

The Benefit of More Bad Choices in Observational Learning

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 others' 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 faced with deciding between two possible actions - one “good” action and one “bad” action. In this paper, we consider the case when these agents instead have more than two actions, where again only one of these is good. We show that sequential observational learning in such settings has substantially different properties compared to the binary action case and further show that increasing the number of “bad” choices from 1 to 2, can improve the agents' learning.

I. INTRODUCTION

Consider that two items, of which only one is “good”, are up for sale in a recommendation-based market where agents arrive sequentially and decide which of the two items to buy, with their choice serving as a recommendation for later agents. The identity of the good or equivalently the bad item is unknown to the agents. Each agent then makes a pay-off optimal decision by using its own prior knowledge of the items' qualities and by observing the choices of its predecessors. Such models of “observational learning” were first studied by [1]–[3] under a Bayesian learning framework wherein each agent has some prior knowledge in the form of a privately observed signal about the pay-off-relevant true state of the world, which in this case is the identity of the good item and has two possible realizations. In such models, 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 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.

In this paper, we study a Bayesian learning model similar to [1]–[3], and others [4], [5], except we consider that agents are faced with deciding among more than two alternatives, where only one of these is a “good” alternative and the others are “bad”. Once again, we assume that the identity of the good item corresponds to an unknown state of the world; agents receive private noisy signals indicating this state and again learn from observing other agents. The objective is to study

the information dynamics and cascades generated by agents' Bayes' optimal action sequence and to compare and contrast them with the dynamics of the binary true state models of [1]–[3] and many others.

An important motivation of this paper is to answer the following: how is learning affected in such models when more alternatives are introduced to the existing binary set of true states? For fairness of comparison, we assume both the original and the altered models maintain the same private signal quality. Now, if the information dynamics for a model with more than two true states were to behave similarly to the model with binary true states, then an intuition is that cascades will take longer to occur in the former model. This is because, unlike in the latter model, the statistic for the cascading action in the former model is required to dominate not one but all the remaining actions' statistics. Longer times to cascade would mean more information being revealed, leading to better learning for subsequent agents. However, our analysis reveals that the information dynamics for more than two true states substantially differs from that in the basic binary true state model, making the above intuition difficult to apply.

In related work, [6] also considers more than two true states and provides conditions such as directionally unbounded private beliefs, that guarantee learning. Our model maintains assumptions of [1]–[3], *i.e.*, discrete bounded private signals, which always leads to a positive probability that learning fails. Another work with multiple true states is [7], which considers non-Bayesian learning of the true state under repeated interactions of agents over a social network. Our work remains with the Bayesian model in [1]–[3], where each agent sequentially takes a one-time action and can observe all prior actions. Other variations of the basic model include relaxing the assumptions of i.i.d. binary valued signals [4], assuming agents don't observe all predecessors' actions [5], [8], allowing for imperfect observations [9]–[11], and others [12]–[14].

II. MODEL

We consider a model in which there is a countable sequence of agents, indexed $t = 1, 2, \dots$ where the index represents both the time and the order of actions. Each agent t takes an action $A_t \in \mathcal{A} = \{a_1, a_2, \dots, a_n\}$ of choosing to buy one among n items, which are indexed by $i = 1, 2, \dots, n$ where $n \geq 2$. While it is common knowledge that only one among the n items is “good” and all the rest are “bad”, the identity (index)

This work was supported in part by the NSF under grants CNS-1908807, ECCS-2030251 and ECCS-2216970.

of the good item is not known to the agents *a priori*. Let $\omega \in \Omega = \{1, 2, \dots, n\}$ denote the true identity of the good item. For simplicity, all possibilities of ω are assumed to be equally likely.

The agents are Bayes-rational utility maximizers where the pay-off received by each agent t , denoted by π_t , depends on the quality of the item he chooses to buy as follows. The agent gains the amount x if the chosen item is good, i.e., if $A_t = a_\omega$, and $-y$ if the chosen item is bad, where $x > 0$ and $y \geq 0$. He also incurs a fixed cost $C > 0$ for buying the item. The agent's net pay-off is then the gain minus the cost of buying the item, which is

$$\pi_t = \begin{cases} x - C, & \text{if } A_t = a_\omega, \\ -y - C, & \text{if } A_t \neq a_\omega. \end{cases} \quad (1)$$

Note that since ω is equiprobable, the *ex ante* expected pay-off for any agent is equal for all actions. Thus, to begin with, an agent is indifferent to all the actions.

