and the corresponding likelihood r.v.s by

$$\mathsf{L}_i^t \triangleq W_i^t(Z_1^t), \tag{6}$$
$$\mathsf{D}_i^t \triangleq H_i^t(Z_1^{t-1}, X_t), \tag{7}$$

with $\mathsf{L}_i^0 = \mathsf{D}_i^0 = 1$, for $i \in \{1,2\}$. Note that player t can compute the likelihoods D_i^t, for $i = 1,2$. The likelihoods $\mathsf{L}_i^0 = \mathsf{D}_i^0 = 1$, for $i \in \{1,2\}$, can be computed by a "genie" who observes the decisions of the first t players. The above probabilities and likelihoods are certain quite complicated functions of $\mathcal{P}$. We let $\hat{W}_i^t$ and $\hat{H}_i^t$ be the corresponding probabilities with $\mathcal{P}$ replaced by $\mathcal{Q}$. We also define $\hat{\mathsf{L}}_i^t$ and $\hat{\mathsf{D}}_i^t$ in the same way, but with W_i^t and H_i^t replaced by $\hat{W}_i^t$ and $\hat{H}_i^t$, respectively. With these definitions, the mismatched MAP estimate, denoted by $Z_t = \mathsf{MAP}_{\mathcal{Q}}(Z_1^{t-1}, X_t)$, is:

$$\mathsf{MAP}_{\mathcal{Q}}(Z_1^{t-1}, X_t) \triangleq \begin{cases} 1, & \hat{\mathsf{D}}_1^t > \hat{\mathsf{D}}_2^t, \\ 2, & \hat{\mathsf{D}}_1^t < \hat{\mathsf{D}}_2^t, \\ X_t, & \hat{\mathsf{D}}_1^t = \hat{\mathsf{D}}_2^t. \end{cases} \tag{8}$$

Thus, the mismatch aspect in our setup lies in the fact that the t^{th} player estimator function depends on $\mathcal{Q}$, and more importantly, independent of $\mathcal{P}$. We mention here that the above estimator in fact provides a family of possible estimators indexed by $\mathcal{Q}$. For instance, $\mathcal{Q} \equiv 1$ corresponds to majority decisions, for which rational players guesses are simply the majority color of their own private and previously announced signals.

Players try to guess the urn majority color. We define the probability of incorrect guess by the t^{th} player, as follows,

$$\mathsf{P}_{e,t}(\mathcal{P}, \mathcal{Q}) \triangleq \mathbb{P}\left(Z_t \neq \theta\right). \tag{9}$$

Our goal is to understand the asymptotic learning rate at which the above error probability decays to zero as a function of t (i.e., learning occurs). To motivate and define our objective precisely, we recall the following recent result which deals with the matched case where $\mathcal{P} = \mathcal{Q}$.

Theorem 1 [Peres et al., 2018, Theorem 1.1] Let

$$\kappa(\gamma) \triangleq \left[1 + \frac{\gamma - 1}{\log \gamma}\left(\log \frac{\gamma - 1}{\log \gamma} - 1\right)\right]^{-1}. \tag{10}$$

Then,

$$\inf_{\mathcal{P}} \limsup_{t \to \infty} t \cdot \mathsf{P}_{e,t}(\mathcal{P}, \mathcal{P}) = \kappa(\gamma). \tag{11}$$

Moreover, one can be arbitrarily close to the optimum by taking,

$$p_t^\star = (1 + \varepsilon) \cdot \frac{(1+\gamma)\kappa(\gamma)}{t} \wedge 1, \tag{12}$$

for $t \geq 1$ and arbitrary $\epsilon > 0$.

Theorem 1 states that the optimal learning rate is $\Theta(1/t)$, and furthermore provides the exact leading constant in (10). To achieve this optimal learning rate, the revealing probabilities should also decay as $\Theta(1/t)$. The intuitive reasoning behind these findings can be found in [Peres et al., 2018, Sec. 1.2]. With this result in mind, in the mismatch case where $\mathcal{Q} \neq \mathcal{P}$, we focus on the following scenario. We assume that $\mathcal{P} = \mathcal{P}^\star$, where $\mathcal{P}^\star$ is defined in (12). In other words, the underlying revealing probabilities sequence is the *optimal* one, while players assume a (possibly) *different* sequence of revealing probabilities $\mathcal{Q}$. For simplicity of demonstration, we opted to focus on this special case since we found it to be the most natural one. Nonetheless, our techniques apply also for the more general case where $\mathcal{P} \neq \mathcal{P}^\star$, and at the end of the following section we cover with this case too. It is then interesting to understand whether such a mismatch has any effect whatsoever on the achieved learning rate. In particular, does asymptotic learning always occur? or, perhaps there is a sequence of revealing probabilities $\mathcal{Q}$ for which learning is impossible. To answer these questions, we aim to characterize the *polynomial learning rate*, defined as follows,

