2024 IEEE 63rd Conference on Decision and Control (CDC)
December 16-19, 2024. MiCo, Milan, Italy

Do More Bad Choices Benefit Social Learning?

Pawan Poojary and Randall Berry

2024 IEEE 63rd Conference on Decision and Control (CDC) | USGov | DOI: 10.1109/CDC56724.2024.10886576

***Abstract*— Online markets can enable agents to learn from the actions of others. Such social learning can lead agents to eventually "follow the crowd" and ignore their own private information. This type of behavior has been well studied for agents faced with two possible actions - one "good" action and one "bad" action. In this paper, we consider a scenario where agents have more than two actions and only one of these is good. We show that sequential 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, improves the agents' learning. Whereas, if they are increased from 1 to more than 2, we find that learning can be improved if the private signals are sufficiently strong.**

I. Introduction

In this paper, we consider a sequential Bayesian social learning problem similar to [1]–[3] and others (e.g. [4], [5]). Namely, a sequence of agents must choose one from a set of possible items, where only one item is "good" and the remainder are "bad." Each agent receives a noisy signal indicating the identity of the "good" item and also observes the choices of all preceding agents. Based on this information, agent's seek to make a Bayesian optimal decision. The goal is to characterize the social learning dynamics that occur. Prior works, such as [1]–[3] have studied such problems when the agents are faced with only two alternatives: i.e., one "good" and one "bad." We depart from this prior work by considering agents that are faced with more than two alternatives, where still only one is "good." The objective is to study the information dynamics generated by agents' Bayes' optimal action sequence and to compare and contrast them with the dynamics of a binary model.

An important motivation of this paper is to understand how learning is affected when more "bad" alternatives are introduced compared to a binary model. For fairness of comparison, we assume both the binary and non-binary models maintain the same private signal quality, i.e., the probability that the signal matches the good item is the same. A key feature of the binary models in [1]–[3] is that *informational cascades* or *herding* occur, which are cases where it is optimal for an agent to ignore his own private signal and follow the actions of the past agents. Moreover, these herds may lead to sub-optimal outcomes in which agents cascade to the "bad" item. Intuitively, having more bad alternatives, would lessen the chance of herding towards any one of these as signals indicating these alternatives would occur less frequently and the condition for a cascade to start would now require that evidence for one action dominates the evidence for *all* other actions. We will show that this intuition is not trivial to apply as the learning dynamics with more than two alternatives differ substantially from those in the binary alternatives case.

In related work, [6] also considers more than two alternatives and provides conditions such as directionally unbounded private beliefs, that gaurantee learning. Our model maintains the assumptions of [1]–[3], i.e., discrete bounded private signals, which always leads to a positive probability that learning fails. Another work with multiple alternatives is [7], which considers non-Bayesian learning of the good alternative 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 do not observe all previous agents' actions [5], [8], allowing for imperfect observations [9], [10], and others [11].

The paper is organized as follows. We describe our model in Section II. We analyze this model and identify several properties in Section III. In Sections IV and V, we compare learning between two models, that differ in the number of possible true states. We conclude in Section VI.

II. Model

Assume there is a countable sequence of agents, indexed $t = 1, 2, \dots$ where the index represents both the agent 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 $n \geq 2$ items, which are indexed by $i = 1, 2, \dots, n$. While it is common knowledge that only one among the n items is "good" and all the rest are "bad", the identity (index) 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, i.e.,

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

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

P. Poojary and R. Berry are with the Department of Electrical and Computer Engineering, Northwestern University, Evanston, IL 60208, USA. pawanpoojary2018@u.northwestern.edu, rberry@ece.northwestern.edu

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 $S_t \in \{s_1, s_2, \ldots, s_n\}$. 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 $1 - p$. Hence, given the true value ω,

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

Here, ${}^1\!/_n < p < 1$, which implies that the signal is informative but not revealing. Moreover, the sequence of private signals $\{S_1, S_2, \ldots\}$ 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 $\{A_1, A_2, \ldots, A_{t-1}\}$ that are observed. Note that the models in [1]–[3] are special cases for this model when $n = 2$.