To incorporate agents' private beliefs about the new items, every agent t receives a private signal $S_t \in \{s_1, s_2, \dots, s_n\}$. This signal, as shown in Figure 1 for $n = 3$, partially reveals the information about the true identity of the good item ω through a n -ary symmetric channel (n -ary SC) with crossover probability $1 - p$. Given the true value ω , S_t is distributed as

$$\mathbb{P}(S_t = s_k | \omega) = \begin{cases} p, & \text{if } k = \omega, \\ \frac{(1-p)}{n-1}, & \text{if } k \neq \omega. \end{cases} \quad (2)$$

Here, $1/n < p < 1$ which implies that the signal is informative but not revealing. Moreover, the sequence of private signals $\{S_1, S_2, \dots\}$ is assumed to be *i.i.d.* given the true value ω . Each agent t takes a *rational* action A_t that depends on his private signal S_t and the past actions $\{A_1, A_2, \dots, A_{t-1}\}$. Lastly, the models in [1]–[3] become special cases, with $n = 2$.




Fig. 1: Transition diagram of 3-ary SC through which agents receive their private signals. Transitions with solid and dashed arrows occur with probabilities p and $(1 - p)/2$, respectively.

III. OPTIMAL DECISION, SUFFICIENT STATISTICS AND CASCADES

For the t^{th} agent, the history of past actions $H_{t-1} := \{A_1, A_2, \dots, A_{t-1}\}$ and its private signal S_t form its information set $\{S_t, H_{t-1}\}$. As the first agent does not have any observation history, he always follows his private signal, i.e., he chooses item i if and only if the signal is s_i . For agent $t \geq 2$, the Bayes' optimal action, A_t is chosen such that it provides the greatest expected pay-off given the information set $\{S_t, H_{t-1}\}$. Let $\gamma_t^i(S_t, H_{t-1}) \triangleq \mathbb{P}(\omega = i | S_t, H_{t-1})$ denote the posterior probability that item i is the good item.

Further, let $\gamma_t := (\gamma_t^1, \gamma_t^2, \dots, \gamma_t^n)$ be the posterior distribution over Ω . Now, it follows from (1) that for any two actions a_i and a_j , with $i \neq j$, their respective π_t 's are symmetric over the set Ω . This implies that a_i is optimum over a_j only if $\gamma_t^i > \gamma_t^j$. Thus, a Bayes' optimal decision rule is given by

$$A_t = \begin{cases} a_i, & \text{if } M_t = \{i\}, \\ \text{follows } S_t, & \text{if } |M_t| > 1 \text{ and } S_t \in \{s_i\}_{i \in M_t}, \\ a_{\tau(M_t)}, & \text{if } |M_t| > 1 \text{ and } S_t \notin \{s_i\}_{i \in M_t}. \end{cases} \quad (3)$$

Here, $M_t := \arg \max_{i \in \Omega} \gamma_t^i$ denotes the index set of the optimal action(s). Note in (3) that when $|M_t| > 1$, a *tie* is said to occur and the agent is indifferent to the actions $\{a_i\}_{i \in M_t}$. Our decision rule in this case is to follow the private signal S_t , when $S_t \in \{s_i\}_{i \in M_t}$, i.e., when following the private signal is optimal. Otherwise, we select an action from the optimal set $\{a_i\}_{i \in M_t}$ as per a deterministic tie-breaking rule $\tau(\cdot)$, and denote the tie-winning action by $a_{\tau(M_t)}$.

Note that when $n = 2$, as in [1], [9], [10], there exists only a single possibility of a tie, which is between actions a_1 and a_2 . As $S_t \in \{s_1, s_2\}$, following the private signal S_t when in a tie, is never sub-optimal, unlike the third case in (3), which exists as a possibility only when $n > 2$.

