

Observational Learning with Fake Agents

Pawan Poojary and Randall Berry

ECE Department, Northwestern University, Evanston IL 60208

Email: pawanpoojary2018@u.northwestern.edu, rberry@ece.northwestern.edu

Abstract—It is common in online markets for agents to learn from other’s actions. Such observational learning can lead to *herding* or *information cascades* in which agents eventually “follow the crowd”. Models for such cascades have been well studied for Bayes-rational agents that choose pay-off optimal actions. In this paper, we additionally consider the presence of *fake agents* that seek to influence other agents into taking one particular action. To that end, these agents take a fixed action in order to influence the subsequent agents towards their preferred action. We characterize how the fraction of such fake agents impacts behavior of the remaining agents and show that in certain scenarios, an increase in the fraction of fake agents in fact reduces the chances of their preferred outcome.

I. INTRODUCTION

Consider a new item that is up for sale in a recommendation-based market where agents arrive sequentially and decide whether to buy the item, with their choice serving as a recommendation for later agents. Every agent makes a pay-off optimal decision based on his own prior knowledge of the item’s quality/utility and by observing the choices of his predecessors. An *informational cascade* or *herding* occurs when it is optimal for an agent to ignore his own private signal and follow the actions of the past agents. Subsequent agents follow suit and from the onset of a cascade, the agents’ actions do not reveal any information conveyed to them by their private signals. Thus the phenomenon of cascading could prevent the agents from learning the socially optimal (right) choice, *i.e.* the agents could herd to a wrong cascade.

Cascades were first studied in [1]–[3] as a Bayesian learning model in which a homogeneous sequence of rational agents take pay-off optimal actions. In this paper, in addition to rational agents, we introduce randomly arriving fake agents that always take (or *fake*) a fixed action, regardless of their pay-off, in order to influence the outcome of a cascade. For example, this could model a setting where a poor quality item is posted for sale on a website that records buyers’ decisions, and fake agents intentionally buy (or appear to buy) this item to make it seem worth buying to future buyers. The model could also represent a setting where rational actions are manipulated or where fake actions are inserted into a recommendation system. The objective is to study the impact of varying the amount of these fake agents on the probability of their preferred cascade. Our main result shows a counter-intuitive phenomenon: the probability with which this preferred cascade occurs is not monotonically increasing in the fraction of fake

agents, ϵ . In fact, in some cases the presence of fake agents can reduce the chance of their preferred cascade.

As in [1]–[3], we consider a setting in which all agents can observe all prior actions (although in our case, each action could be either rational or fake). Many other variations of this basic model have been studied such as [4], where observations depend on an underlying network structure and [5], where agents are allowed to select their observations at a cost.

Closer to our work is the model in [6], which assumes that the recording of actions for subsequent agents is subject to an error that is unbiased towards each of the possible actions. In our setting, an action being either fake or rational depending on the agent-type could equivalently be perceived as an error while recording a rational action (as in [6]); except that in our case, the error is biased only towards a preferred action¹.

There is also a body of work that considers agents similar to our fake agents, who only take a single action regardless of the true state. This includes the *crazy agents* considered in [7], *stubborn agents* in [8] and *zealots* in [9]. While [7] relaxes the assumption of binary signals, it does not consider how changing the fraction of crazy agents affects the cascade probability, which is the main focus of our work. In [8] a non-Bayesian model for learning was considered instead of the Bayesian model considered here. Other related cascade models are investigated in [10]–[17].

The paper is organized as follows. We describe our model in Section II. We analyze this model and identify the resulting cascade properties in Section III. In Section IV, we present our Markov process formulation and in Section V, we consider the limiting scenario where the proportion of fake agents approaches unity. We conclude in Section VI.

II. MODEL

We consider a model similar to [1] in which there is a countable sequence of agents, indexed $i = 1, 2, \dots$ where the index represents both the time and the order of actions. Each agent i takes an action A_i of either buying (Y) or not buying (N) a new item that has a true value (V) which could either be good (G) or bad (B). For simplicity, both possibilities of V are assumed to be equally likely. To incorporate agents’ private beliefs about the new item, every agent i receives a private signal $S_i \in \{H \text{ (high)}, L \text{ (low)}\}$. This signal, as shown in Figure 1(a), partially reveals the information about the true

This work was supported in part by the NSF under grants CNS-1701921 and CNS-1908807.