$$\mathsf{E}(\mathcal{P}, \mathcal{Q}) \triangleq \liminf_{t \to \infty} -\frac{\log \mathsf{P}_{e,t}(\mathcal{P}, \mathcal{Q})}{\log t}. \tag{13}$$

To lower bound $\mathsf{E}(\mathcal{P}, \mathcal{Q})$ we upper bound the error probability $\mathsf{P}_{e,t}(\mathcal{P}, \mathcal{Q})$, which in turn can be used to upper bound the expected total number of errors:

$$\mathsf{TE}_t \triangleq \mathbb{E}\left[\sum_{\ell=1}^{t} \mathbb{1}\left[Z_\ell \neq \theta\right]\right]. \tag{14}$$

Accordingly, a positive polynomial decaying learning rate implies that $\mathsf{TE}_t = o(t)$, i.e., the number of errors is negligible compared to the total number of players participated thus far. It is clear that Theorem 1 implies that $\mathsf{E}(\mathcal{P}^\star, \mathcal{P}^\star) = 1$, and in fact that $0 \leq \mathsf{E}(\mathcal{P}^\star, \mathcal{Q}) \leq 1$, for any $\mathcal{Q}$. We would like to understand when it is possible or impossible to obtain a positive polynomial decaying learning rate, i.e., when $\mathsf{E}(\mathcal{P}^\star, \mathcal{Q}) > 0$. Note that an interesting question is to characterize the specific constant in front of the polynomial decaying term by evaluating $\lim_{t\to\infty} t^{\mathsf{E}(\mathcal{P},\mathcal{Q})} \cdot \mathsf{P}_{e,t}(\mathcal{P}, \mathcal{Q})$, which we leave as an open question for future research. Finally, note that Theorem 1 gives a simple lower bound on $\mathsf{P}_{e,t}(\mathcal{P}^\star, \mathcal{Q})$ because of the trivial inequality $\mathsf{P}_{e,t}(\mathcal{P}^\star, \mathcal{Q}) \geq \mathsf{P}_{e,t}(\mathcal{P}^\star, \mathcal{P}^\star)$. Indeed, for *each* player seeking to minimize the error probability, its best action is to output the (matched) MAP decision. It is interesting to note that this observation also follows from

$$\mathsf{P}_{e,t}(\mathcal{P}^\star,\mathcal{Q}) \geq \frac{\beta}{\alpha+\beta}\cdot p_t^\star = \frac{\kappa(\gamma)}{t}, \tag{15}$$

for t large enough, and the first inequality follows because $\frac{\beta}{\alpha+\beta}p_t^\star$ is the error probability when player t acts on its private information only (i.e., revealer), while we ignore the error resulted when player t is rational. This lower bound, however, is not tight as we show in the following section. For the rest of this paper, we let $\mathcal{P}_t \triangleq (p_1,p_2,\ldots,p_t)$, and $\|\cdot\|$ denotes the ℓ_1-norm.

3 MAIN RESULTS

In this section, we study the probabilistic setting described in the previous section. According to Theorem 1, to achieve the optimal learning rate the revealing probabilities should decay as $\Theta(t^{-1})$. Note that it is clear that the revealing probabilities cannot decay to zero too quickly. Indeed, if for example $\|\mathcal{P}\|<\infty$, then by the Borel-Cantelli lemma there will be only a finite number of revealers almost surely, which is equivalent to the situation where no revealers exist. This in turn implies a non-vanishing error probability. The situation is somewhat similar when the revealing probabilities decay to zero too slowly. Intuitively, in this case, it can be shown that the error probability is dominated by the probability that the t^{th} player is a revealer and announce a wrong decision, namely, $\frac{\beta}{\alpha+\beta}p_t$. Therefore, if for example $p_t=\Theta(t^{-c})$, for some $c\in(0,1)$, then $\mathsf{E}(\mathcal{P},\mathcal{P}) = c < \mathsf{E}(\mathcal{P}^\star,\mathcal{P}^\star)$.