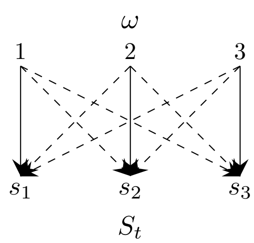

Fig. 1: Transition diagram of 3*-ary SC through which agents receive their private signals. Transitions with solid and dashed arrows occur with probabilities* p *and* $(1-p)/2$*, respectively.*

III. Optimal decision, sufficient statistics and cascades

The t^{th} agent's information set is given by $\{S_t, H_{t-1}\}$, where S_t is its private signal and $H_{t-1} := \{A_1, A_2, \ldots, A_{t-1}\}$ is the history of past actions. For each agent t, the Bayes' optimal action, A_t is chosen such that it provides the greatest expected pay-off given the information set $\{S_t, H_{t-1}\}$. For the first agent, H_0 is the empty set as he does not have any any observation history. Hence, his optimal action is to follow his private signal, i.e., he chooses item i if and only if the signal is s_i. For $t \geq 2$, let $\gamma_t^i(S_t, H_{t-1}) \triangleq \mathbb{P}(\omega = i \,|\, S_t, H_{t-1})$ denote the agent's posterior probability that item i is the good item. Further, let $\gamma_t := (\gamma_t^1, \gamma_t^2, \ldots, \gamma_t^n)$ be the posterior distribution over Ω. Now, it follows from (1) that the pay-off from any action will be the same if it corresponds to the true state ω and if it does not. This implies that a_i is optimum over a_j only if $\gamma_t^i > \gamma_t^j$. This is shown in Figure 2, which depicts regions of optimality for each action within the $(\gamma^i)_{i\in\Omega}$ simplex, for $\Omega = \{1, 2, 3\}$. Thus, a Bayes' optimal decision rule is given by

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

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 determisitic 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 applicable 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 for breaking 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].

A. Cascade conditions

Definition 1: An information cascade is said to occur when an agent's decision becomes a fixed action, regardless 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 $S_t \in \{s_1, s_2, \ldots, s_n\}$. The inequality is strict 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 the cascade 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

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

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

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

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 $\gamma_t^i = 1/(1 + \sum_{j\neq i} l_{t-1}^{j,i} \beta_t^{j,i})$. Using this expression, it can be shown that

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

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

Figure 2 depicts the cascade regions (in darker shades) described by Lemma 1 in the simplex of beliefs $(\gamma^1, \gamma^2, \gamma^3)$ for state space $\Omega = \{1, 2, 3\}$. Here, given any information observed, it can be shown that the belief $(\gamma^i)_{i\in\Omega}$ and the corresponding likelihood ratios for that information, $\{l^{i,j}\}_{i,j\in\Omega}$, are related as $\gamma^i/\gamma^j = l^{i,j}$ for all $i, j \in \Omega$.

Now, 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 $r = 0, 1, 2, \ldots$ 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], [10].

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

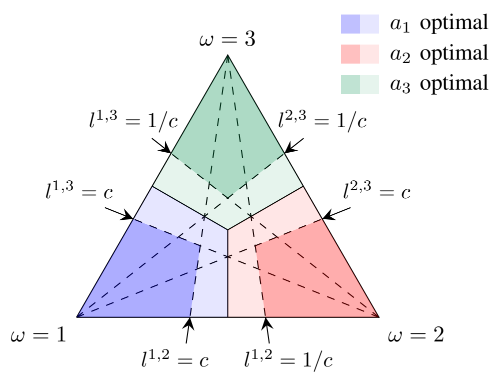

Fig. 2: Simplex of beliefs $(\gamma^i)_{i\in\Omega}$ for state space $\Omega = \{1, 2, 3\}$, which relate to the corresponding likelihood ratios $\{l^{i,j}\}_{i,j\in\Omega}$ as $\gamma^i/\gamma^j = l^{i,j}$. Each dashed line depicts a set of beliefs having a constant likelihood ratio, as indicated. Shaded regions depict optimal actions under uncertainty. Darker shades highlight the cascade regions for the respective actions.