¹This change requires a different analysis approach than that in [6] as the underlying Markov process now typically has an uncountable state-space.

value of the item through a binary symmetric channel (BSC) with crossover probability $1 - p$, where $1/2 < p < 1$. Each agent i takes a *rational* action A_i that depends on his private signal S_i and the past observations $\{O_1, O_2, \dots, O_{i-1}\}$ of actions $\{A_1, A_2, \dots, A_{i-1}\}$. Next, we modify the information structure in [1] by assuming that at each time instant, an agent could either be *fake* with probability (w.p.) $\epsilon \in [0, 1)$ or *ordinary* w.p. $1 - \epsilon$. An ordinary agent i *honestly* reports his action, i.e. $O_i = A_i$. On the contrary, a fake agent *always* reports a Y , reflecting his intention of influencing the successors into buying the new item, regardless of its true value. This implies that at any time i , with probability $1 - \epsilon$, the reported action O_i satisfies $O_i = A_i$ and with probability ϵ , $O_i = Y$. Refer to Figure 1(b).

An equivalent model is where a fake agent always takes an action Y , while $O_i = A_i$ for all agents. This yields the same information structure as the model above. However we chose the former model mainly to simplify our notation.




Fig. 1: (a) The BSC through which agents receive private signals. (b) The channel through which agents' actions are corrupted.

III. OPTIMAL DECISION AND CASCADES

For the n^{th} agent, let the history of past observations be denoted by $\mathcal{H}_{n-1} = \{O_1, O_2, \dots, O_{n-1}\}$. As the first agent does not have any observation history, he always follows his private signal i.e., he buys if and only if the signal is H . For the second agent onwards, the Bayes' optimal action for every agent n , A_n is chosen according to the hypothesis ($V = G$ or B) that has the higher posterior probability given the information set $\mathbb{I}_n = \{S_n, \mathcal{H}_{n-1}\}$. Let $\gamma_n(S_n, \mathcal{H}_{n-1}) \triangleq \mathbb{P}(G|S_n, \mathcal{H}_{n-1})$ denote the posterior probability for the item being good, $V = G$. Then the Bayes' optimal decision rule is:

$$A_n = \begin{cases} Y, & \text{if } \gamma_n > 1/2, \\ N, & \text{if } \gamma_n < 1/2, \\ \text{follows } S_n, & \text{if } \gamma_n = 1/2. \end{cases} \quad (1)$$

Note that when $\gamma_n = 1/2$, an agent is indifferent between the actions. Similar to [6], our decision rule in this case follows the private signal S_n , unlike [1] where they employ random tie-breaking.

Definition 1: An information cascade is said to occur when an agent's decision becomes independent of his private signal

It follows from (1) that, agent n cascades to a Y (N) if and only if $\gamma_n > 1/2$ ($< 1/2$) for all $S_n \in \{H, L\}$. The other case being $\gamma_n > 1/2$ for $S_n = H$ and less than $1/2$ for $S_n = L$; in which case, agent n follows S_n . A more intuitive way to express this condition is by using the information contained

in the history $\mathcal{H}_{n-1}$ observed by agent n in the form of the likelihood ratio $l_{n-1}(\mathcal{H}_{n-1}) \triangleq \mathbb{P}(\mathcal{H}_{n-1}|B)/\mathbb{P}(\mathcal{H}_{n-1}|G)$ as follows [7].

Lemma 1: Agent n cascades to a Y (N) if and only if $l_{n-1} < \frac{1-p}{p}$ ($l_{n-1} > \frac{p}{1-p}$) and otherwise follows its private signal S_n .

This lemma follows by expressing γ_n in terms of l_{n-1} using Bayes' law, and then using the condition on γ_n for a Y (N) cascade. If agent n cascades, then the observation O_n does not provide any additional information about the true value V to the successors than what is contained in $\mathcal{H}_{n-1}$. As a result, $l_{n+i} = l_{n-1}$ for all $i = 0, 1, 2, \dots$ and hence they remain in the cascade, which leads us to the following property, also exhibited by prior models [1]–[3], [6].

Property 1: Once a cascade occurs, it lasts forever.