The optimal scaling of the revealing probabilities suggests that with high probability there should be $\|\mathcal{P}_t\|\sim\log t$ revealers, as $t\to\infty$. Accordingly, in terms of mismatch, it makes sense that "wise" players will assume that the revealing probabilities decay at the same order, but perhaps with a different constant in front, e.g., $q_t=\rho\cdot p_t$, and $\rho\neq 1$. Nonetheless, as mentioned above, rather than modeling the imperfect knowledge of players about the revealing probabilities, our mismatched MAP can also model the situation where players intentionally employ suboptimal strategies (e.g., in order to reduce computational complexity), such as when $\mathcal{Q}\equiv 1$, which results in majority decisions. While it is clear from the above that in the matched case it is strictly worse to assume that the revealing probabilities $\mathcal{P}$ decay to zero too quickly/slowly, it is a-priori unclear if this is true for $\mathcal{Q}$ as well. The following result shows that if the *assumed* revealing probabilities $\mathcal{Q}$ are either too small or too large, then asymptotic learning does *not* occur at all! In particular, $\mathsf{E}(\mathcal{P}^\star,\mathcal{Q})=0$. We have the following result.

Theorem 2 (Too quick/slow) For any sequence of mismatched revealing probabilities $\mathcal{Q}$ such that,

$$\liminf_{t\to\infty}\frac{\|\mathcal{Q}_t\|}{\log t}=0 \quad \text{or} \quad \limsup_{t\to\infty}\frac{\|\mathcal{Q}_t\|}{\log t}=\infty, \tag{16}$$

we have $\mathsf{E}(\mathcal{P}^\star,\mathcal{Q})=0$.

Theorem 2 implies, for example, that if $q_t=o(t^{-1})$ or $q_t=\omega(t^{-1})$, then the error probability cannot decay to zero polynomially fast. In fact, our proof gives a general lower bound on the probability of error which implies for instance that when $\|\mathcal{Q}\|<\infty$, then the probability of error is lower bounded by a constant. Indeed, as will be seen in the proofs, to analyze the error probability one needs to track the dynamics of the likelihood ratio $\mathsf{R}_t\triangleq\frac{\mathsf{L}_1^t}{\mathsf{L}_2^t}$. In particular, $\mathsf{R}_t<\beta/\alpha$ ($\mathsf{R}_t>\alpha/\beta$) implies that the t^{th} player MAP decision is "2" ("1"). Accordingly, given that $\theta=1$, the main observation in the proof of Theorem 2, is based on the realization that when $\|\mathcal{Q}_t\|=O(1)$, even in the worst-case scenario where the majority of the decisions before player t were $\hat{\theta}_i=1$, only a *finite* number of wrong decisions (namely, $\hat{\theta}_i=2$) suffice to mislead player t and output $\hat{\theta}_t=2$. This happens to be the case because of the fact that the likelihood ratio depends on $\mathcal{Q}$ only through $\|\mathcal{Q}\|$, and therefore, it cannot diverge. The intuitive explanation for the obtained result when $\|\mathcal{Q}_t\|\gg\log t$ is given after Theorem 3.

We next consider the more interesting case where $q_t=\rho\cdot p_t$, for $t\in\mathbb{N}$ and $\rho\in\mathbb{R}_+$, or, more generally, $\|\mathcal{Q}_t\|/\|\mathcal{P}_t^\star\|\to\rho$, as $t\to\infty$. To present our main result we establish first some notation. Let

$$\rho_0\triangleq\frac{\log\gamma}{\gamma-1}, \tag{17}$$

and $\rho_1\triangleq\gamma\cdot\rho_0$. Also, define

$$\delta(\gamma,\rho)\triangleq\frac{\gamma\log\gamma-\rho(\gamma-1)\left[1+\log\frac{\gamma\log\gamma}{\rho(\gamma-1)}\right]}{(1+\gamma)\log\gamma}. \tag{18}$$

Theorem 3 (Multiplicative mismatch) For any sequence of mismatched revealing probabilities $\mathcal{Q}$ such that,

$$\limsup_{t\to\infty}\frac{\|\mathcal{Q}_t\|}{\|\mathcal{P}_t^\star\|}=\rho\in\mathbb{R}_+, \tag{19}$$

we have:

- If $\rho\leq\rho_0$ or $\rho\geq\rho_1$,

$$\mathsf{E}(\mathcal{P}^\star,\mathcal{Q})=0.$$

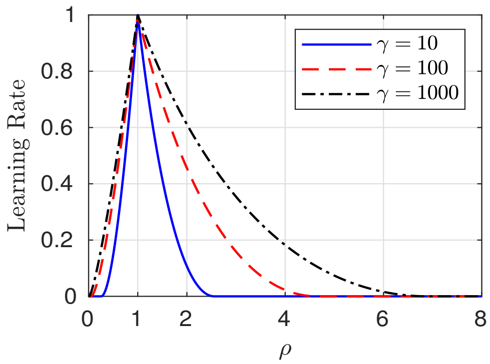

Figure 1: The learning rate $\mathsf{E}(\mathcal{P}^\star, \mathcal{Q})$ in Theorem 3 as a function of ρ, for different values of γ. The maximal learning rate is achieved at $\rho = 1$ for any γ, as expected. Also, as γ increases the range of ρ's for which the learning rate is positive increases.

- If $\rho_0 \leq \rho \leq 1$,

$$\mathsf{E}(\mathcal{P}^\star, \mathcal{Q}) = (1+\gamma)\left[\delta(\gamma,\rho) - \frac{\gamma-1}{\gamma+1}(1-\rho)\right]\kappa(\gamma).$$

- If $1 \leq \rho \leq \rho_1$,

$$\mathsf{E}(\mathcal{P}^\star, \mathcal{Q}) = (1+\gamma)\delta(\gamma,\rho)\kappa(\gamma).$$

It can be checked that $\mathsf{E}(\mathcal{P}^\star, \mathcal{Q})$ is continuous in (ρ, γ). Theorem 3 suggests a clear phase transition in the behaviour of the learning rate (see Fig. 1 for a numerical illustration). To wit, even when players assume that the revealing probabilities decay at the same order of the optimal revealing assignments, albeit with a different constant, there are regimes where the learning rate is zero (i.e., when $\rho \leq \rho_0$ and $\rho \geq \rho_1$). We next give a heuristic explanation as to why the learning rate is zero in these regimes. We start with the case where $\rho \geq \rho_1$. First, note that there are two main sources for wrong action: 1) the t^{th} player is irrational, which happens with probability $p_t^\star$, and his draw is of minority color type, or, 2) the t^{th} player is rational, which happens with probability $1-p_t^\star$, but his mismatched MAP estimate is wrong. Therefore, it is clear that the error probability can be lower bounded by $\mathsf{P}_{e,t}(\mathcal{P}^\star, \mathcal{Q}) \geq (1-p_t^\star)\cdot \mathbb{P}_1\left(\mathsf{MAP}_{\mathcal{Q}}(Z_1^{t-1}, X_t) = 2\right)$. We show in the proof that the MAP error probability can be further lower bounded by the probability that the likelihood ratio $\mathsf{R}_{t-1} \triangleq \mathsf{L}_1^{t-1}/\mathsf{L}_2^{t-1}$, at time $t-1$, is less than β/α, namely, we have $\mathsf{P}_{e,t}(\mathcal{P}^\star, \mathcal{Q}) \geq (1-p_t^\star)\cdot\mathbb{P}_1\left(\mathsf{R}_{t-1} < \beta/\alpha\right)$. To further lower bound the probability term on the r.h.s. of the above inequality, we construct a particular trajectory that ensures that the likelihood ratio R_{t-1} stays always below the threshold β/α. The main observation is then that when the likelihood ratio is bellow β/α, the corresponding log-likelihood ratio process $\log \mathsf{R}_t$ as a function of t, performs a random walk with a *downward* drift *when* $\rho \geq \rho_1$, and thus intuitively the probability that the likelihood ratio will stay bellow some fixed value is high. Technically speaking, we show that this random walk is a supermartingale, and by using well-known tail probability bounds for such processes, we show that $\mathbb{P}_1\left(\mathsf{R}_{t-1} < \beta/\alpha\right)$ can not decay too fast. Establishing this result, and proving a certain monotonicity property of the error probability w.r.t. $\mathcal{Q}$, we show that $\mathsf{E}(\mathcal{P}^\star, \mathcal{Q}) = 0$ when $\|\mathcal{Q}_t\| \gg \log t$ as well, as claimed in Theorem 2. On other hand, when $\rho < \rho_1$, the previously mentioned random walk has an upward drift, and thus, the probability that this walk stays always bellow β/α is intuitively small, and in fact, decays at the polynomial rate given in Theorem 3.