B. 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.

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

where the tuple $\{n_t^r\}_{r\in\Omega}$ denotes the number of private signals of each type revealed till time t. If agent $t+1$ also follows its private signal, which happens to be any s_j, then it trivially follows that

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

and the updated likelihood ratios, $\{l_{t+1}^{i,j}\}$ relate to $\{n_{t+1}^r\}_{r\in\Omega}$ as per (7). Further, given tuple $\{n_t^r\}_{r\in\Omega}$, 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.

Ordering of $\{n_t^r\}$	A_{t+1}	$\{n_t^r\}$ updates
$\lvert I\rvert \geq 2$, $K \neq \emptyset$	a_{i^*}	$n_{t+1}^r = n_t^r + \delta(\lvert K\rvert)$, $r \in \{i^*\} \cup K$
	any a_j, $j \in J \cup I\backslash\{i^*\}$	$n_{t+1}^j = n_t^j + 1$
$\lvert I\rvert = 1$, $J, K \neq \emptyset$	a_i	$n_{t+1}^r = n_t^r + \delta(\lvert K\rvert)$, $r \in \{i\} \cup K$
	any a_j, $j \in J$	$n_{t+1}^j = n_t^j + 1$
$\lvert I\rvert = 1$, $J = \emptyset$	a_i- cascade	No further updates
Otherwise	any a_j, $j \in \Omega$	$n_{t+1}^j = n_t^j + 1$

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. Others, which are increments of $\delta(\lvert K\rvert) < 1$, imply only a partial revelation of multiple signals through the action.

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 $\{n_t^r\}_{r\neq i}$. Otherwise, any action a_k where n_t^k is maximal in $\{n_t^r\}_{r\neq i}$ 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.

So far, the relation in (7) and Properties 2 and 3 rely on the assumption that all agents till time t are fully revealing. However, this may not be the case at all times as we show in the next discussion that for certain orderings of $\{n_t^r\}_{r\in\Omega}$, the action of agent $t+1$ may not fully reveal its private signal. 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.

$$\begin{aligned} &\{n_t^i\}_{i\in I} > \{n_t^j\}_{j\in J} > \{n_t^k\}_{k\in K} \;\; \text{such that} \\ &n_t^{i_1} = n_t^{i_2} \;\text{ for any }\; i_1, i_2 \in I, \text{ and } 0 < (n_t^i - n_t^j) \leq 1 \\ &\text{and } (n_t^i - n_t^k) > 1 \;\text{ for any }\; i \in I,\, j \in J,\, k \in K. \end{aligned} \tag{9}$$

Here, the cardinality of sets I, J, K determines the specific ordering of $\{n_t^r\}$ considered. Table I summarizes the public updates on $\{n_t^r\}$, for all of their possible orderings that agent $t+1$ may observe.

The first ordering of $\{n_t^r\}$ that is shown in Table I is when $\lvert I\rvert \geq 2$, $K \neq \emptyset$ in (9). 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 $i^* := \tau(I)$ denote the index of the tie-winning action. Thus agent $t+1$ follows S_{t+1} only when it belongs to $\{s_j\}_{j\in J\cup I\backslash\{i^*\}}$, 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

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

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

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

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

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

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, $\delta(|K|) = 1$.

The second ordering of $\{n_t^r\}$ that is shown in Table I is when $|I| = 1$, $J, K \neq \emptyset$ in (9). In this case, as I is a singleton set, say $I = \{i\}$, the only change with respect to the first ordering 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-winnning 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, as shown for the second ordering in Table I.

The third ordering of $\{n_t^r\}$, shown in Table I, where $|I| = 1$, $J = \emptyset$ in (9) essentially refers to the ordering in Property 3. Thus, agent $t+1$ onwards, all agents cascade to action a_i and $\{n_t^r\}$ stops updating. For all other orderings of the tuple $\{n_t^r\}$, 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. With the updates in Table I for the respective orderings of $\{n_t^r\}$, equation (7) and Properties 2 and 3 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.