On the other hand, if agent n does not cascade, then Property 1 and Lemma 1 imply that all the agents until and including n follow their own private signals ignoring the observations of their predecessors. Mutual independence between the private signals results in the likelihood ratio l_n depending only on the number of Y 's (denoted by n_Y) and N 's (denoted by n_N) in the observation history $\mathcal{H}_n$. Specifically, $l_n = \left(\frac{1-p}{p}\right)^{h_n}$ where $h_n = \eta n_Y - n_N$ is the difference between the number of Y 's weighted by η and the number of N 's. Here, the weight $\eta \triangleq \log(\frac{a}{1-b})/\log(\frac{p}{1-p})$, where for all $i = 1, 2, \dots, n-1$,

$$a = \mathbb{P}(O_i = Y|V = G) \quad \text{and} \quad b = \mathbb{P}(O_i = N|V = B)$$

give the probabilities that the observations follow the true value V , given that these agents follow their own private signals. It can be shown from Figures 1(a) and 1(b) that in the above case, i.e., when A_i follows S_i ,

$$a = p + (1-p)\epsilon \quad \text{and} \quad b = p(1-\epsilon).$$

Thus, agents that have not yet entered a cascade satisfy the following property.

Property 2: Until a cascade occurs, each agent follows its private signal. Moreover, h_n is a sufficient statistic of the information contained in the past observations.

Note that if $\epsilon = 0$ (no fake agents) then $a = b = p$ and $\eta = 1$, in which case h_n is the same as in [6]. The expression for h_n shows that, due to the presence of fake agents, the dependence of an agent's decision on a Y in his observation history reduces by a factor of η , whereas the dependence on a N remains unaffected. This is to be expected because, unlike a N which surely comes from an honest agent, a Y incurs the possibility that the agent could be fake. Further, this reduced dependence on Y is exacerbated with an increase in the possibility of fake agents, as η reduces with an increase in ϵ .

Using the expression for l_n in Lemma 1, it follows that until a cascade occurs, $-1 \leq h_n \leq 1$ for all such times n , and the update rule for h_n is given by

$$h_n = \begin{cases} h_{n-1} + \eta & \text{if } O_n = Y, \\ h_{n-1} - 1 & \text{if } O_n = N. \end{cases} \quad (2)$$

Whereas, once $h_n > 1$ (< 1), a Y (N) cascade begins and h_n stops updating (Property 1).

IV. MARKOVIAN ANALYSIS OF CASCADES

Let us, for the sake of discussion, consider the realization $V = B$. Similar arguments hold when $V = G$. It follows from the previous section that the process $\{h_n\}$ is a discrete-time Markov process taking values in $[-1, 1]$ before getting absorbed into the left wall (< -1) causing a N cascade or the right wall (> 1) causing a Y cascade. More specifically, equation (2) shows that, until a cascade occurs, $\{h_n\}$ is a random walk that shifts to the right by η w.p. $\mathbb{P}(O_n = Y|B) = 1 - b$ or to the left by 1 w.p. $\mathbb{P}(O_n = N|B) = b$. Note that the walk starts at state 0 since the first agent has no observation history.

The random walk is depicted in Figure 2 where $p_f \triangleq \mathbb{P}(O_n = Y|V)$ denotes the probability of a Y being observed. Depending on the item's true value, $p_f = a$ for $V = G$ whereas $p_f = 1 - b$ for $V = B$. Note that in the special case where η satisfies $1/\eta = r$ for some $r = 1, 2, \dots$, the process $\{h_n\}$ is equivalent to a Markov Chain with finite state-space $\mathcal{A} = \{-r-1, -r, \dots, -1, 0, 1, \dots, r, r+1\}$, and with $-r-1$ and $r+1$ being absorption states corresponding to N and Y cascades respectively. In this case, absorption probabilities can be obtained by solving a system of linear equations. In this paper, our main focus is on the more common case of non-integer values of $1/\eta$ resulting in $\{h_n\}$ taking possibly uncountable values in $[-1, 1]^2$.




Fig. 2: Transition diagram of the random walk.