Remark 1: The third case in (3) exists as a possibility only when $n > 2$. Thus, $\tau(\cdot)$ is used only for $n > 2$.

An example of such a possibility is when $n = 3$, and an agent sees a tie between actions a_1 and a_2 , while he receives the signal s_3 . Our decision rule in (3), which is to break ties by following S_t , only if doing so is optimal, can be viewed as a generalization of similar decision rules in [9], [10] to models with $n \geq 2$, where following S_t when in a tie may not always be optimal. Another choice to break ties is to employ a randomized tie-breaking rule, given by a distribution over the optimal action set $\{a_i\}_{i \in M_t}$, as done in [1].

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

It follows from (3) that, agent t cascades to an action a_i if and only if $\gamma_t^i > \gamma_t^j$ for all $j \neq i$ and for any $S_t \in \{s_1, s_2, \dots, s_n\}$. This is because, when an agent cascades, there cannot be a tie between actions, as this implies that there always exists a different private signal, that if received, would alter the agents' optimal action. A more intuitive way to present this condition is to first express the information contained in the history H_{t-1} observed by agent t in the form of a public likelihood ratio of true state $\omega = i$ to $\omega = j$, for every $i, j \in \Omega$, $i \neq j$ defined as

$$l_{t-1}^{i,j}(H_{t-1}) \triangleq \frac{\mathbb{P}(H_{t-1} | \omega = i)}{\mathbb{P}(H_{t-1} | \omega = j)}. \quad (4)$$

Similarly, we express the information contained in the private signal S_t in the form of agent t 's private likelihood ratio of true state $\omega = i$ to $\omega = j$, for every i, j , defined as

$$\beta_t^{i,j}(S_t) \triangleq \frac{\mathbb{P}(S_t | \omega = i)}{\mathbb{P}(S_t | \omega = j)} = \begin{cases} c, & \text{if } S_t = s_i, \\ 1/c, & \text{if } S_t = s_j, \\ 1, & \text{o.w.} \end{cases} \quad (5)$$

with $c := (n-1)p/(1-p)$, which follows from Figure 1 or equivalently from (2). Next, using Bayes' rule, γ_t^i can be expressed in terms of the public and private likelihood ratios as $\gamma_t^i = 1/(1 + \sum_{j \neq i} l_{t-1}^{j,i} \beta_t^{j,i})$, which can be used to show

$$\gamma_t^i \stackrel{(<)}{>} \gamma_t^j \Leftrightarrow l_{t-1}^{j,i} \beta_t^{j,i} \stackrel{(<)}{>} 1, \quad \text{for any } i \neq j. \quad (6)$$

As a result, the condition on $\{\gamma_t^j\}_{j \in \Omega}$ for an a_i cascade to occur translates to $l_{t-1}^{j,i} > 1/\beta_t^{j,i}$ for all $j \neq i$ and for any S_t . By using the values of $\beta_t^{j,i}$ from (5), the cascade condition simplifies to give the following Lemma.

Lemma 1: Agent t cascades to an action a_i if and only if $l_{t-1}^{i,j} > c$ for all $j \neq i$.

If agent t cascades, then the action A_t does not provide any additional information about the true value ω to the successors over what is contained in H_{t-1} . As a result, $\{l_{t+r}^{i,j}\} = \{l_{t-1}^{i,j}\}$ for all $r = 0, 1, 2, \dots$ and hence they remain in the a_i cascade, which leads us to the following property, also exhibited by prior models, e.g. [1]–[3], [9]–[11].