The reason for the learning rate being zero for $\rho \leq \rho_0$ is similar. Contrary to the case where $\rho \geq \rho_1$, in this regime, it can be shown that bellow $\log(\beta/\alpha)$ the log-likelihood ratio process performs a random walk with an upward drift, and thus the approach used before for lower bounding the error probability is not going to work. It turns out, however, that above $\log(\alpha/\beta)$, the log-likelihood ratio process has a downward drift. Moreover, we can show that in this case the walk cannot diverge, or, more precisely, go beyond a certain *finite* value. Among other things, this implies that it takes only a finite number of timestamps to drive the log-likelihood ratio bellow $\log(\beta/\alpha)$, which in turn entails that the error probability is finite as well. Specifically, to lower bound the error probability, we show that it is suffice to look at all trajectories for which the private signals of the last $t^\star \in \mathbb{N}$ revealers are opposite to the majority (e.g., $X_i = 2$, for if $\theta = 1$). Indeed, this way we can assure that the likelihood ratio decreases by a multiplicative factor of β/α. Accordingly, since we argue that the maximal value that likelihood ratio can attain is finite, it is clear that there exists a finite value of $t^\star$ which will drive R_t bellow β/α (note that $t^\star$ revealers decrease the likelihood by an exponential factor of $(\beta/\alpha)^{t^\star}$). Finally, when $\rho_0 \leq \rho \leq 1$, the random walk has now an upward drift, and thus, the probability that it will go bellow β/α is small, and in fact, decaying at the polynomial rate given in Theorem 3.

The above results characterize $\mathsf{E}(\mathcal{P}^\star, \mathcal{Q})$. As mentioned in the previous section, the same techniques exactly can be used to derive the learning rate $\mathsf{E}(\mathcal{P}, \mathcal{Q})$ for any $\mathcal{P}$, and we present our main findings bellow. Proof sketches can be found in Appendix C. First, as was mentioned at the beginning of this section whenever $\mathcal{P}$ is such that $\|\mathcal{P}\| < \infty$, then Borel-Cantelli lemma implies that there will be only a finite number of revealers almost surely, which is equiv-

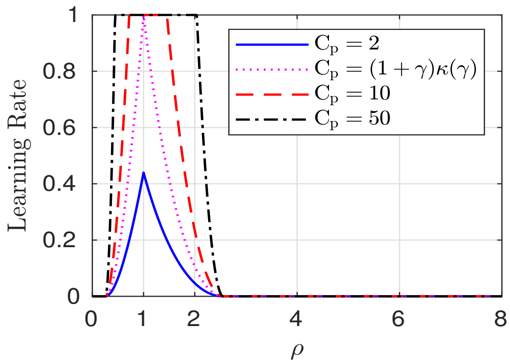

Figure 2: The learning rate $\mathsf{E}(\mathcal{P},\mathcal{Q})$ in Theorem 5 as a function of ρ, for different values of $\mathsf{C_p}$, and $\gamma = 10$.

alent to the situation where no revealers exist. This in turn implies a non-vanishing trivial error probability, and asymptotic learning does not occur. In fact, if $\|\mathcal{P}_t\| = o(\log t)$, then $\mathsf{E}(\mathcal{P},\mathcal{Q}) = 0$, as we show in Appendix C.

Theorem 4 For any sequence of revealing probabilities $\mathcal{P}$ such that $\|\mathcal{P}_t\| = o(\log t)$, and any sequence of mismatched revealing probabilities $\mathcal{Q}$, it holds that $\mathsf{E}(\mathcal{P},\mathcal{Q}) = 0$.

Next, we consider the case where $\|\mathcal{P}_t\| / \log t \to \mathsf{C_p}$, as $t \to \infty$, which happens to be the case, for example, when $p_t = \mathsf{C_p}/t \wedge 1$, for some $\mathsf{C_p} \in \mathbb{R}_+$. Recall the definitions of ρ_0, ρ_1, and $\delta(\gamma,\rho)$ in (17)–(18). We have the following result.