A. Error thresholds

In the absence of fake agents ($\epsilon = 0$) as in [1], $\eta = 1$ and so cascading to a $Y(N)$ cascade requires at least two consecutive Y 's (N 's). However, in the presence of fake agents, even a single N after a Y could trigger a N cascade. On the other hand, as ϵ increases, a greater number of consecutive Y 's (≥ 2) are required to cause a Y cascade. This is characterized in the following lemma.

Lemma 2: Let $\alpha = p/(1 - p)$. For $r = 1, 2, \dots$ define the increasing sequence of thresholds $\{\epsilon_r\}_{r=1}^\infty$, where the r^{th} threshold ϵ_r is given as:

$$\epsilon_r = \frac{\alpha - \alpha^{\frac{1}{r}}}{\alpha^{\frac{1}{r}+1} - 1}. \quad (3)$$

Define $\mathcal{I}_r \triangleq [\epsilon_r, \epsilon_{r+1})$ as the r^{th} ϵ -interval. Then for $\epsilon \in \mathcal{I}_r$, at least $r + 1$ consecutive Y 's are necessary for a Y cascade to begin.

The proof follows by noting that, given that if a cascade has not begun, when an N is observed, then the rightmost position that the random walk can be in, is in state 0. From here, at

²For example, if η was chosen uniformly at random, then with probability 1 it would fall into this case.




Fig. 3: Thresholds ϵ_r for the indicated values of r versus p .

least $r + 1$ consecutive Y 's would be needed only if $\eta \leq 1/r$, thereby giving the r^{th} threshold ϵ_r . It follows that $\epsilon \in \mathcal{I}_r$ implies $\frac{1}{r+1} < \eta \leq \frac{1}{r}$.

Remark 1: For $\epsilon \in \mathcal{I}_r$, starting from state 0, $r+1$ consecutive Y 's start a Y cascade.

Moreover, given ϵ , the integer r such that $\epsilon \in \mathcal{I}_r$ can be readily obtained from Lemma 2 as $r = \lfloor 1/\eta \rfloor$. Here, recall that η is a function of ϵ .

Figure 3 shows the thresholds ϵ_r varying with p , for different values of r . For a fixed ϵ , we see that as the signal quality (p) improves, more consecutive Y 's are required for a Y cascade to begin. This is because, an increase in p increases the information contained in a Y , but not as much as the corresponding increase in the information contained in a N . Further, note that as $\epsilon \rightarrow 1$, $r \rightarrow \infty$ which implies that infinitely many consecutive Y 's are required for a Y cascade to begin. Equivalently, the information contained in a Y observation becomes negligible. We further investigate learning in this asymptotic scenario in Section V.

B. Y cascade probability

In this subsection, we will compute the probability of absorption of $\{h_n\}$ to a Y cascade (right wall). The following iterative process depicted in Figure 4 describes all possible sequences that can lead to a Y cascade. For the rest of the paper, we assume $\epsilon \in \mathcal{I}_r \setminus \{\epsilon_r\}$ for some $r = 1, 2, \dots$. We initialize Stage 1 with $r_1 = r + 1$. Now, starting from state 0, consider the sequences shown in Stage 1 in Figure 4. The first sequence of r_1 consecutive Y 's, denoted by Y^{r_1} , clearly hits the right wall (Remark 1). The rest of the sequences, each of length $r_1 + 1$, are simply permutations of each other that contain only a single N . Two N 's or more are not possible without hitting the left wall. So each of these r_1 distinct sequences results in the same net right shift, which ends somewhere in the region $[0, \eta]$. From here, it would take either r or $r + 1$ consecutive Y 's to hit the right wall. Let this value be denoted by r_2 . The sequences that would then be possible from this point onwards could again be enumerated

exactly as in the first stage, except that r_2 now replaces r_1 . This forms Stage 2. Now, unless there are r_2 consecutive Y 's, the second stage would again end in the region $[0, \eta]$, and then the process continues to the next stage. Here, r_n denotes the number of consecutive Y 's required to hit the right wall in the n^{th} stage. Note that by considering $\epsilon \in \mathcal{I}_r \setminus \{\epsilon_r\}$, we avoid the case of integer values of $1/\eta$ which could result in certain pathological sequences with non-zero probability, that are not enumerated in Figure 4.