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

A. Information dynamics until a cascade

Recall that an agent is said to follow his private signal if he takes action a_i only when the signal is s_i . This implies that his action A_t “fully” reveals the private signal S_t to future agents. Assume, without loss of generality, that all agents till some time t follow their private signals. This is a valid assumption, as the first two agents are known to always follow their private signals. Then, due to the mutual independence of the signals $\{S_k\}_{k \leq t}$ given ω , it follows from (4) and the updates in (5) that the public likelihood ratios $\{l_t^{i,j}\}$ can be expressed in terms of the number of s_i 's (denoted by n_t^i), for each $i \in \Omega$, revealed by the observation history H_t as follows.

$$l_t^{i,j} = c^{n_t^i} \left(\frac{1}{c}\right)^{n_t^j} \quad \text{for all } i, j \in \Omega. \quad (7)$$

If agent $t+1$ also follows its private signal which happens to be any s_j , then it trivially follows that

$$n_{t+1}^r = \begin{cases} n_t^r + 1, & \text{if } r = j, \\ n_t^r, & \text{o.w.} \end{cases} \quad (8)$$

Now, given tuple $\{n_t^r\}_{r \in \Omega}$, which denotes the number of private signals of each type revealed till time t , if agent $t+1$ receives a private signal s_i , then it can be shown from (6) and (7) that an optimal action must satisfy the following property.

Property 2: If agent $t+1$ receives a private signal s_i , then action a_i is optimal only if $n_t^i + 1$ upper bounds the tuple $\{n_t^r\}$. Otherwise, any action a_k where n_t^k is maximal in $\{n_t^r\}$ is optimal.

Further, by applying (7) to Lemma 1, the condition on tuple $\{n_t^r\}$ for agent $t+1$ to cascade to an action a_i is as follows.

Property 3: Agent $t+1$ cascades to an action a_i if and only if $n_t^i > 1 + \max\{n_t^r\}_{r \neq i}$. Once any cascade occurs, $\{n_t^r\}$ stops updating.

We now explore the other cases for the tuple $\{n_t^r\}_{r \in \Omega}$ that agent $t+1$ may observe. Consider the following general ordering of the tuple $\{n_t^r\}_{r \in \Omega}$ for some I, J, K that are mutually exclusive and exhaustive in Ω .

$$\begin{aligned} &\{n_t^i\}_{i \in I} > \{n_t^j\}_{j \in J} > \{n_t^k\}_{k \in K} \quad \text{such that} \\ &n_t^{i_1} = n_t^{i_2} \quad \text{for any } i_1, i_2 \in I, \text{ and } 0 < (n_t^{i_1} - n_t^{j_1}) \leq 1 \\ &\text{and } (n_t^i - n_t^k) > 1 \quad \text{for any } i \in I, j \in J, k \in K. \end{aligned}$$

CASE (1): $|I| \geq 2, K \neq \emptyset$.

Here, if agent $t+1$ receives a signal in $\{s_k\}_{k \in K}$, a tie between actions $\{a_i\}_{i \in I}$ occurs, in which case let $\tau(I) = i^*$ denote the index of the tie-winning action. Then, based on the private signal S_{t+1} observed, agent $t+1$ acts as per:

$$A_{t+1} = \begin{cases} a_{i^*}, & \text{if } S_{t+1} \in \{s_{i^*}\} \cup \{s_k\}_{k \in K}, \\ a_j, & \text{if } S_{t+1} \in \{s_j\}_{j \in J \cup I \setminus \{i^*\}}. \end{cases} \quad (9)$$

Note in (9) that the agent follows S_{t+1} only when it belongs to $\{s_j\}_{j \in J \cup I \setminus \{i^*\}}$, in which case $\{n_t^r\}$ updates as per (8). Otherwise, agent takes action a_{i^*} which reveals not just the signal s_{i^*} but also “equally” reveals the signals $\{s_k\}_{k \in K}$. Thus, if action a_{i^*} is taken, the public likelihood ratios would not update as per (5), but instead would update as

$$l_{t+1}^{i^*,j} = \begin{cases} l_t^{i^*,j} \left(\frac{c+|K|}{|K|+1}\right), & \forall j \in J \cup I \setminus \{i^*\}, \\ l_t^{i^*,j}, & \forall j \in K. \end{cases} \quad (10)$$