Theorem 5 For any sequence of revealing probabilities $\mathcal{P}$ such that $\|\mathcal{P}_t\| / \log t \to \mathsf{C_p}$, for some $\mathsf{C_p} \in \mathbb{R}_+$, and mismatched revealing probabilities $\mathcal{Q}$ such that,

$$\limsup_{t\to\infty} \frac{\|\mathcal{Q}_t\|}{\|\mathcal{P}_t\|} = \rho \in \mathbb{R}_+, \tag{20}$$

we have:

- If $\rho \le \rho_0$ or $\rho \ge \rho_1$, then $\mathsf{E}(\mathcal{P},\mathcal{Q}) = 0$.
- If $\rho_0 \le \rho \le 1$,
$$\mathsf{E}(\mathcal{P},\mathcal{Q}) = 1 \wedge \left[\mathsf{C_p} \cdot \left(\delta(\gamma,\rho) - \frac{\gamma-1}{\gamma+1}(1-\rho)\right)\right].$$
- If $1 \le \rho \le \rho_1$, then $\mathsf{E}(\mathcal{P},\mathcal{Q}) = 1 \wedge [\mathsf{C_p} \cdot \delta(\gamma,\rho)]$.

From Theorem 5 it can be seen that conceptually the learning rate behaves similarly to the learning rate when $\mathcal{P} = \mathcal{P}^\star$. In particular, the learning exhibits the same phase transitions as in Theorem 3. Fig. 2 presents a numerical calculation of the rate in Theorem 5, for several values of $\mathsf{C_p}$. Note that for $\rho = 1$, when $\mathsf{C_p} < (1+\gamma)\kappa(\gamma)$, we get that $\mathsf{E}(\mathcal{P},\mathcal{Q}) < 1$, while, for any $\mathsf{C_p} \ge (1+\gamma)\kappa(\gamma)$, we get that $\mathsf{E}(\mathcal{P},\mathcal{Q}) = 1$. This might seem counterintuitive because Theorem 1 claims that $\mathsf{C_p} = (1+\gamma)\kappa(\gamma)$ is the optimal value minimizing the error probability, while the above suggests that any value $\mathsf{C_p} \ge (1+\gamma)\kappa(\gamma)$ suffices. Note, however, that while indeed any value $\mathsf{C_p} \ge (1+\gamma)\kappa(\gamma)$ gives a unit polynomial learning rate, the choice of $\mathsf{C_p} = (1+\gamma)\kappa(\gamma)$ minimizes the leading coefficient in front of the decaying term, namely, $\limsup_{t\to\infty} t \cdot \mathsf{P}_{e,t}(\mathcal{P}^\star,\mathcal{P}^\star) < \limsup_{t\to\infty} t \cdot \mathsf{P}_{e,t}(\mathcal{P},\mathcal{P})$, for any $\mathcal{P}$ with $\mathsf{C_p} > (1+\gamma)\kappa(\gamma)$. This explains also why $\mathsf{E}(\mathcal{P},\mathcal{Q})$ in Theorem 5 is increasing as a function of $\mathsf{C_p}$. Specifically, it is seen that when $\mathsf{C_p} > (1+\gamma)\kappa(\gamma)$, there are values of ρ for which $\mathsf{E}(\mathcal{P}^\star,\rho\cdot\mathcal{P}^\star) < \mathsf{E}(\mathcal{P},\rho\cdot\mathcal{P})$. Indeed, in case of mismatch, taking $\mathcal{P}^\star$ to be the underlying revealing probabilities might be sub-optimal, and choosing a different set of probabilities which combat the mismatch results in a higher rate. Finally, we consider the case where $\|\mathcal{P}_t\| = \omega(\log t)$, for which we have the following result.