Now, if the ordering of $\{n_t^r\}$ in (9) is such that $|K| \neq \emptyset$, then the following are implied. First, action a_k for any $k \in K$ can never be optimal at $t+1$ since $n_t^k < \max\{n_t^r\}_{r\in\Omega} - 1$. Second, as per Table I, if n_t^k increments by $\delta(|K|)$ at time $t+1$, then so does n_t^i for some $i \in I$, where I is the index set of signals with the maximum count. This ensures that even at time $t+1$, $n_{t+1}^k < \max\{n_{t+1}^r\}_{r\in\Omega} - 1$, which leads to the following remark.

Remark 2: At any given time t, if $n_t^k < \max\{n_t^r\}_{r\in\Omega} - 1$ for some $k \in \Omega$, then action a_k will never be taken by any of the subsequent agents.

Note that for $n = 2$, the first two orderings in Table I cannot occur and so all agents until a cascade fully reveal their private signals.

Remark 3: For $n = 2$, until a cascade occurs, each agent follows its private signal, thus fully revealing it.

The works in [9], [10], study the $n = 2$ model with the agents publicly observing a noisy version of the past actions. These models also satisfy Remark 3. 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.

C. Cascade Probabilities and Welfare

An a_i-cascade is *correct* if $a_i = a_\omega$, that is, if the cascade action is optimal for the realized value of ω. Otherwise, it 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 ${}^{(n)}\mathbb{P}^{\omega}_{a_i\text{-cas}}$. Here, the superscript (n) refers to the cardinality of Ω. Then, the probability of a wrong cascade conditioned on ω, denoted by ${}^{(n)}\mathbb{P}^{\omega}_{\text{wrong-cas}}$, can be expressed as

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

and the unconditional probability of a wrong cascade is

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

Further, let the t^{th} agent's *welfare* refer to its pay-off averaged (in expectation) over $\omega \in \Omega$. It can be shown that this welfare as $t \to \infty$ relates to the wrong cascade probability as

$$\lim_{t\to\infty} \mathbb{E}[\pi_t] = (x - C) - (x + y)\left[{}^{(n)}\mathbb{P}_{\text{wrong-cas}}\right]. \tag{15}$$

Eq. (15) implies that better learning, i.e., a lower probability of wrong cascade, results in a higher asymptotic welfare.

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). Due to this symmetry, we have for $n = 2$:

$${}^{(2)}\mathbb{P}^{1}_{a_2\text{-cas}} = {}^{(2)}\mathbb{P}^{2}_{a_1\text{-cas}} = {}^{(2)}\mathbb{P}_{\text{wrong-cas}} \tag{16}$$

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 $(\gamma_t^1, \ldots, \gamma_t^n)$. Hence, the conditional wrong cascade probability ${}^{(n)}\mathbb{P}^{\omega}_{\text{wrong-cas}}$ given in (13) may or may not be equal for any distinct $\omega_1, \omega_2 \in \Omega$.

Remark 4: For models with $n > 2$, 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 $n = 2$ & $n = 3$

In this section, we compare the probability of learning the true value $\omega \in \Omega = \{1, \ldots, n\}$ between a model with $n = 2$ and one with $n = 3$, both having 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 models M_2 and M_3, let $\{s_1^{'}, s_2^{'}\}$ and $\{s_1, s_2, s_3\}$ denote their respective sets of private signals. Further, recall from Remark 1 that a complete description of model M_3 additionally requires defining the tie-breaking rule $\tau(\cdot)$. In this model, $\tau(\cdot)$ is used to resolve a tie when there are two optimal actions and neither corresponds to following the private signal. We now state the following proposition.