Observe in (10) that the ratio $l_t^{i^*,j}$, for all $j \in K$, remain unchanged as signals: s_{i^*} and $\{s_j\}_{j \in K}$ are equally revealed by action a_{i^*} . Now, if the relation between $l_t^{i,j}$ and the pair (n_t^i, n_t^j) given in (7) for all i, j has to be ensured for time $t+1$, then the tuple $\{n_t^r\}_{r \in \Omega}$ should be updated as

$$n_{t+1}^r = \begin{cases} n_t^r + \delta(|K|), & \text{if } r \in \{i^*\} \cup K, \\ n_t^r, & \text{if } r \in J \cup I \setminus \{i^*\}, \end{cases} \quad (11)$$

where $\delta(\cdot)$ is a function of $|K|$, which is the number of signals other than s_{i^*} that result in action a_{i^*} , and is given by

$$\delta(|K|) := \log \left(\frac{c+|K|}{|K|+1} \right) / \log c \in \left(\frac{1}{|K|+1}, 1 \right). \quad (12)$$

Note the range of $\delta(|K|)$ in (12) for any $c > 1$ or equivalently for any $p \in (1/n, 1)$. As $\delta(|K|) < 1$, this implies that action a_{i^*} only “partially” reveals the signals: s_{i^*} and $\{s_j\}_{j \in K}$. Only in the special case when $K = \emptyset$, a_{i^*} fully reveals s_{i^*} . This can be verified by (12), where in this case, $\delta(|K|) = 1$.

CASE (2): $|I| = 1, J, K \neq \emptyset$.

In this case, as I is a singleton set, say $I = \{i\}$, the only change with respect to Case (1) is that if agent $t+1$ receives a signal in $\{s_k\}_{k \in K}$, there is no tie and a_i is the sole optimal action. Thus, it trivially follows that the index of the tie-winning action $i^* = i$. Moreover, $I \setminus \{i^*\} = \emptyset$. Substituting these values in (11) yields the updates for the tuple $\{n_t^r\}$ on observing action a_i . Refer to Table I for the public updates of tuple $\{n_t^r\}$ in this case.

For all orderings of the tuple $\{n_t^r\}$ other than in Cases

(1), (2) and the cascade ordering in Property 3, it can be shown that agent $t + 1$ follows any private signal that it receives, and updates as per (8). Action A_{t+1} thereby fully reveals the agent's private signal. Table I summarizes the public updates on $\{n_t^r\}$, for all of their possible orderings. With these updates, equation (7) and Property 2 now hold for any agent t , regardless of whether all agents $k < t$ fully reveal their private signals. The following property thereby follows.

Property 4: The tuple $\{n_t^r\}_{r \in \Omega}$, updated as per Table I, is a sufficient statistic of the information contained in the public history H_t .

Ordering of $\{n_t^r\}$	A_{t+1}	$\{n_t^r\}$ updates
Case (1)	a_{i^*} any $a_j, j \in J \cup I \setminus \{i^*\}$	$n_{t+1}^r = n_t^r + \delta(K), r \in \{i^*\} \cup K$ $n_{t+1}^j = n_t^j + 1$
Case (2)	a_i any $a_j, j \in J$	$n_{t+1}^r = n_t^r + \delta(K), r \in \{i\} \cup K$ $n_{t+1}^j = n_t^j + 1$
$n_t^i > 1 + \max_{j \neq i} \{n_t^j\}$	a_i - cascade	No further updates
Otherwise	any a_j	$n_{t+1}^j = n_t^j + 1$

TABLE I: Public updates on $\{n_t^r\}$ given the observed action A_{t+1} for varying orderings of $\{n_t^r\}$. Only those n_t^r 's that get updated are specified. Only updates of the form $n_{t+1}^j = n_t^j + 1$ imply a fully revealing action. Other update forms, where $\delta(|K|) < 1$, imply only a partial revelation of multiple signals through the action.

Note that for $n = 2$, Cases (1) and (2) cannot occur and so all agents until a cascade fully reveal their private signals.

Remark 2: For $n = 2$, until a cascade occurs, each agent follows its private signal, thus fully revealing it.

The works in [9], [10], [11] study the $n = 2$ model with the agents publicly observing a noisy version of the past actions. These models also satisfy Remark 2. However, due to noise, the observations until a cascade only partially reveal the agents' private signals. Interestingly, this feature of partially revealing observations occurs in our model for $n > 2$, without considering any observation noise.