Let S_n denote the probability of hitting the right wall given that the sequence has not terminated before the n^{th} stage. Thus, the following recursion holds:

$$S_n = p_f^{r_n} [1 + r_n(1 - p_f)S_{n+1}], \text{ for } n = 1, 2, \dots \quad (4)$$

and the probability of a Y cascade, denoted by $\mathbb{P}_{Y\text{-cas}}$ is:

$$\mathbb{P}_{Y\text{-cas}}(\epsilon) = S_1, \text{ for } \epsilon \in \mathcal{I}_r \setminus \{\epsilon_r\}. \quad (5)$$

Here, while $r_1 = r + 1$ for $\epsilon \in \mathcal{I}_r$, successive values of r_i for $i = 2, 3, \dots$ can be obtained from r_1 using the updates:

$$r_n = \begin{cases} r, & \text{if } \sum_{i=1}^{n-1} (r_i \eta - 1) + r \eta > 1, \\ r + 1, & \text{o.w.} \end{cases} \quad (6)$$

Since (4) is an infinite recursion, to compute $\mathbb{P}_{Y\text{-cas}}$ in practice, we truncate the process to a finite number of iterations M . To this end, we first assume that $S_{M+1} = 1$. Next, we use (4) to successively compute S_k while k counts down from M to 1, performing a total of M iterations. We denote the obtained value as $\mathbb{P}_{Y\text{-cas}}^M$. We now state in the following theorem that $\mathbb{P}_{Y\text{-cas}}^M$ is in fact a tight upper bound to $\mathbb{P}_{Y\text{-cas}}$ as $M \rightarrow \infty$. Moreover, the difference $\mathbb{P}_{Y\text{-cas}}^M - \mathbb{P}_{Y\text{-cas}}$ decays to zero at least as fast as $\{0.5^M\}$, in the number of iterations M .

Theorem 1: Let $\epsilon \in \mathcal{I}_r$ for some $r = 1, 2, \dots$, with p_f denoting the probability of a Y . Then, for any $M = 1, 2, \dots$,

$$0 \leq \mathbb{P}_{Y\text{-cas}}^M(\epsilon) - \mathbb{P}_{Y\text{-cas}}(\epsilon) \leq k^M,$$

where $k \triangleq (r + 1)(1 - p_f)p_f^r$ and for any $p \in (0.5, 1)$ and $\epsilon \in [0, 1]$, k satisfies $0 < k \leq 1/2$.

Proof: Recall that $S_{M+1} = 1$ while computing $\mathbb{P}_{Y\text{-cas}}^M$, which implies that $\mathbb{P}_{Y\text{-cas}}^M$ is the net probability of: (a) sequences that terminate in a Y cascade by the M^{th} stage and (b) sequences that do not terminate by the M^{th} stage. Thus, the




Fig. 4: An enumeration of all possible sequences that would lead to a Y cascade.




Fig. 5: Probability of Y cascade as a function of ϵ for $V = B$ and $p = 0.7$.

difference: $\mathbb{P}_{Y\text{-cas}}^M - \mathbb{P}_{Y\text{-cas}}$ is upper-bounded by the probability of (b). Accounting for all possible sequence combinations that persist through stages $1, 2, \dots, M$ yields the probability k . Next, for a fixed p and for $\epsilon \in \mathcal{I}_r$, k is maximized only if ϵ is such that $p_f = \frac{r}{1+r}$, which may not be satisfied for any ϵ in $\mathcal{I}_r$. Assuming k is maximized, the maximal value of k would be $(\frac{r}{1+r})^r$. As this maximal value decreases in r , evaluating it at $r = 1$ yields $k \leq 1/2$ for any r . ■

Figure 5 shows a plot of $\mathbb{P}_{Y\text{-cas}}$ with respect to ϵ , for the case $V = B$ with $p = 0.7$. The plot uses $M = 10$ which gives an error of less than 10^{-3} . It can be seen that in the r^{th} ϵ -interval $\mathcal{I}_r$, $\mathbb{P}_{Y\text{-cas}}$ increases with ϵ , but with infinitely many discontinuities (where $\mathbb{P}_{Y\text{-cas}}(\epsilon)$ decreases). Despite the discontinuities, $\mathbb{P}_{Y\text{-cas}}$ achieves the minimum and maximum values at the edge points of $\mathcal{I}_r$, i.e., ϵ_r^+ and ϵ_{r+1}^- , respectively. Further, note that the relatively larger drops in $\mathbb{P}_{Y\text{-cas}}(\epsilon)$ observed in Figure 5 (marked by $\times$ and $\circ$) occur exactly at the threshold points $\{\epsilon_r\}_{r=1}^\infty$. Here, counter to expectation, a slight increase in ϵ beyond ϵ_r causes a significant decrease in the probability of a Y cascade. We now quantify $\mathbb{P}_{Y\text{-cas}}$ as ϵ tends to each threshold point ϵ_r . For this, we state the following lemma:

Lemma 3: For all $i \geq 2$, $r_i \rightarrow r$ as $\epsilon \rightarrow \epsilon_r$.

Proof: We consider two cases: $\epsilon \rightarrow \epsilon_r^+$ and $\epsilon \rightarrow \epsilon_r^-$. The proof outline for the first case is as follows: assume $r_j = r$ for all $2 \leq j \leq i - 1$ and note that $r_1 = r + 1$ as $\epsilon_r^+ \in \mathcal{I}_r$. Then it follows from (6) that $r_i = r$ only if $\eta > (r + \frac{1}{i})^{-1}$. Now as $\epsilon \rightarrow \epsilon_r^+$, $\eta \rightarrow 1/r$ and hence the condition for $r_i = r$ is satisfied. Using this argument inductively shows that $r_i = r$ for all $i \geq 2$. The second case is proved similarly. ■

It follows from Lemma 3 that as $\epsilon \rightarrow \epsilon_r$, the recursion in (4) results in the same infinite computation to obtain S_i as for S_{i+1} , for all $i \geq 2$. Thus, all S_i for $i \geq 2$ have the same value which satisfies: $S_i = p_f^r [1 + r(1 - p_f)S_i]$. Solving this equation for $i = 2$ gives the value for S_2 which when used in equation (4) for $n = 1$ yields S_1 , i.e., $\mathbb{P}_{Y\text{-cas}}$. However, note that while solving equation (4), $r_1 = r + 1$ for $\epsilon = \epsilon_r^+$ whereas $r_1 = r$ for $\epsilon = \epsilon_r^-$. This corresponds to the following two different

values of $\mathbb{P}_{Y\text{-cas}}$ as $\epsilon \rightarrow \epsilon_r$:

$$\mathbb{P}_{Y\text{-cas}}(\epsilon_r^+) = p_f^{r+1} \frac{1 + (1 - p_f)p_f^r}{1 - r(1 - p_f)p_f^r}, \quad (7)$$

$$\mathbb{P}_{Y\text{-cas}}(\epsilon_r^-) = p_f^r \frac{1}{1 - r(1 - p_f)p_f^r}. \quad (8)$$

Hence, the fractional decrease in $\mathbb{P}_{Y\text{-cas}}$ that occurs abruptly at ϵ_r , defined as $\delta_r = [\mathbb{P}_{Y\text{-cas}}(\epsilon_r^-) - \mathbb{P}_{Y\text{-cas}}(\epsilon_r^+)] / \mathbb{P}_{Y\text{-cas}}(\epsilon_r^-)$ is:

$$\delta_r = (1 - p_f)(1 - p_f^{r+1}). \quad (9)$$

It can be verified that as $r \rightarrow \infty$, $\delta_r \rightarrow 0$. This is depicted in Figure 5 where the sequences $\{\mathbb{P}_{Y\text{-cas}}(\epsilon_r^-)\}$ and $\{\mathbb{P}_{Y\text{-cas}}(\epsilon_r^+)\}$, marked by $\times$ and $\circ$ respectively, converge to a limiting value as $r \rightarrow \infty$.

Property 3: If the possibility of fake agents in the history equals the r^{th} ϵ -threshold, $r = 2, 3, \dots$, then a further increase in fake agents reduces the chances of a Y cascade by a factor of δ_r , rather than increasing it.

Next, we consider the cascade behaviour for low values of ϵ and state the following theorem.