Theorem 6 Let $\mathcal{P}$ be such that $\|\mathcal{P}_t\| / \log t \to \infty$. If the mismatched revealing probabilities $\mathcal{Q}$ is such that,

$$\limsup_{t\to\infty} \frac{\|\mathcal{Q}_t\|}{\|\mathcal{P}_t\|} = \rho \in \mathbb{R}_+, \tag{21}$$

with $\rho \le \rho_0$ or $\rho \ge \rho_1$, or

$$\liminf_{t\to\infty} \frac{\|\mathcal{Q}_t\|}{\|\mathcal{P}_t\|} = 0 \quad \text{or} \quad \limsup_{t\to\infty} \frac{\|\mathcal{Q}_t\|}{\|\mathcal{P}_t\|} = \infty,$$

then $\mathsf{E}(\mathcal{P},\mathcal{Q}) = 0$. Otherwise, if $\rho_0 < \rho < \rho_1$, then

$$\mathsf{E}(\mathcal{P},\mathcal{Q}) = \lim_{t\to\infty} -\frac{\log p_t}{\log t}.$$

Theorem 6 states that if the number of revealers is significantly bigger than $\log t$, then the error probability is dominated by the probability that the t^{th} player is a revealer and announce a wrong decision, namely, $\frac{\beta}{\alpha+\beta}p_t$. Therefore, if for example $p_t = \Theta(t^{-c})$, for some $c \in (0,1)$, then $\mathsf{E}(\mathcal{P},\mathcal{Q}) = c$, as long as the mismatch is not too "severe", namely, $\|\mathcal{Q}_t\| / \|\mathcal{P}_t\| \to \rho$, and $\rho_0 < \rho < \rho_1$. Otherwise, learning does not occur and $\mathsf{E}(\mathcal{P},\mathcal{Q}) = 0$.

Finally, we mention here that in Appendix D we list many interesting directions for future study of both theoretical and practical significance.

Acknowledgements

This research was supported by the European Research Council (ERC) under grant agreement 639573.

References

[Acemoglu et al., 2011] Acemoglu, D., Dahleh, M. A., and Ozdaglar, I. L. A. (2011). Bayesian learning in social networks. *The Review of Economic Studies*, 78:1201–1236.

[Anderson and Holt, 1997] Anderson, L. and Holt, C. (1997). Information cascades in the laboratory. *American Economic Review*, 87(5):847–62.

[Anderson and Holt, 1996] Anderson, L. R. and Holt, C. A. (1996). Classroom games: Information cascades. *Journal of Economic Perspectives*, 10(4):187–193.

[Anunrojwong and Sothanaphan, 2018] Anunrojwong, J. and Sothanaphan, N. (2018). Naive bayesian learning in social networks. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC'18, pages 619–636.

[Banerjee, 1992] Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–818.

[Bernardo and Welch, 2001] Bernardo, A. E. and Welch, I. (2001). On the evolution of overconfidence and entrepreneurs. *Journal of Economics & Management Strategy*, 10(3):301–330.

[Bikhchandani et al., 1992] Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026.

[Cheng et al., 2018] Cheng, Y., Hann-Caruthers, W., and Tamuz, O. (2018). A deterministic protocol for sequential asymptotic learning. *arXiv:1802.06871*.

[David and Jon, 2010] David, E. and Jon, K. (2010). *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, New York, NY, USA.

[Fan et al., 2015] Fan, X., Grama, I., and Liu, Q. (2015). Exponential inequalities for martingales with applications. *Electron. J. Probab.*, 20:22 pp.

[Freedman, 1975] Freedman, D. A. (1975). On tail probabilities for martingales. *Ann. Probab.*, 3(1):100–118.

[Hann-Caruthers et al., 2017] Hann-Caruthers, W., Martynov, V. V., and Tamuz, O. (2017). The speed of sequential asymptotic learning. *arXiv:1707.02689*.

[Harel et al., 2018] Harel, M., Mossel, E., Strack, P., and Tamuz, O. (2018). Rational groupthink. *arXiv:1412.7172*.

[Huck and Oechssler, 2000] Huck, S. and Oechssler, J. (2000). Informational cascades in the laboratory: Do they occur for the right reasons? *Journal of Economic Psychology*, 21(6):661–671.

[Kahneman and Tversky, 1973] Kahneman, D. and Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80(4):237–251.

[Mossel and Tamuz, 2017] Mossel, E. and Tamuz, O. (2017). Opinion exchange dynamics. *Probab. Surveys*, 14:155–204.

[Peres et al., 2018] Peres, Y., Racz, M. Z., Sly, A., and Stuhl, I. (2018). How fragile are information cascades? *arXiv:1711.04024*.

[Smith and Srensen, 1996] Smith, L. and Srensen, P. (1996). Pathological outcomes of observational learning. *Econometrica*, 68.

[Weizscker, 2010] Weizscker, G. (2010). Do we follow others when we should? a simple test of rational expectations. *The American Economic Review*, 100(5):2340–2360.