Proposition 1: For any private signal quality $p \in (1/2, 1)$ and tie-breaking rule $\tau(\cdot)$,

$$^{(3)}\mathbb{P}_{\text{wrong-cas}} < {}^{(2)}\mathbb{P}_{\text{wrong-cas}}. \tag{17}$$

We prove (17) 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 Fig. 3a, 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 3a, regardless of the tie-breaking rule $\tau(\cdot)$, 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: We begin by enumerating all possible action sequences in M_3, that lead to an a_2 or a_3 cascade, which are wrong cascades given $\omega = 1$. We aim to show that these sequences, under the coupling in Fig. 3a, result in an a_2 cascade in M_2. The shortest such sequences are a_2, a_2 and a_3, a_3; both trivially lead to an a_2 cascade in M_2, considering that the first two actions in any sequence are always fully revealing. The next possibility is a sequence that starts with two dissimilar actions, say a_i, a_j with $i \neq j$. If the third action is a_k, with $k \neq i, j$, then as each private signal in $\{s_1, s_2, s_3\}$ is revealed once, the public belief on ω is reset to being uniform over Ω. The same argument applies to any number of successive permutations of (a_1, a_2, a_3). We consider sequences in M_3 that begin in this manner separately in Subsection IV-A, and for the sake of this proof consider that the first three actions are not all dissimilar. Thereby, we are left with sequences that begin with a_i, a_j, a_j

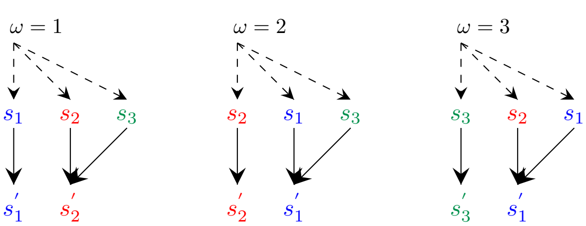

(a) $\omega = 1$ *(b)* $\omega = 2$ *(c)* $\omega = 3$

Fig. 3: 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 $\Omega = \{1, 3\}$.

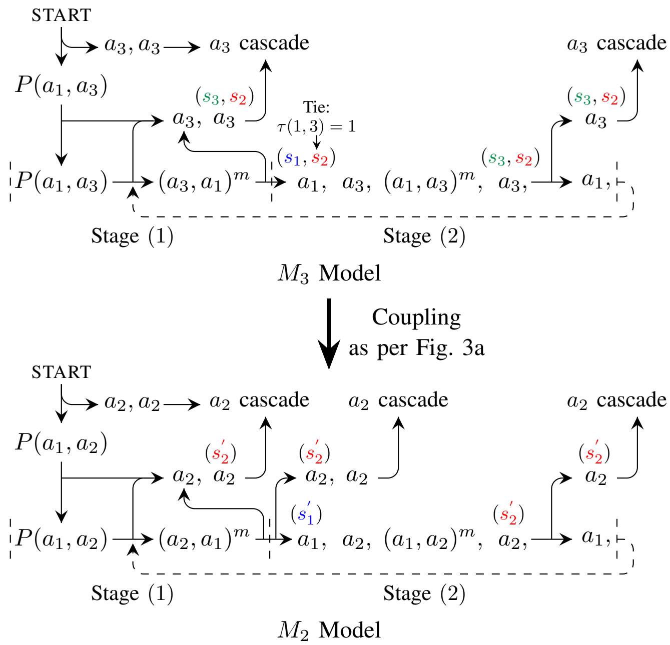

Fig. 4: An enumeration of all possible (a_1, a_3)-sequences that lead to an a_3-cascade in M_3 and the corresponding sequence of actions, generated in M_2 as a result of the coupling between private signals of the two models, given in Fig. 3a. In M_3, we assume $\tau(1, 3) = 1$. Refer to Appendix I for the case $\tau(1, 3) = 3$.

or a_i, a_j, a_i, where again $i \neq j$ and for which the following lemma applies.

Lemma 2: Any action sequence in M_3 that starts with a_i, a_j, a_j or a_i, a_j, a_i will subsequently comprise only of actions a_i and a_j, and will almost surely end in either an a_i or a_j cascade.