B. Cascade Probabilities

An a_i -cascade is *correct* if $a_i = a_\omega$ and otherwise is *wrong*. A correct cascade implies that the agents eventually learn the true value ω . Now, given the true value $\omega \in \Omega$, let the probability that an a_i cascade begins be denoted by $^{(n)}\mathbb{P}_{a_i\text{-cas}}^\omega$. Here, the superscript (n) refers to the cardinality of Ω . Then, the probability of a wrong cascade conditioned on ω , denoted by $^{(n)}\mathbb{P}_{\text{wrong-cas}}^\omega$, can be expressed as

$$^{(n)}\mathbb{P}_{\text{wrong-cas}}^\omega := \sum_{i \neq \omega} ^{(n)}\mathbb{P}_{a_i\text{-cas}}^\omega, \quad (13)$$

and the unconditional probability of a wrong cascade is

$$^{(n)}\mathbb{P}_{\text{wrong-cas}} := \frac{1}{n} \sum_{\omega} ^{(n)}\mathbb{P}_{\text{wrong-cas}}^\omega. \quad (14)$$

For $n = 2$, recall from Remark 1 that the tie-breaker $\tau(\cdot)$ is not involved. For this reason, any agent t 's decision rule in (3)

is commutative with respect to the ordering of the posteriors (γ_t^1, γ_t^2) . As a result of this symmetry, we have for $n = 2$:

$$^{(2)}\mathbb{P}_{a_2\text{-cas}}^1 = ^{(2)}\mathbb{P}_{a_1\text{-cas}}^2 = ^{(2)}\mathbb{P}_{\text{wrong-cas}} \quad (15)$$

Whereas, for $n > 2$, as $\tau(\cdot)$ is involved, which is a deterministic tie-breaking rule, (3) is non-commutative with respect to the ordering of the posteriors $(\gamma_t^1, \dots, \gamma_t^n)$. Hence, the conditional wrong cascade probability $^{(n)}\mathbb{P}_{\text{wrong-cas}}^\omega$ given in (13) may or may not be equal for any distinct pair of $\omega_1, \omega_2 \in \Omega$.

Remark 3: For $n > 2$, the conditional wrong cascade probability given by (13) may not necessarily be equal among all ω 's in Ω . Whereas, for $n = 2$, (13) is always equal for any ω .

IV. COMPARISON BETWEEN $n = 2$ AND $n = 3$

In this section, we compare the probability of learning the true value $\omega \in \Omega$ between a model with $n = 2$ and one with $n = 3$, when they have the same private signal quality p . Let M_2 and M_3 denote the respective models. Here, we assume $p \in (1/2, 1)$ to ensure that the private signal is informative in both models. To differentiate between M_2 and M_3 , let $\{s'_1, s'_2\}$ and $\{s_1, s_2, s_3\}$ denote their respective sets of private signals.

Proposition 1: For any private signal quality $p \in (1/2, 1)$ and tie-breaking rule $\tau(\cdot)$, $^{(3)}\mathbb{P}_{\text{wrong-cas}} < ^{(2)}\mathbb{P}_{\text{wrong-cas}}$.

We prove Proposition 1 using a sequence of claims that follow. First, for the sake of discussion, consider the realization $\omega = 1$ under both models. For this ω , we construct a coupling, depicted in Figure 2a, through which signals in $\{s'_1, s'_2\}$ can be generated from the signals in $\{s_1, s_2, s_3\}$ of the M_3 -model.

Claim 1: Given $\omega = 1$ and the coupling in Figure 2a, any a_2 or a_3 (wrong) cascade in the M_3 -model is sufficient for an a_2 (wrong) cascade to occur in the M_2 -model.