Theorem 2: Given the private signal quality $p \in (0.5, 1)$, and the item's true value $V \in \{G, B\}$, there exists some $\underline{\epsilon} = f(V, p) > 0$ such that

$$\mathbb{P}_{Y\text{-cas}}(\epsilon) < \mathbb{P}_{Y\text{-cas}}(0), \quad \forall \epsilon \in (0, \underline{\epsilon}). \quad (10)$$

From the point of view of the fake agents, the above theorem implies that if they occur with a probability of less than $\underline{\epsilon}$, then the effect that their presence has on the honest buyers is opposite to what they had intended. That is, they reduce the chances of a Y cascade instead of increasing it. On the other hand, for honest buyers, if $V = B$, then they benefit from the presence of fake agents. Otherwise if $V = G$, then they are harmed by the presence of fake agents.

The proof for Theorem 2 follows by noting that as $\epsilon \rightarrow 0$, the limiting value of $\mathbb{P}_{Y\text{-cas}}$ can be obtained from (7) with $r = 1$ and $p_f \rightarrow 1 - p$ for $V = B$, whereas $p_f \rightarrow p$ for $V = G$, which gives

$$\lim_{\epsilon \rightarrow 0} \mathbb{P}_{Y\text{-cas}}(\epsilon) = \begin{cases} (1 - p)^2 \frac{1 + p(1 - p)}{1 - p(1 - p)}, & \text{for } V = B, \\ p^2 \frac{1 + p(1 - p)}{1 - p(1 - p)}, & \text{for } V = G. \end{cases} \quad (11)$$

However, in the absence of fake agents ($\epsilon = 0$) as in [1], $h_n = n_Y - n_N$ and Y cascade starts when $h_n = 2$. In this case, $\mathbb{P}_{Y\text{-cas}}(0) = p_f^2 / [p_f^2 + (1 - p_f)^2]$, with $p_f = 1 - p$ for $V = B$, whereas $p_f = p$ for $V = G$. By comparing this expression with the one in (11) for the corresponding values of V ,

$$\lim_{\epsilon \rightarrow 0} \mathbb{P}_{Y\text{-cas}}(\epsilon) < \mathbb{P}_{Y\text{-cas}}(0), \quad \forall p \in (0.5, 1), V \in \{G, B\}.$$

Thus, there exists some $\underline{\epsilon} > 0$ such that (10) holds true.

V. LEARNING IN THE LIMIT AS $\epsilon \rightarrow 1$

As the probability of an agent being fake tends to 1, the information contained in a Y observation becomes negligible. As a result, an agent would need to observe infinitely many consecutive Y 's in his history for him to be convinced of

starting a Y cascade. Hence, one would expect that if $V = G$, learning would never occur, whereas if $V = B$, then learning would always occur. However, recall that as $\epsilon \rightarrow 1$, the occurrence of Y 's becomes increasingly frequent, i.e. $p_f \rightarrow 1$; for both $V = B$ and G . This motivates studying $\mathbb{P}_{Y\text{-cas}}$ in this limiting scenario. First, recall that in the process of enumerating all sequences leading to a Y cascade, for $\epsilon \in \mathcal{I}_r$, in each stage $i \geq 2$, r_i is either r or $r + 1$. However, $\epsilon \rightarrow 1$ implies $r \rightarrow \infty$, in which case $r \approx r + 1$. As a result, the expressions obtained in (7) and (8) yield the same limiting value for $\mathbb{P}_{Y\text{-cas}}$ as $\epsilon_r \rightarrow 1$, i.e., $r \rightarrow \infty$. In particular,

$$\lim_{\epsilon \rightarrow 1} \mathbb{P}_{Y\text{-cas}}(\epsilon) = \lim_{r \rightarrow \infty} \mathbb{P}_{Y\text{-cas}}(\epsilon_r^+) \stackrel{(a)}{=} \lim_{r \rightarrow \infty} \mathbb{P}_{Y\text{-cas}}(\epsilon_r^-),$$

where Step (a) can be proved by recalling that $\delta_r \rightarrow 0$ as $r \rightarrow \infty$. Further, the limiting probability of a Y cascade, in terms of $\alpha = p/(1 - p)$, is given by (refer to [18] for proof):

$$\lim_{\epsilon \rightarrow 1} \mathbb{P}_{Y\text{-cas}}(\epsilon) = \lim_{r \rightarrow \infty} p_f^r \frac{1}{1 - r(1 - p_f)p_f^r} = \frac{1}{e^t - t}; \quad (12)$$

where $t = \frac{1}{\alpha - 1} \log \alpha$ for $V = G$, and $t = \frac{\alpha}{\alpha - 1} \log \alpha$ for $V = B$.