The above lemma follows from Remark 2. Now, let (a_i, a_j)-sequences refer to all such sequences, i.e., which comprise only of actions a_i and a_j, where $i \neq j$. These sequences in M_3 are instances where agents sequentially attempt to learn which among the states $\{i, j\}$ is the true state ω. There are three classes of such sequences, namely, (a_1, a_3), (a_1, a_2) and (a_2, a_3)-sequences. To prove Claim 1 for these sequences, let us for the sake of discussion, enumerate all possible (a_1, a_3)-sequences in the M_3-model, that would result in a wrong cascade, i.e., an a_3-cascade. Here, we choose the tie-breaking rule $\tau(1, 3) = 1$ and defer the alternate choice, $\tau(1, 3) = 3$ to Appendix I. Figure 4 depicts these enumerations and shows the corresponding sequence of actions, that are generated in the M_2-model as a result of the coupling between the private signals of the two models, defined as per Fig. 3a. 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 and s_3 in Fig. 3a 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 the sequences in Figure 4, 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. Such actions only occur for the first two orderings in Table I, where for the M_3 model, the only possible value that $|K|$ can take is 1. Thus, a partially revealed signal s_r results in the update: $n^r_{t+1} = n^r_t + \delta(1)$, 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. In M_2, as per Remark 3, all actions until a cascade are fully revealing.

Remark 5: In M_3, if a private signal s_r is "partially" revealed at time $t+1$ (indicated above action), then $n^r_{t+1} = n^r_t + \delta(1)$, where $\delta(1) \in (0.5, 1)$ as per (12). Whereas, in any model, if s_r is "fully" revealed, then $n^r_{t+1} = n^r_t + 1$.

Stage (1) of the sequence in M_3 could start with a_3, a_3, which would directly result in a_3 cascade. Else, it 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, an a_3 would be fully revealing. Hence, 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, i.e., a correct casacde, which we do not intend to enumerate. So, the viable choice is that a_3 occurs. Let $\{n^1, n^2, n^3\}$ refer to the private signal counts of each type revealed till time t, where we abuse notation by dropping t from the subscript as it can be inferred from the context of the discussion. At this point, $n^3 > n^1 > (n^2+1)$ with $n^3 - n^1 < 1$, i.e., the second ordering in Table I 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. 4. Otherwise, an a_3 starts an a_3 cascade. In this way, all sequences that lead to a a_3 cascade are enumerated.

Observe that an a_3 cascade in M_3 gaurantees an a_2 cascade in M_2, thus proving Claim 1. This is depsite our choice of tie-breaking rule $\tau(1,3) = 1$, which by partially revealing signal s_1 when in a tie, favours an a_1 (correct) cascade to occur in M_2. Whereas, the other choice, $\tau(1,3) = 3$, for which Claim 1 is similarly proved by the enumerations in Fig. 6 in Appendix I does not reveal s_1 when in a tie, and hence does not favour a correct cascade in M_2. This makes $\tau(1,3) = 1$ more challenging among the two choices of tie-breaking rules for proving Claim 1. Similar arguments hold for (a_1, a_2)-sequences, that result in an a_2-cascade in M_3 due to the symmetry of signals s_2 and s_3 in Fig. 3a, whereas the arguments for (a_2, a_3)-sequences are trivial. In this manner, all action sequences in M_3 that end in a wrong cascade are accounted, thus proving Claim 1. ■

For the realization $\omega = 2$, we consider a different coupling, shown in Figure 3b. 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 3b, regardless of $\tau(\cdot)$, an a_1 or a_3 (wrong) cascade in M_3 is sufficient for an a_1 (wrong) cascade to occur 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 3c, 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 3c, regardless of $\tau(\cdot)$, 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 ${}^{(3)}\mathbb{P}^3_{\text{wrong-cas}} < {}^{(2)}\mathbb{P}^2_{\text{wrong-cas}}$. This inequality along with Claims 1 and 2 being valid yield the following relation, in which the equality holds due to Remark 4.