Proof: In the M_3 -model, let (a_i, a_j) -sequences refer to all sequences comprising only of actions a_i and a_j , where $i \neq j$. To prove Claim 1, let us enumerate all possible (a_1, a_3) -sequences in M_3 , that would result in an a_3 -cascade. We aim to show that these sequences, under the coupling in Fig. 2a, result in an a_2 cascade in M_2 . Here, we choose the tie-breaking rule $\tau(1, 3) = 1$. Fig. 3 depicts these enumerations and shows the corresponding sequence of actions, that are generated in M_2 as a result of the coupling between the private signals of the two models, defined as per Fig. 2a. The arguments that follow similarly hold for (a_1, a_2) -sequences that result in an a_2 -cascade in M_3 due to the associated symmetry of signals s_2




Fig. 2: Coupling between the private signals in models M_3 and M_2 for different values of ω , such that signals in M_2 are generated through signals in M_3 . For $\omega = 3$, we assume M_2 has state space $\Omega = \{1, 3\}$.




Fig. 3: An enumeration of all possible (a_1, a_3) -sequences that lead to an a_3 -cascade in M_3 and the corresponding sequence of actions, generated in M_2 as a result of the coupling between private signals of the two models, given in Fig. 2a. In M_3 , we assume $\tau(1, 3) = 1$.

and s_3 in Fig. 2a and by noting that both a_3 and a_2 cascades are wrong, given that $\omega = 1$. Lastly, any (a_2, a_3) -sequence in M_3 trivially results in an a_2 -cascade in M_2 at time $t = 3$, hence need not be considered further.

In any sequence in Fig. 3, the function $P(\cdot)$ denotes any permutation of its arguments, which are a set of actions. The notation $(a_i, a_j)^m$ depicts the sub-sequence (a_i, a_j) , successively repeated $m \geq 0$ number of times. In M_3 , actions that only partially reveal multiple private signals are highlighted by indicating these signals above them. As per Table I, such actions only occur under Cases 1 and 2, where in M_3 , the only possible value that $|K|$ can take is 1. Thus, a partially revealed signal s_r results in the update: $n_{t+1}^r = n_t^r + \delta(1)$, where $\delta(1) \in (0.5, 1)$. All other actions in M_3 until a cascade are fully revealing, i.e., any such action a_i fully reveals signal s_i , resulting in the update: $n_{t+1}^i = n_t^i + 1$. In M_2 , as per Remark 2, all actions until a cascade are fully revealing.

Stage (1) in M_3 either starts with a_3, a_3 , resulting directly in an a_3 cascade, or begins with $P(a_1, a_3)$ and then either terminates in an a_3 cascade through a_3, a_3 or continues further with another $P(a_1, a_3)$. At this point, a fully revealing pattern $(a_3, a_1)^m$ is possible, until we observe an a_1 , which begins Stage (2). Here, an a_1 results not only from receiving s_1 but also from s_2 . This is because receiving an s_2 would cause a tie between actions a_1 and a_3 , which a_1 would win as $\tau(1, 3) = 1$. Next, if again an a_1 occurs, it would begin an a_1 cascade which we do not want to enumerate. So, the viable choice is that a_3 occurs. At this point, $n_3 > n_1 > (n_2 + 1)$ with $n_3 - n_1 < 1$. Here, Case (1) applies, thus an a_1 is fully revealing whereas an a_3 is not. So, a fully revealing pattern $(a_1, a_3)^m$ is possible, until we observe an a_3 , which partially

reveals s_3 and s_2 . At this point, $n_3 = n_1 + 1$ and so an a_1 then equalizes n_1 with n_3 such that $n_3 = n_1 > (n_2 + 1)$, which ties back to a point in Stage (1) as shown in Fig. 3. Otherwise, an a_3 starts an a_3 cascade. In this way, all sequences that lead to a a_3 cascade are enumerated.

Applying the possible sequences of private signals, that realise the action sequence in M_3 , to the coupling in Figure 2a generates the corresponding possible sequences of private signals and resulting actions in M_2 , depicted in Fig. 3. Observe that an a_3 cascade in M_3 guarantees an a_2 cascade in M_2 . Similar arguments hold for the alternate choice of tie-breaking rule $\tau(1, 3) = 3$, or if an a_2 cascade occurs in M_3 (by symmetry), thus proving Claim 1. ■

For the realization $\omega = 2$, we consider a different coupling, shown in Figure 2b. Then, by using similar arguments as done for $\omega = 1$, the following claim can be proven.