Fig. 6: Probability of Y cascade versus private signal quality for the indicated values of ϵ .

Figure 6 illustrates (12) and the corresponding probability when $\epsilon = 0.9$ as a function of the signal quality. For both $V = B$ and $V = G$, a better signal quality leads to improved learning even when fake agents have overwhelmed the ordinary agents. Also note that for $V = B$, for a weak signal quality, the incorrect cascade is more likely than the correct one (while for $V = G$ this is never true).

VI. CONCLUSIONS AND FUTURE WORK

We studied the effect of randomly arriving fake agents that seek to influence the outcome of an information cascade. We focussed on the impact of varying the amount of fake agents on the probability of their preferred cascade and identified scenarios where the presence of fake agents reduces the chances of their preferred cascade. In particular, the chances of the preferred cascade decrease abruptly as the fraction of agents increase beyond specific thresholds. We quantified the correct cascade probability as a function of the private signal quality in the asymptotic regime where the fraction of fake agents approaches unity. Studying the effects of multiple types of fake agents and non-Bayesian rationality are possible directions for future work.

REFERENCES

- [1] S. Bikhchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of political Economy*, vol. 100, no. 5, pp. 992–1026, 1992.
- [2] A. V. Banerjee, "A simple model of herd behavior," *The quarterly journal of economics*, vol. 107, no. 3, pp. 797–817, 1992.
- [3] I. Welch, "Sequential sales, learning, and cascades," *The Journal of finance*, vol. 47, no. 2, pp. 695–732, 1992.
- [4] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *The Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.
- [5] Y. Song, "Social learning with endogenous network formation," *arXiv preprint arXiv:1504.05222*, 2015.
- [6] T. N. Le, V. G. Subramanian, and R. A. Berry, "Information cascades with noise," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 3, no. 2, pp. 239–251, 2017.
- [7] L. Smith and P. Sørensen, "Pathological outcomes of observational learning," *Econometrica*, vol. 68, no. 2, pp. 371–398, 2000.
- [8] D. Acemoglu, G. Como, F. Fagnani, and A. Ozdaglar, "Opinion fluctuations and disagreement in social networks," *Mathematics of Operations Research*, vol. 38, no. 1, pp. 1–27, 2013.
- [9] M. Mobilia, "Does a single zealot affect an infinite group of voters?" *Physical review letters*, vol. 91, no. 2, p. 028701, 2003.
- [10] I. H. Lee, "On the convergence of informational cascades," *Journal of Economic theory*, vol. 61, no. 2, pp. 395–411, 1993.
- [11] Y. Wang and P. M. Djurić, "Social learning with bayesian agents and random decision making," *IEEE Transactions on Signal Processing*, vol. 63, no. 12, pp. 3241–3250, 2015.
- [12] T. Le, V. Subramanian, and R. Berry, "Quantifying the utility of noisy reviews in stopping information cascades," in *IEEE CDC Proceeding*, 2016.
- [13] Y. Peres, M. Z. Rácz, A. Sly, and I. Stuhl, "How fragile are information cascades?" *arXiv preprint arXiv:1711.04024*, 2017.
- [14] T. N. Le, V. G. Subramanian, and R. A. Berry, "Bayesian learning with random arrivals," in *2018 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2018, pp. 926–930.
- [15] I. Bistriz and A. Anastasopoulos, "Characterizing non-myopic information cascades in bayesian learning," in *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 2018, pp. 2716–2721.
- [16] Y. Cheng, W. Hann-Caruthers, and O. Tamuz, "A deterministic protocol for sequential asymptotic learning," in *2018 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2018, pp. 1735–1738.
- [17] G. Schoenebeck, S. Su, and V. G. Subramanian, "Social learning with questions," *CoRR*, vol. abs/1811.00226, 2018. [Online]. Available: <http://arxiv.org/abs/1811.00226>
- [18] P. Poojary and R. Berry, "Observational learning with fake agents," *arXiv preprint arXiv:2005.05518*, 2020.