$$ {}^{(3)}\mathbb{P}^{\omega}_{\text{wrong-cas}} < {}^{(2)}\mathbb{P}^1_{\text{wrong-cas}} = {}^{(2)}\mathbb{P}^2_{\text{wrong-cas}}, \;\; \forall \omega \in \{1, 2, 3\}. \tag{18} $$

It then follows from the inequality in (18), that the unconditional wrong cascade probabilites for $n = 3$ and $n = 2$, defined by (14), are related as per the Proposition in (17). Thus, despite M_3 having partially revealing actions, while no such possibility exists in M_2 (Remark 3), we show that learning in M_3 is strictly better than in M_2.

A. Proof of Claim 1 for action sequences in M_3 that begin with any number of successive permutations, $P(a_1, a_2, a_3)$.

If any sequence in M_3 begins with $P(a_1, a_2, a_3)$, then since it reveals the private signals $P(s_1, s_2, s_3)$, the only possible permutations in M_3 which do not result in an a_2 cascade in M_2 are $P(a_1, a_2), a_3$ and $P(a_1, a_3), a_2$, both of which yield the sequence $P(a_1, a_2), a_2$ in M_2, through the coupling of Fig. 3a. Note that so far, these sequences (in both models) have fully revealed their private signals. It can be observed that while the number of private signals of each type revealed are equal in M_3 (no information bias), the number of s_2' in M_2 is one greater than the number of s_1', which indicates an information bias in M_2 towards the wrong state, i.e., $\omega = 2$. Due to this information bias in M_2, it can be shown that any (a_1, a_3)-sequence in M_3 illustrated in Figure 4 for $\tau(1,3) = 1$ and in Figure 6 for $\tau(1,3) = 3$, when prefixed with $P(a_1, a_2, a_3)$, results in the corresponding sequence in M_2 to always end in an a_2-cascade. Similar arguments follow for (a_1, a_2)-sequences that end in an a_2-cascade in M_3, when they are prefixed with $P(a_1, a_2, a_3)$, again due to the symmetry of signals s_2 and s_3 in Fig. 3a. Whereas, if any (a_2, a_3)-sequence in M_3 is prefixed with $P(a_1, a_2, a_3)$, an a_2-cascade in M_2 surely occurs by $t = 4$. We have thereby accounted for each (a_i, a_j)-sequence class being prefixed with a $P(a_1, a_2, a_3)$.

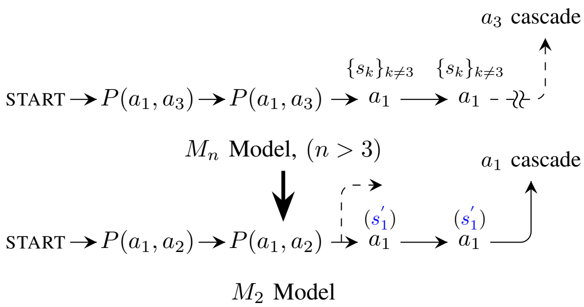

Fig. 5: An example that counters Claim 1 when comparing M_n for $n > 3$, with M_2. Given $\omega = 1$, an (a_1, a_3)-sequence in M_n that leads to an a_3 (wrong) cascade, can result in the corresponding sequence generated in M_2 to end in an a_2 (correct) cascade.

Lastly, if a sequence in M_3 begins with more than one successive $P(a_1, a_2, a_3)$, then the information bias created in M_2 by the first $P(a_1, a_2, a_3)$ gaurantees an a_2 cascade in M_2 by the second $P(a_1, a_2, a_3)$ in M_3. Thus, Claim 1 is proved for all (a_i, a_j)-sequences in M_3 ending in a wrong cascade, that are prefixed with one or more successive $P(a_1, a_2, a_3)$.