Claim 2: Given $\omega = 2$ and the coupling in Figure 2b, an a_1 or a_3 cascade in M_3 is sufficient for an a_1 cascade in M_2 .

Lastly, for the realization $\omega = 3$, we assume that model M_2 has a state space $\Omega = \{1, 3\}$. This is required to make $\omega = 3$ a common possibility under both models. It also follows that we should consider the coupling in Figure 2c, which unlike the couplings for other ω 's, maps $\{s_1, s_2, s_3\}$ to $\{s'_1, s'_3\}$. Then, by using similar arguments as done for $\omega = 1$, the following claim can be proven.

Claim 3: For $\omega = 3$ and the coupling in Figure 2c, an a_1 or a_2 (wrong) cascade in M_3 is sufficient for an a_1 (wrong) cascade in M_2 , that has a state space $\Omega = \{1, 3\}$.

Note that when $\omega = 3$, the chances of an a_1 cascade in M_2 with state space $\{1, 3\}$ is equal to the chances of an a_1 cascade in M_2 with state space $\{1, 2\}$ when $\omega = 2$. This is due to the fact that for $n = 2$, there are no state index-dependent changes in agent's decision making. Thus, Claim 3 also implies that $(3)\mathbb{P}_{\text{wrong-cas}}^3 < (2)\mathbb{P}_{\text{wrong-cas}}^2$. This inequality along with Claims 1 and 2 being valid yield the following relation, in which the equality holds due to Remark 3.

$$(3)\mathbb{P}_{\text{wrong-cas}}^\omega < (2)\mathbb{P}_{\text{wrong-cas}}^1 = (2)\mathbb{P}_{\text{wrong-cas}}^2, \quad \forall \omega \in \{1, 2, 3\}. \quad (16)$$

It then follows from the inequality in (16), that the unconditional wrong cascade probabilities for $n = 3$ and $n = 2$, defined by (14), are related as per Proposition 1. Thus, despite M_3 having partially revealing actions, while no such possibility exists in M_2 (Remark 2), we show that learning in M_3 is strictly better than in M_2 .

V. CONCLUSIONS AND FUTURE WORK

We considered the impact of increasing the number of actions in a Bayesian social learning setting and showed that compared to a setting with only two actions, different behavior can emerge. Further, we showed that increasing the number of actions from two to three results in strictly improved learning. This is driven in part because we are keeping the signal quality for the good action fixed as the number of actions increases. One avenue of future study would be to allow this to decrease as the number of actions grows. Extending our analysis to more than three actions is another direction of interest.

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] L. Smith and P. Sørensen, "Pathological outcomes of observational learning," *Econometrica*, vol. 68, no. 2, pp. 371–398, 2000.
- [5] 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.
- [6] N. Kartik, T. Liu, and D. Rappoport, "Beyond unbounded beliefs: How preferences and information interplay in social learning," Tech. Rep., 2022.
- [7] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, "Non-bayesian social learning," *Games and Economic Behavior*, vol. 76, no. 1, pp. 210–225, 2012. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0899825612000851>
- [8] Y. Song, "Social learning with endogenous observation," *Journal of Economic Theory*, vol. 166, pp. 324–333, 2016.
- [9] 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.
- [10] P. Poojary and R. Berry, "Observational learning with fake agents," in *2020 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2020, pp. 1373–1378.
- [11] P. Poojary and R. Berry, "Welfare effects of ex-ante bias and tie-breaking rules on observational learning with fake agents," in *2023 21st International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, 2023, pp. 334–341.
- [12] I. Bistriz, N. Heydaribeni, and A. Anastasopoulos, "Informational cascades with nonmyopic agents," *IEEE Transactions on Automatic Control*, vol. 67, no. 9, pp. 4451–4466, 2022.
- [13] P. Poojary and R. Berry, "Observational learning with negative externalities," in *2022 IEEE International Symposium on Information Theory (ISIT)*, 2022, pp. 1495–1496.
- [14] P. Poojary and R. Berry, "Observational learning in mean-field games with imperfect observations," in *2023 59th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2023, pp. 1–8.