V. Comparison between $n = 2$ & $n > 3$

We briefly discuss whether learning in models: M_2 and M_n, for any $n > 3$, can be compared using arguments similar to those in Section IV. For sake of discussion, consider the realization $\omega = 1$ for both models. Assume a coupling similar to Fig. 3a, but for $n > 3$ such that signals $\{s_1^{'}, s_2^{'}\}$ of M_2 can be generated from $\{s_1, \ldots, s_n\}$ of M_n. Fig. 5 shows an (a_1, a_3)-sequence in M_n, assuming the tie-breaking rule $\tau(1, 3) = 1$, and a possible corresponding sequence in M_2. Here, when any signal is partially revealed, we assume that the increment, $\delta(|K|)$, which is monotonic and increasing in signal quality, p, satisfies $\delta(|K|) \leq 0.5$. This holds only if p is less than a threshold, say κ. Now, observe in M_n that after starting with two consecutive $P(a_1, a_3)$'s, there could be an a_1, a_1. But, since $\delta(|K|) \leq 0.5$, this does not result in an a_1 cascade. Let this sequence eventually lead to an a_3 (wrong) cascade. However, the corresponding sequence in M_2 ends in an a_1 (correct) cascade if the the two successive a_1's in M_n were caused by two s_1's. This contradicts Claim 1. This can be avoided if $\delta(|K|) > 0.5$, in which case we expect the arguments in Section IV to hold true. Therefore, learning in M_n, for any $n > 3$ can arguably be better that in M_2 if the private signal quality p is higher than threshold κ.

VI. Conclusions and Future Work

We studied 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, very different learning behavior can emerge. In this setting, until a cascade, an action may partially reveal more than one private signals, while only one among them would be the actual underlying signal. This contrasts the setting with two possible actions, where each action until a cascade always fully reveals its underlying signal. Further, depsite these differences, we showed by applying a coupling method that increasing the number of actions from two to three results in strictly improved learning. Increasing actions from two to more than three improves learning if signals are sufficiently strong. One avenue of future study would be to allow this signal strength to decrease as the number of actions grows. Extending our analysis to other tie-breaking rules 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 & 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.

[8] Y. Song, "Social learning with endogenous network formation," *arXiv preprint arXiv:1504.05222*, 2015.

[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] I. Bistritz, N. Heydaribeni, and A. Anastasopoulos, "Informational cascades with nonmyopic agents," *IEEE Transactions on Automatic Control*, vol. 67, no. 9, pp. 4451–4466, 2022.

Appendix I
Illustration to prove Claim 1 for (a_1, a_3)-sequences in model M_3 and rule, $\tau(1, 3) = 3$.

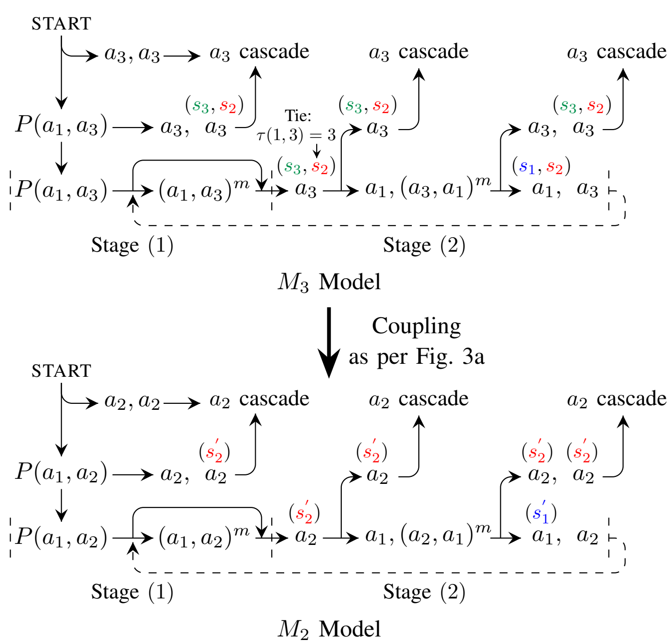

Fig. 6: An enumeration of all possible (a_1, a_3)-sequences that lead to an a_3-cascade in M_3 for $\tau(1, 3) = 3$, and the corresponding sequence of actions, generated in M_2.